EXPONENTIAL COLLAPSE OF SOCIAL BELIEFS OVER WEAKLY-CONNECTED HETEROGENEOUS NETWORKS

Vincenzo Matta

DIEM, University of Salerno
via Giovanni Paolo II, I-84084, Fisciano, Italy
e-mail: vmatta@unisa.it

Augusto Santos, Ali H. Sayed *

EPFL, School of Engineering
CH-1015 Lausanne, Switzerland
e-mails: {augusto.santos, ali.sayed}@epfl.ch

ABSTRACT

We consider a distributed social learning problem where a network of agents is interested in selecting one among a finite number of hypotheses. The data collected by the agents might be heterogeneous, meaning that different sub-networks might observe data generated by different hypotheses. For example, some sub-networks might be receiving (or even intentionally generating) data from a fake hypothesis and will bias the rest of the network via social influence. This work focuses on a two-step diffusion algorithm where each agent: i) first updates individually its *belief function* using its *private* data; ii) then computes a new belief function by *exponentiating* a linear combination of the *log*-beliefs of its neighbors. We obtain analytical formulas that reveal how the agents' detection capability and the network topology interplay to influence the asymptotic beliefs of the agents. Some interesting behaviors arise, such as the "*mind-control*" effect or the "*truth-is-somewhere-in-between*" effect.

Index Terms— Social learning, weakly-connected networks, Bayesian update, diffusion strategy.

1. INTRODUCTION

In the network era, understanding the fundamental laws that govern the mechanism of social learning is a challenge of paramount importance. The social learning problem can be broadly described as one involving interacting agents constructing their beliefs about a phenomenon of interest (learning process) by updating continuously their individual knowledge through the exchange of information (social interaction) with neighboring agents [1–6].

Two relevant models for social learning are the Bayesian model [1, 7], where the agents implement Bayes' rule to update their beliefs by relying on some prior knowledge [1,8,9], and the non-Bayesian model, where agents interact with their neighbors and aggregate their beliefs into their own [10–13].

This work focuses on non-Bayesian learning. Several useful distributed implementations have been proposed for non-Bayesian learning including, e.g., consensus implementations [13] and diffusion-type implementations, either with a linear combination of beliefs [14] or with the linear combination of log-beliefs [4, 15]. The vast majority of works focus on strongly-connected networks. For such networks, it is known that all agents are able to reach a common belief about the most likely state of nature. In contrast, the works [16, 17] considered the case of *weakly-connected* graphs, which arise over many popular social networks. Over these graphs, the topology consists of multiple sub-networks where one first category of sub-networks (called *sending* sub-networks) feeds information in one direction to other network components (called *receiving* sub-networks) without receiving back (or being interested in) any information from them. This behavior is not uncommon over social networks. For example, a celebrity user may have a large number of followers but may not be interested in following the beliefs of most of them. It was shown in [17] that, over weakly-connected graphs, a sending sub-network plays a domineering role by influencing significantly the beliefs of the receiving agents (*mind-control* effect), and *irrespective of the local observations sensed by the latter*. In particular, receiving agents can be led to incorrect decisions, and they can also be made to disagree on their inferences among themselves (*discord* effect).

The main contribution of the present work consists of extending the weak-graph analysis conducted in [17] to the class of log-belief combination algorithms (and, as a result, we end up extending the analysis for log-belief combination algorithms conducted in [4, 15] to the case of weak graphs). We characterize the limiting beliefs at each agent through analytical formulas that depend in a transparent manner on inferential descriptors (Kullback-Leibler divergences), and network descriptors (limiting combination matrix). Some revealing behaviors emerge from these formulas, which are examined in Sec. 3 further ahead.

2. MAIN RESULT

We assume that the agents gather streaming observations, $\{\boldsymbol{\xi}_{k,i}\}$, where k is the agent index ($k = 1, 2, \ldots, N$) and i

*This work was supported in part by NSF grant CCF-1524250.

is the time index ($i = 1, 2, \ldots$). The data are assumed to be independent across time, whereas they can be dependent across agents. Moreover, we work under the *heterogeneous* setting where the data $\boldsymbol{\xi}_{k,i}$ observed by agent k at time i are generated according to some distribution $f_k(\cdot)$. Following the model proposed in [13,14], we assume that these distributions are unknown. The main goal of the network agents is learning about a certain state of nature within a finite collection of states Θ. To this end, each agent k possesses a family of likelihood functions, $\{L_k(\cdot|\theta)\}_{\theta\in\Theta}$, which, we remark, need not to coincide with the true distributions $f_k(\cdot)$. In a nutshell, the choice of a particular state of nature will be determined by the nominal likelihood function that gives the best match with the observations. In order to rule out pathological cases, we assume that all distributions and all likelihood functions of a given agent share a common support. However, the supports may vary across different agents.

The learning algorithm considered in this work is:

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

$$\boldsymbol{\mu}_{k,i}(\theta) = \frac{\exp\left\{\sum_{\ell\in\mathcal{N}_k} a_{\ell k}\ln\boldsymbol{\psi}_{\ell,i}(\theta)\right\}}{\sum_{\theta'\in\Theta}\exp\left\{\sum_{\ell=1}^{N} a_{\ell k}\ln\boldsymbol{\psi}_{\ell,i}(\theta')\right\}}. \quad (2)$$

Each computation in the above algorithm is repeated for each admissible hypothesis, $\theta \in \Theta$. We see that agent k first updates its local belief, $\boldsymbol{\mu}_{k,i-1}(\theta)$, by incorporating the local likelihood $L_k(\boldsymbol{\xi}_{k,i}|\theta)$, which is based on the fresh *private* observation, $\boldsymbol{\xi}_{k,i}$, available to agent k at time i. This step leads to an intermediate belief $\boldsymbol{\psi}_{k,i}(\theta)$.

Then, agent k aggregates the intermediate beliefs received from its neighbors. Differently from [17], the second step in (2) combines linearly the *logarithm* of the intermediate beliefs, $\ln\boldsymbol{\psi}_{\ell,i}(\theta)$, in the neighborhood of agent k, i.e., for all $\ell \in \mathcal{N}_k$. Finally, exponentiation and normalization are required to give back an admissible belief.

Before stating our main results, it is necessary to introduce some useful quantities. First we introduce the combination matrix, A, which collects the combination weights $[a_{\ell k}]$, and describes the *connection features* of our problem. We assume a limiting matrix, $A^\bullet \triangleq \lim_{n\to\infty} A^n$, exists.

Next we introduce a quantity that describes the *inference features* of our problem, namely, the Kullback-Leibler (KL) divergence between the true distribution $f_k(\cdot)$ and the likelihood function $L_k(\cdot|\theta)$, for all $\theta \in \Theta$:

$$D_k(\theta) \triangleq \mathbb{E}_k\left[\ln\frac{f_k(\boldsymbol{\xi}_k)}{L_k(\boldsymbol{\xi}_k|\theta)}\right], \quad k = 1, 2, \ldots, N, \quad (3)$$

where the expectation is computed under the statistical model $f_k(\cdot)$, and we write $\boldsymbol{\xi}_k$ in place of $\boldsymbol{\xi}_{k,i}$ to highlight identical distribution across time.

Finally, we introduce a quantity that combines the network centrality and the inference capability:

$$\boxed{\mathscr{D}_k(\theta) \triangleq \sum_{\ell=1}^{N} a^\bullet_{\ell k} D_\ell(\theta)} \quad (4)$$

We remark that $\mathscr{D}_k(\theta)$ depends on the particular agent k, and represents the convex combination, through the limiting weights $[a^\bullet_{\ell k}]$, of the KL divergences between the true distributions over the various agents, $f_\ell(\cdot)$, and the tentative likelihood functions, $L_\ell(\cdot|\theta)$.

Motivated by the results from [14, 17], the following lemma, whose proof is omitted for space constraints, and the subsequent theorem, generalize the convergence results from [15] to the case of heterogeneous data and weak graphs, and the convergence results from [4] to the case of weak graphs. In what follows, we always assume that, at time $i = 0$, the agents have no strong confidence to reject any hypothesis, and, hence, they assign $\mu_{k,0}(\theta) > 0$ for all $\theta \in \Theta$ and $k = 1, 2, \ldots, N$.

Lemma 1 *Assume that, for all $\theta \in \Theta$, and $k = 1, 2, \ldots, N$:*

$$\mathbb{E}_k\left|\ln\frac{f_k(\boldsymbol{\xi}_k)}{L_k(\boldsymbol{\xi}_k|\theta)}\right| < \infty. \quad (5)$$

Then, for all $\theta, \theta' \in \Theta$, we have:

$$\boxed{\lim_{i\to\infty}\frac{1}{i}\ln\frac{\boldsymbol{\mu}_{k,i}(\theta)}{\boldsymbol{\mu}_{k,i}(\theta')} \overset{\text{a.s.}}{=} \mathscr{D}_k(\theta') - \mathscr{D}_k(\theta)} \quad (6)$$

where $\overset{\text{a.s.}}{=}$ denotes almost-sure convergence. □

The result of Lemma 1 is useful to determine the asymptotic behavior of the beliefs at the individual agents. In order to understand why, let us consider the case that $\frac{1}{i}\ln\frac{\boldsymbol{\mu}_{k,i}(\theta)}{\boldsymbol{\mu}_{k,i}(\theta')}$ converges to a *strictly positive* quantity. This implies that $\boldsymbol{\mu}_{k,i}(\theta')$ converges to zero (since the numerator is bounded by 1). This property is exploited to prove the next theorem.

Theorem 1 (Belief collapse) *Assume that, for each $k = 1, 2, \ldots, N$, the function $\mathscr{D}_k(\theta)$ admits a unique minimizer:*

$$\vartheta^\bullet_k = \operatorname*{argmin}_\theta \mathscr{D}_k(\theta). \quad (7)$$

Then:

$$\boxed{\lim_{i\to\infty}\boldsymbol{\mu}_{k,i}(\vartheta^\bullet_k) \overset{\text{a.s.}}{=} 1, \qquad \text{for all } k = 1, 2, \ldots, N} \quad (8)$$

Proof. Since by assumption $\vartheta^\bullet_k$ is the unique minimizer of $\mathscr{D}_k(\theta)$, from (6) we can write, for all $\theta \neq \vartheta^\bullet_k$:

$$\lim_{i\to\infty}\frac{1}{i}\ln\frac{\boldsymbol{\mu}_{k,i}(\vartheta^\bullet_k)}{\boldsymbol{\mu}_{k,i}(\theta)} \overset{\text{a.s.}}{=} \mathscr{D}_k(\theta) - \mathscr{D}_k(\vartheta^\bullet_k) > 0. \quad (9)$$

In light of the observation preceding the claim of this theorem, Eq. (9) implies that $\boldsymbol{\mu}_{k,i}(\theta) \to 0$ for all $\theta \neq \vartheta^\bullet_k$. On the other hand, we know that it is not possible that the belief function vanishes at *all* θ's because $\boldsymbol{\mu}_{k,i}(\theta)$ is a probability measure over Θ. As a result, we conclude that, at each agent k, the belief function converges necessarily to 1 at $\vartheta_k^\bullet$. $\square$

We remark that, since the quantity $\vartheta_k^\bullet$ depends upon the agent index k, different agents can in principle be in *discord*, since they can coalesce towards different hypotheses.

In this conference article we examine some interesting nontrivial behaviors that emerge over weakly-connected networks when the assumption of a unique minimizer for $\mathscr{D}_k(\theta)$ holds. This assumption is not the most general case that can be addressed. However, in view of (3), coincidence of the average divergences for two distinct values of θ requires some ad-hoc combination of: i) conditions determined by the network topology (through the limiting matrix weights, $a_{\ell k}^\bullet$) and ii) conditions determined by the difficulty of the detection problem (through the KL divergences). As a result, the occurrence of a non-unique minimizer seems to be the exception rather than the rule.

3. WEAKLY CONNECTED NETWORKS

We divide the network into a *sending* component, $\mathcal{S} \triangleq \bigcup_{s=1}^{S} \mathscr{N}_s$ (i.e., internally composed of S disjoint components $\mathscr{N}_s$), and a receiving component, $\mathcal{R} \triangleq \bigcup_{r=1}^{R} \mathscr{N}_{S+r}$ (i.e., internally composed of R components $\mathscr{N}_{S+r}$). We assume that each of the sending sub-networks is *strongly* connected, whereas each of the receiving subnetworks is assumed to be *connected* (not necessarily *strongly* connected). Communication from the sending component to the receiving component is permitted. In particular, we assume that each receiving sub-network is connected to at least one sending agent. In contrast, communication in the reverse direction is forbidden and communication from one sending sub-network to another is forbidden. Finally, communication among the R receiving sub-networks is permitted.

According to the above description, the combination matrix A corresponding to a weakly-connected network admits the following block representation:

$$\left[\begin{array}{cccc|cccc} A_1 & 0 & \cdots & 0 & A_{1,S+1} & A_{1,S+2} & \cdots & A_{1,S+R} \\ 0 & A_2 & \cdots & 0 & A_{2,S+1} & A_{2,S+2} & \cdots & A_{2,S+R} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_S & A_{S,S+1} & A_{S,S+2} & \cdots & A_{S,S+R} \\ \hline 0 & 0 & \cdots & 0 & A_{S+1} & A_{S+1,S+2} & \cdots & A_{S+1,S+R} \\ 0 & 0 & \cdots & 0 & A_{S+2,S+1} & A_{S+2} & \cdots & A_{S+2,S+R} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & A_{S+R,S+1} & A_{S+R,S+2} & \cdots & A_{S+R} \end{array}\right] \tag{10}$$

Figure 1 offers a graphical illustration of the weakly-connected paradigm. It has been shown in [16] that the limiting combination matrix has the following structure:

$$A^\bullet = \left[\begin{array}{c|c} E & EW \\ \hline 0 & 0 \end{array}\right], \tag{11}$$

where $E = \mathsf{blockdiag}\left\{p^{(1)}\mathbb{1}_{N_1}^\top, p^{(2)}\mathbb{1}_{N_2}^\top, \ldots, p^{(S)}\mathbb{1}_{N_S}^\top\right\}$ is

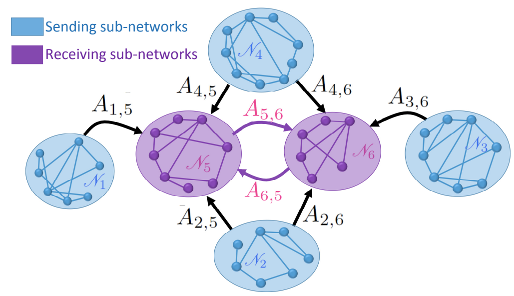

Fig. 1: One example of a weakly connected network.

a block diagonal matrix that stacks the Perron eigenvectors $p^{(s)}$ associated with each sending sub-network A_s, and $W = T_{SR}(I - T_{RR})^{-1}$, with T_{SR} and T_{RR} being the upper-right block and the lower-right block in (10), respectively. Let us now focus on a generic agent belonging to a *receiving* network. Using (11), the result of Theorem 1 specializes, for all $k \in \mathcal{N}_s$, with $s = 1, 2, \ldots, S$:

$$\vartheta_k^\bullet = \underset{\theta}{\operatorname{argmin}} \sum_{\ell \in \mathscr{N}_s} p_\ell^{(s)}\, D_\ell(\theta), \tag{12}$$

and for all $k \in \mathcal{R}$, into:

$$\vartheta_k^\bullet = \underset{\theta}{\operatorname{argmin}} \sum_{s=1}^{S} \sum_{\ell \in \mathscr{N}_s} [EW]_{\ell k}\, D_\ell(\theta). \tag{13}$$

Note that, given an agent k in the receiving sub-network $\mathscr{N}_{S+r}$, then $[EW]_{\ell k}$ is strictly positive for all agents ℓ in the sending sub-networks $\mathscr{N}_s$ that are connected to $\mathscr{N}_{S+r}$.

Equation (13) reveals a remarkable analogy with what was proved in [17]: the limiting belief, at the *receiving* agents, is determined only by attributes of the *sending* network. However, there are also some nontrivial novel behaviors emerging from our results. One relevant aspect pertains to the "mind control" problem. Let us initially refer to the simplest case of one sending network and one receiving network. For example, the sending network might represent some big media network, whereas the receiving network might represent the audience of such media. Assume that the sending network is interested in controlling the audience by leading them to certain (maybe fake) conclusions. To this end, all the agents of the sending network generate fake data according to a distribution $f_\ell(\cdot) = L_\ell(\cdot|\theta_{\text{fake}})$, which, in view of (3), means that $D_\ell(\theta_{\text{fake}}) = 0$ for all $\ell = 1, 2, \ldots, S$. Substituting this result into (13), we see that $\vartheta_k^\bullet = \theta_{\text{fake}}$ for all k in the receiving network, and we conclude that the sending network is successful in controlling the "mind" of the receiving agents, irrespective of the receiving agents' attributes (in centrality and detectability capacities).

However, the mind control effect can become more intricate in the general case when there are more sending sub-networks. This can be illustrated even with the simplest

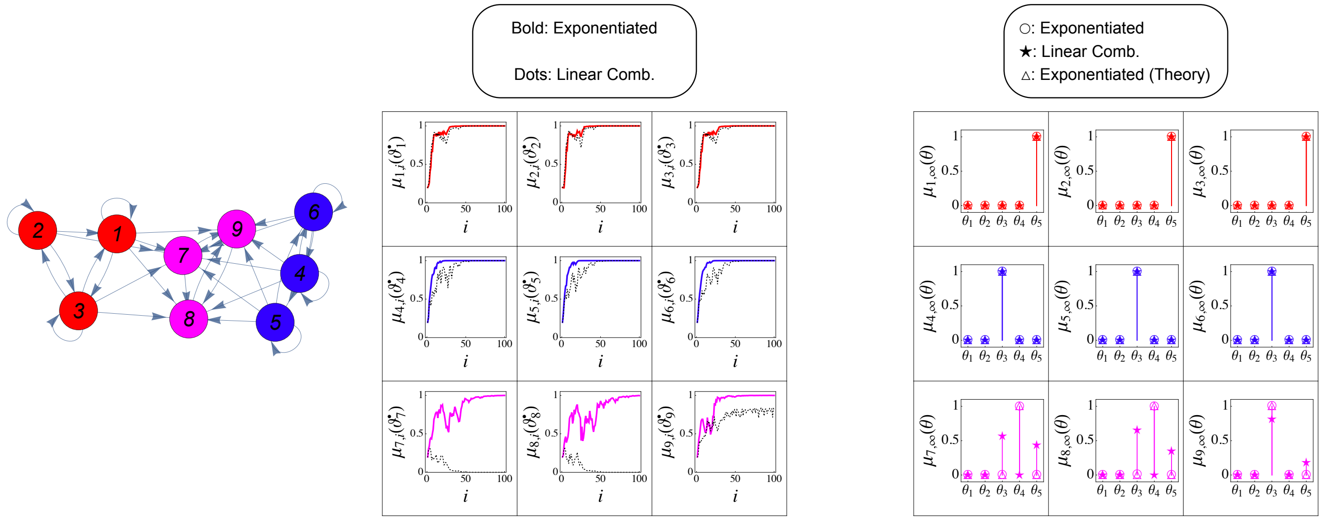

(a) Network topology. (b) Beliefs at distinct agents (time-evolution). (c) Limiting beliefs at distinct agents.

Fig. 2: Example over a weakly-connected network.

example of two sending agents, one receiving agent, and three hypotheses, $\theta_1, \theta_2, \theta_3$. For the sake of concreteness, assume that the data of sending agent 1 are drawn according to $L_1(\cdot|\theta_1)$, whereas the data of sending agent 2 are drawn according to $L_2(\cdot|\theta_2)$, which reflects the situation where the two agents have conflicting requirements (i.e., data from distinct hypotheses). Finally, assume that both sending agents are connected to the receiving agent. Using this model to compute the average divergence perceived by the receiving agent 3, we get: $\mathscr{D}_3(\theta_1) = 0.5\, D_2(\theta_1)$, $\mathscr{D}_3(\theta_2) = 0.5\, D_1(\theta_2)$, and $\mathscr{D}_3(\theta_3) = 0.5\, [D_1(\theta_3) + D_2(\theta_3)]$. Under this situation, if $\mathscr{D}_3(\theta_1)$ is the minimum, then network 1 gains the mind control. Conversely, if $\mathscr{D}_3(\theta_2)$ is the minimum, then network 2 gains the mind control. However, there is a third possibility: if $\mathscr{D}_3(\theta_3)$ is the minimum, then the receiving agent chooses an option that is not promoted by either of the sending sub-networks. It is useful to explore how this particular (and perhaps unexpected) behavior is possible. To this aim, we resort to the following example. Assume that the receiving agent wants to bet on a soccer match between team A and team B. Assume also that a group of (sending) agents promotes victory of A, whereas another group promotes B. The receiving agent decides to manage these conflicting observations according to the philosophy that "*the truth is somewhere in between*", namely, he bets on draw. In other words, if the discrepancy between the two hypotheses promoted by the sending networks is too high, the receiving agent can be driven to opt for the hypothesis that best reconciles the discordant solicitations received from the environment. Such enhanced capability of *fighting against mind control* is in contrast with what is shown in [17] for the case of linear combination of beliefs, where no credit is given to a draw, since the belief function of the receiving agent assigns some credit to the victory of A, and some (complementary) credit to the victory of B. In a sense, this way of handling the conflict amounts to a form of *soft* decision.

4. ILLUSTRATIVE EXAMPLE

We consider the weakly-connected topology displayed in the leftmost panel of Fig. 2, where we have two sending subnetworks (red and blue), and one receiving sub-network (magenta). To keep things simple, the distributions of the data are chosen within the same class of nominal likelihood functions used by the agents. In particular, these likelihood functions correspond to Gaussian distributions with possible expectations $\theta_1 < \theta_2 < \ldots < \theta_5$. The true hypotheses are: θ_5 for the red sending sub-network, θ_3 for the blue sending subnetwork, and θ_2 for the receiving sub-network.

Panel (b) illustrates the time-evolution of the beliefs of each agent, for the exponentiated log-belief combination algorithm, and for the linear-combination algorithm from [17]. In panel (c), we display the corresponding limiting beliefs. We see that the simulated beliefs for the log-belief algorithm match the theoretical predictions. In particular, we observe that the sending agents converge to the true parameters that generate their own data. For what concerns the receiving agents, two important features emerge. First, they all collapse to a particular hypothesis. Second, a peculiar form of discord is observed, where agent 9 converges to θ_3 (the true parameter of the blue sending sub-network), while agents 7 and 8 converge to θ_4, which is *neither the true parameter of the receiving network, nor a parameter of the sending network*.

5. REFERENCES

[1] C. Chamley, *Rational Herds: Economic Models of Social Learning*, Cambridge Univ. Press, Cambridge, UK, 2004.

[2] A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi, "Information heterogeneity and the speed of learning in social networks," *Columbia Business School Research Paper*, pp. 13–28, May 2013.

[3] A. Nedić and A. Olshevsky, "Distributed optimization over time-varying directed graphs," *IEEE Trans. Autom. Control*, vol. 60, no. 3, pp. 601–615, March 2015.

[4] 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, Nov 2017.

[5] C. Chamley, A. Scaglione, and L. Li, "Models for the diffusion of beliefs in social networks: An overview," *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 16–29, May 2013.

[6] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione, "Binary opinion dynamics with stubborn agents," *ACM Trans. Econ. Comput.*, vol. 1, no. 4, pp. 19:1–19:30, December 2013.

[7] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dyn. Games Appl.*, vol. 1, no. 1, pp. 3–49, 2011.

[8] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *Rev. Econ. Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

[9] 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.

[10] M. H. DeGroot, "Reaching a consensus," *J. Amer. Statist. Assoc.*, vol. 69, no. 345, pp. 118–121, 1974.

[11] L. G. Epstein, J. Noor, and A. Sandroni, "Non-Bayesian learning," *BE J. Theor. Econ.*, vol. 10, no. 1, pp. 1–20, 2010.

[12] D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi, "Spread of (mis)information in social networks," *Games and Economic Behavior*, vol. 70, no. 2, pp. 194–227, 2010.

[13] 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.

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

[15] A. Lalitha, T. Javidi, and A. D. Sarwate, "Social learning and distributed hypothesis testing," *IEEE Trans. Inf. Theory*, vol. 64, pp. 6161–6179, September 2018.

[16] B. Ying and A. H. Sayed, "Information exchange and learning dynamics over weakly connected adaptive networks," *IEEE Trans. Inf. Theory*, vol. 62, no. 3, pp. 1396–1414, March 2016.

[17] H. Salami, B. Ying, and A. H. Sayed, "Social learning over weakly connected graphs," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 3, no. 2, pp. 222–238, June 2017.