Online Graph Learning from Social Interactions

Valentina Shumovskaia*, Konstantinos Ntemos*, Stefan Vlaski[†] and Ali H. Sayed*

* École Polytechnique Fédérale de Lausanne (EPFL)

[†] Imperial College London

Abstract—Social learning algorithms provide models for the formation of opinions over social networks resulting from local reasoning and peer-to-peer exchanges. Interactions occur over an underlying graph topology, which describes the flow of information and relative influence between pairs of agents. For a given graph topology, these algorithms allow for the prediction of formed opinions. In this work, we study the inverse problem. Given a social learning model and observations of the evolution of beliefs over time, we aim at identifying the underlying graph topology. The learned graph allows for the inference of pairwise influence between agents, the overall influence agents have over the behavior of the network, as well as the flow of information through the social network. The proposed algorithm is online in nature and can adapt dynamically to changes in the graph topology or the true hypothesis.

Index Terms—Graph learning, inverse modeling, online learning, social learning.

I. INTRODUCTION

Graphs provide a useful tool to model and exploit relations in high-dimensional data, such as social networks [1], [2], roadway networks [3], [4], and molecular data [5], [6], among others. A common observation in all these settings is the fact that the graph topology impacts data distribution and evolution. Hence, knowledge of the graph has the potential to improve the performance of inference tasks. However, the underlying graph structure is unknown in many applications and needs to be estimated through observations. A number of solutions for graph learning [7]–[18] have already been proposed in the literature, where algorithms for graph inference have been developed for particular models, describing the relationship between observations and graphs.

For instance, graph learning for the heat diffusion process is studied in [9], [11], [12], [19], while learning under structural constraints, such as connectivity [8] and sparsity, appears in [8], [13], [18], and approaches based on examining the precision matrix appear in [20], [21]. Most of these works consider static graphical models. This is in contrast to graphs with dynamic properties [9] where the connectivity among agents can change over time.

In this work, we develop an algorithm for graph learning in the social learning setting where agents react to streaming data and also to information shared with their neighbors. Our study focuses on the social learning paradigm studied in earlier works [22]–[31]. The combination weights are unknown and need to be inferred through observations acquired during the agents' learning process.

Social learning refers to the problem of distributed hypothesis testing, where each agent aims at learning an underlying true hypothesis (or state) through its own observations and from information shared by its neighbors. Social learning studies can be categorized into Bayesian [30], [31] and non-Bayesian [22]–[29]. Non-Bayesian approaches have gained increased interest due to their appealing scalability traits. In these approaches, at every time instant, agents follow a twostage process. First, every agent updates its belief (which is a probability distribution over the possible hypotheses) based on its current received observation. Then, it fuses the shared beliefs from its neighbors. The main focus of these studies is to prove that agents' beliefs across the network converge to the true hypothesis after sufficient repeated interactions.

In this work, we are interested in revealing the underlying influence pattern. More specifically, we formulate the inverse problem, where given the evolution of beliefs, the objective is to identify the relative influence between pairs of agents, captured by the graph topology. We are interested in studying a dynamic setting where both the graph topology and the true hypothesis can change over time. Therefore, we consider that the agents follow the adaptive social learning protocol [26].

We describe the system model in Section II, while Section III describes the algorithm and examines the steady-state performance of the graph learning process. In Section IV, we provide experiments and illustrate the robustness of the proposed method against dynamic changes on the graph topology.

II. SOCIAL LEARNING MODEL

We consider a set \mathcal{N} of agents connected by a graph $\mathcal{G} = \langle \mathcal{N}, \mathcal{E} \rangle$, where \mathcal{E} represents the links between agents. Two agents that are linked can exchange information directly with each other. The set of neighbors of an agent $k \in \mathcal{N}$ including itself, is denoted by \mathcal{N}_k .

All agents aim at learning the true hypothesis θ^* , belonging to a set of all possible hypotheses denoted by Θ (whose cardinality is at least two). To this end, each agent k has access to observations $\zeta_{k,i} \in \mathbb{Z}_k$ at every time $i \ge 1$. Agent k also has access to the likelihood functions $L_k(\zeta_{k,i}|\theta)$, for all $\theta \in \Theta$. The signals $\zeta_{k,i}$ are independent over both time and space, and are also identically distributed (i.i.d.) over time. We will use the notation $L_k(\theta)$ instead of $L_k(\zeta_{k,i}|\theta)$ for brevity. At each time *i*, agent *k* keeps a *belief vector* $\boldsymbol{\mu}_{k,i}$, which is a probability distribution over the possible states. The belief

Emails: {valentina.shumovskaia, konstantinos.ntemos, ali.sayed}@epfl.ch, s.vlaski@imperial.ac.uk. This work was supported in part by SNSF grant 205121-184999.

component $\mu_{k,i}(\theta)$ quantifies the confidence of agent k that θ is the true state. Therefore, at time i, each agent's true state estimator is as follows:

$$\widehat{\boldsymbol{\theta}}_{k,i}^{\circ} = \arg \max_{\boldsymbol{\theta} \in \Theta} \boldsymbol{\mu}_{k,i}(\boldsymbol{\theta}).$$
(1)

To avoid technicalities, where agents discard a particular state a priori, we impose the following assumption on initial beliefs.

Assumption 1. (Positive initial beliefs). For all hypotheses $\theta \in \Theta$, all agents $k \in \mathcal{N}$ start with positive initial belief $\mu_{k,0}(\theta) > 0$.

At every time instant *i*, every agent *k* updates its belief by using a two-stage process. First, it incorporates information from the received observation $\zeta_{k,i}$ and then it fuses the information from its neighbors. More specificially, in this work we consider the *adaptive social learning* rule [26], which has been shown to have favorable transient and steady-state performance in terms of convergence rate and probability of error. Under this protocol, agents update their beliefs in the following manner:

$$\boldsymbol{\psi}_{k,i}(\boldsymbol{\theta}) = \frac{L_k^{\delta}(\boldsymbol{\zeta}_{k,i}|\boldsymbol{\theta})\boldsymbol{\mu}_{k,i-1}^{1-\delta}(\boldsymbol{\theta})}{\sum_{\boldsymbol{\theta}'\in\boldsymbol{\Theta}} L_k^{\delta}(\boldsymbol{\zeta}_{k,i}|\boldsymbol{\theta}')\boldsymbol{\mu}_{k,i-1}^{1-\delta}(\boldsymbol{\theta}')}, \quad k \in \mathcal{N}$$
(2)

$$\boldsymbol{\mu}_{k,i}(\boldsymbol{\theta}) = \frac{\prod_{\ell \in \mathcal{N}_k} \boldsymbol{\psi}_{\ell,i}^{a_{\ell k}}(\boldsymbol{\theta})}{\sum_{\boldsymbol{\theta}' \in \Theta} \prod_{\ell \in \mathcal{N}_k} \boldsymbol{\psi}_{\ell,i}^{a_{\ell k}}(\boldsymbol{\theta}')}, \quad k \in \mathcal{N}$$
(3)

where $a_{\ell k}$ denotes the *combination weight* assigned by agent k to neighboring agent ℓ , satisfying $0 < a_{\ell k} \leq 1$, for all $\ell \in \mathcal{N}_k$, $a_{\ell k} = 0$ for all $\ell \notin \mathcal{N}_k$, and $\sum_{\ell \in \mathcal{N}_k} a_{\ell k} = 1$. The algorithm is called "adaptive" due to the step-size parameter $\delta \in (0, 1)$, which allows it to track changes in the true hypothesis θ^* . Observe that the numerator in (3) is the weighted geometric mean of the priors $\psi_{\ell,i}(\theta)$ at time i with weights given by the scalars $\{a_{\ell k}\}$.

Let A_{\star} denote the left-stochastic *combination matrix* consisting of all combination weights $a_{\ell k}$. Regarding the network topology, we impose the following assumption [23], [26], [32], which allows information to flow throughout the whole network.

Assumption 2. (Strongly-connected network). The communication graph is strongly connected (i.e., there exists a path with positive weights linking any two agents, and at least one agent in the graph has a self-loop, meaning that there is at least one agent $k \in \mathcal{N}$ with $a_{kk} > 0$).

Finally, we impose assumptions on the agents' observation models. We assume that the agents can collectively identify the underlying true hypothesis [26], [32].

Assumption 3. (Identifiability assumption). For each wrong hypothesis $\theta \neq \theta^*$, there is at least one agent $k \in \mathcal{N}$ that has strictly positive KL-divergence $D_{KL}(L_k(\theta) || L_k(\theta_*)) > 0$.

We also assume the boundedness of the likelihood func-

tions [27].

Assumption 4. (Bounded likelihoods). There is a finite constant b > 0 such that, for all $k \in \mathcal{N}$:

$$\left|\log\frac{L_k(\boldsymbol{\zeta}|\boldsymbol{\theta})}{L_k(\boldsymbol{\zeta}|\boldsymbol{\theta}')}\right| \le b \tag{4}$$

for all θ , $\theta' \in \Theta$, and $\zeta \in \mathcal{Z}_k$.

III. INVERSE MODELING PROBLEM

A. Problem Statement

In our study, we assume that the graph is completely hidden. The assumption is motivated by the fact that in real-world settings, the pattern of interactions among agents is usually unknown to an external observer. In addition, in the social learning strategy, it is common [28], [32] to assume that for each time $i \ge 1$, each agent local observation $\zeta_{k,i}$ is private and external observers do not have access to it. On the other hand, beliefs (i.e., $\psi_{k,i}(\theta)$) are public and exchanged across the network. For this reason, our goal is to infer the graph topology by observing the exchanged beliefs among the agents.

Formally, we assume that at each time step $i \ge 1$ we observe the beliefs of the agents in the network, collected into the set:

$$\mathcal{D}_{i} = \left\{ \psi_{k,i}(\theta), \ k \in \mathcal{N} \right\}$$
(5)

The problem of interest is to recover the combination matrix A_{\star} based on knowledge of $\{\mathcal{D}_i\}_{i>1}$.

B. Likelihood and Beliefs Ratios

We define the matrices Λ_i and \mathcal{L}_i of size $|\mathcal{N}| \times (|\Theta| - 1)$, where each element is a relative measure of log beliefs and likelihood ratios as follows:

$$[\mathbf{\Lambda}_i]_{k,j} \triangleq \log \frac{\boldsymbol{\psi}_{k,i}(\boldsymbol{\theta}_0)}{\boldsymbol{\psi}_{k,i}(\boldsymbol{\theta}_j)}$$
(6)

$$[\mathcal{L}_i]_{k,j} \triangleq \log \frac{L_k(\boldsymbol{\zeta}_{k,i}|\boldsymbol{\theta}_0)}{L_k(\boldsymbol{\zeta}_{k,i}|\boldsymbol{\theta}_j)},\tag{7}$$

In these expressions, we have chosen some $\theta_0 \in \Theta$ as a reference state, while $\theta_j \neq \theta_0$. Due to Assumption 4, \mathcal{L}_i has finite entries. Since we initially start from positive beliefs $\mu_{k,0}(\theta)$, and the likelihoods remain positive with probability one, it follows from the update rules (2)-(3) that $\Lambda_i < \infty$.

Observe that both matrices vary with the time index i. Based on the definitions (6)-(7), some algebra will show that we can transform (2)-(3) into an update relating these matrices:

$$\boldsymbol{\Lambda}_{i} = (1 - \delta) \boldsymbol{A}_{\star}^{\mathsf{T}} \boldsymbol{\Lambda}_{i-1} + \delta \boldsymbol{\mathcal{L}}_{i}.$$
(8)

At every iteration *i*, the quantities $\{\Lambda_i, \Lambda_{i-1}\}$ are known based on knowledge of the beliefs \mathcal{D}_i from (5). On the other hand, the quantity \mathcal{L}_i is not known because the observations $\{\zeta_{k,i}\}$ are private. We wish to devise a scheme that allows us to estimate A_{\star} in (8) from knowledge of $\{\Lambda_i, \Lambda_{i-1}\}$ and from a suitable approximation for \mathcal{L}_i . Before discussing the learning algorithm, however, we establish the following useful property. For simplicity of notation, we will write

$$\mathbb{E}[\cdot] \triangleq \mathbb{E}_{\boldsymbol{\zeta}_{k,t} \sim L_k(\theta^\star), k \in \mathcal{N}, t \leq i}[\cdot]$$
(9)

where the expectation is relative to the randomness in all local observations up to time i.

Lemma 1 (Mean likelihood matrix). Random variables \mathcal{L}_i are i.i.d. over time and space, and their mean matrix $\overline{\mathcal{L}} = \mathbb{E}\mathcal{L}_i$ is independent of time and finite with each entry equal to:

$$\begin{bmatrix} \boldsymbol{\mathcal{L}} \end{bmatrix}_{k,j} = D_{KL} \left(L_k \left(\boldsymbol{\theta}^{\star} \right) || L_k \left(\boldsymbol{\theta}_j \right) \right) - D_{KL} \left(L_k \left(\boldsymbol{\theta}^{\star} \right) || L_k \left(\boldsymbol{\theta}_0 \right) \right).$$
(10)

Proof. Proof is omitted due to space limitations.

C. Algorithm Development

The linear nature of the update for Λ_i in (8) motivates the following instantaneous quadratic loss function for finding A_* :

$$Q'(A; \mathbf{\Lambda}_i, \mathbf{\Lambda}_{i-1}) = \frac{1}{2} \|\mathbf{\Lambda}_i - (1-\delta)A^{\mathsf{T}}\mathbf{\Lambda}_{i-1} - \delta \mathcal{L}_i\|_{\mathrm{F}}^2,$$
(11)

where $\|\cdot\|_{\rm F}$ denotes Frobenius norm. Computation of \mathcal{L}_i requires knowledge of $\zeta_{k,i}$, $k \in \mathcal{N}$, which is assumed to be private for each agent, therefore hidden from the observer. For this reason, we will assume only knowledge of $\overline{\mathcal{L}}$, which is in principle requires knowledge of the true hypothesis θ^* due to Lemma 1. We explain in the sequel how to circumvent this requirement.

Typically, at each time step $i \ge 1$, every agent $k \in \mathcal{N}$ estimates the true state using (1). It can be shown [26, Theorem 2]) that the probability of error $\mathbb{P}(\widehat{\theta}_{k,i}^{\circ} \neq \theta^{\star}) \to 0$ as $i \to \infty$ and $\delta \to 0$. It can be verified that the same conclusion continues to hold if we estimate the underlying hypothesis based on the intermediate belief vectors (which are the quantities that are assumed to be observable):

$$\widehat{\boldsymbol{\theta}}_{k,i} = \arg\max_{\boldsymbol{\theta}\in\Theta} \boldsymbol{\psi}_{k,i}(\boldsymbol{\theta}).$$
(12)

In order to have agreement on the θ_* among the agents, we will estimate a common $\hat{\theta}_i$ by using a majority vote rule. Then, the following conclusion holds.

Lemma 2 (True state learning error: majority vote).

$$\lim_{\delta \to 0} \mathbb{P}\left(\lim_{i \to \infty} \widehat{\boldsymbol{\theta}}_i \neq \theta^*\right) = 0.$$
(13)

Proof. Proof is omitted due to space limitations.

Therefore, we replace (11) by the following loss function:

$$Q(A; \mathbf{\Lambda}_i, \mathbf{\Lambda}_{i-1}, \bar{\mathbf{\mathcal{L}}}_i) = \frac{1}{2} \|\mathbf{\Lambda}_i - (1-\delta)A^{\mathsf{T}}\mathbf{\Lambda}_{i-1} - \delta\bar{\mathbf{\mathcal{L}}}_i\|_{\mathrm{F}}^2.$$
(14)

with $\bar{\mathcal{L}}_i = \mathbb{E}_{\widehat{\theta}_i} \mathcal{L}_i$, and $\mathbb{E}_{\widehat{\theta}_i}$ means that the expectation is computed assuming that the private data $\zeta_{k,i}$ is generated

according to $\hat{\theta}_i : \zeta_{k,i} \sim L_k(\zeta_{k,i}|\hat{\theta}_i)$. Our minimization problem over a horizon of N observations then becomes:

$$\min_{A} J(A) \triangleq \frac{1}{N} \sum_{i=1}^{N} J_i(A), \tag{15}$$

$$J_i(A) \triangleq \mathbb{E}Q(A; \mathbf{\Lambda}_i, \mathbf{\Lambda}_{i-1}, \bar{\mathbf{\mathcal{L}}}_i)$$
(16)

where the statistical properties of Λ_i vary with time. This explains why we are averaging over a time-horizon in (15). We apply stochastic approximation to solve (15), namely, a recursion of the form:

$$\boldsymbol{A}_{i}^{\mathsf{I}} = \boldsymbol{A}_{i-1}^{\mathsf{I}} + \mu(1-\delta) \\ \times \left(\boldsymbol{\Lambda}_{i} - (1-\delta)\boldsymbol{A}_{i-1}^{\mathsf{T}}\boldsymbol{\Lambda}_{i-1} - \delta\bar{\boldsymbol{\mathcal{L}}}_{i}\right)\boldsymbol{\Lambda}_{i-1}^{\mathsf{T}}$$
(17)

In order to examine the steady-state performance of the algorithm, we introduce an independence assumption that is common in the study of adaptive systems [33].

Assumption 5. (Separation principle) Let $\widehat{A}_i = A_{\star} - A_i$ denote the estimation error. Assume the step-size μ is sufficiently small, so that in the limit, $\|\widetilde{A}_i\|_{\mathrm{F}}^2$ reaches a steady state distribution, and \widetilde{A}_i is independent of Λ_i .

Using this condition, we can establish the following steadystate performance for the Online Graph Learning (OGL) algorithm.

Theorem 1 (Steady-state performance). Under Assumptions 1-5, after large enough number of social learning iterations with $\delta \rightarrow 0$ and for sufficiently small μ , the mean squared deviation converges exponentially fast with:

$$\limsup_{i \to \infty} \mathbb{E} \| \widetilde{\boldsymbol{A}}_i \|_{\mathrm{F}}^2 \le \frac{\mu^2 \gamma}{1 - \alpha} = O(\mu), \tag{18}$$

where

$$\alpha = 1 - 2\mu\nu + O(\mu^2)$$

$$\gamma = \delta^2 \kappa |\mathcal{N}| \lambda_{\max}(\mathcal{R}_{\mathcal{L}})$$

$$\nu = (1 - \delta)^2 \lambda_{\min}(\mathcal{R}_{\Lambda})$$

$$\kappa = (1 - \delta)^2 \lambda_{\max}(\mathcal{R}_{\Lambda})$$
(19)

and $\mathcal{R}_{\mathcal{L}} \triangleq \mathbb{E}(\mathcal{L}_i - \bar{\mathcal{L}})(\mathcal{L}_i - \bar{\mathcal{L}})^{\mathsf{T}}$ is independent of *i*, whereas $\mathcal{R}_{\Lambda} \triangleq \lim_{i \to \infty} \mathbb{E} \Lambda_i \Lambda_i^{\mathsf{T}}$ is finite.

Proof. Proof is omitted due to space limitations.

IV. COMPUTER SIMULATIONS

The experiments that follow help illustrate the ability of the proposed algorithm to identify edges and to adapt to situations where the graph topology is dynamic, as well as the hypothesis.

A. Setup

We consider a network of 30 agents with $|\Theta| = 4$ states, where the adjacency matrix is generated according to the Erdos-Renyi model with edge probability p = 0.2. We set Z_k to be a discrete sample space with $|Z_k| = 4$ for $k \in \mathcal{N}$. The step-size of the model is set to $\delta = 0.05$. We define the likelihood functions $L_k(\theta), k \in \mathcal{N}, \theta \in \Theta$ as follows:

$$L_{k}(\boldsymbol{\zeta}|\boldsymbol{\theta}) = \sum_{z \in \mathcal{Z}_{k}} \mathbb{I}[\boldsymbol{\zeta} = z]\beta_{k,z}(\boldsymbol{\theta}),$$

$$\beta_{k,z}(\boldsymbol{\theta}) \ge 0, z \in \mathcal{Z}_{k}$$

$$\sum_{z \in \mathcal{Z}_{k}} \beta_{k,z}(\boldsymbol{\theta}) = 1,$$
 (20)

where the parameters $\beta_{k,z}(\theta)$ are generated randomly. During the graph learning procedure, we use $\mu = 0.01$.

B. Graph learning

We provide a comparison between the true combination matrix and the estimated combination matrix. We plot the combination matrices in Fig. 1. The experiment shows the ability of the algorithm to identify the graph: the recovered combination weights are close to the actual weights. In no-edge places, we observe reasonably small weights on the recovered matrix. These can be removed in post-processing by simple thresholding or more elaborate schemes, such as the *k*-means algorithm [34].

Additionally, in Fig. 2, we plot how the deviation from the true matrix A_{\star} evolves. The deviation is computed as the following quantity:

$$\|\tilde{A}_i\|_{\mathrm{F}}^2 = \|A_{\star} - A_i\|_{\mathrm{F}}^2.$$
 (21)

We provide the error rates for both algorithm variants with known true state θ^* and estimated true state $\hat{\theta}_i$. We see that there is a negligible gap between the learning performances.

The proposed algorithm is robust to changes in the true state and graph topology. In Fig. 3, we regenerate edges at time 15000. The algorithm adapts and converges to the new combination matrix at a linear rate. Thus, we have experimentally illustrated that the algorithm is stable to dynamic network changes, which is a natural setting to consider in practice. These properties hold because the algorithm is online and processes data one by one with a constant learning rate $\mu > 0$.

V. CONCLUSIONS

In this paper, the problem of graph learning through observing social interactions is investigated. We develop an online algorithm that learns the agents' influence pattern via observing agents' beliefs over time. We prove that the proposed algorithm successfully learns the underlying combination weights matrix and demonstrate its performance through analysis and computer simulations. In this way, we are able to discover the pattern of information flow in the network. A distinct feature of the proposed algorithm is the fact that it can track changes in the graph topology as well as in the true hypothesis.

As future work, we aim to investigate *partial information* setting, where the algorithm has access only to the beliefs of a subset of the network agents.





(b) Learned graph.

Fig. 1: True combination matrix and the learned matrix using the Online Graph Learning (OGL) algorithm.



Fig. 2: Error evolution of two algorithm variants: known true state and estimated true state.



Fig. 3: Error evolution when the graph edges are regenerated.

REFERENCES

- J. A. Barnes, "Graph theory and social networks: A technical comment on connectedness and connectivity," *Sociology*, vol. 3, no. 2, pp. 215– 232, 1969.
- [2] M. E. Newman, D. J. Watts, and S. H. Strogatz, "Random graph models of social networks," *Proc. National Academy of Sciences*, vol. 99, no. suppl 1, pp. 2566–2572, 2002.
- [3] Z. Cui, K. Henrickson, R. Ke, and Y. Wang, "Traffic graph convolutional recurrent neural network: A deep learning framework for networkscale traffic learning and forecasting," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 11, pp. 4883–4894, 2019.
- [4] C. Zheng, X. Fan, C. Wang, and J. Qi, "Gman: A graph multi-attention network for traffic prediction," in *Proceedings of the AAAI Conference* on Artificial Intelligence, vol. 34, no. 01, 2020, pp. 1234–1241.
- [5] C. Chen, W. Ye, Y. Zuo, C. Zheng, and S. P. Ong, "Graph networks as a universal machine learning framework for molecules and crystals," *Chemistry of Materials*, vol. 31, no. 9, pp. 3564–3572, 2019.
- [6] S. Kearnes, K. McCloskey, M. Berndl, V. Pande, and P. Riley, "Molecular graph convolutions: moving beyond fingerprints," *Journal of Computer-Aided Molecular Design*, vol. 30, no. 8, pp. 595–608, 2016.
- [7] V. Kalofolias, "How to learn a graph from smooth signals," in Artificial Intelligence and Statistics. PMLR, 2016, pp. 920–929.
- [8] H. E. Egilmez, E. Pavez, and A. Ortega, "Graph learning from data under laplacian and structural constraints," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 6, pp. 825–841, 2017.
- [9] S. Vlaski, H. P. Maretić, R. Nassif, P. Frossard, and A. H. Sayed, "Online graph learning from sequential data," in 2018 IEEE Data Science Workshop (DSW), Lausanne, Switzerland, 2018, pp. 190–194.
- [10] X. Dong, D. Thanou, M. Rabbat, and P. Frossard, "Learning graphs from data: A signal representation perspective," *IEEE Signal Processing Magazine*, vol. 36, no. 3, pp. 44–63, 2019.
- [11] B. Pasdeloup, V. Gripon, G. Mercier, D. Pastor, and M. G. Rabbat, "Characterization and inference of graph diffusion processes from observations of stationary signals," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 4, no. 3, pp. 481–496, 2017.
- [12] D. Thanou, X. Dong, D. Kressner, and P. Frossard, "Learning heat diffusion graphs," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 3, pp. 484–499, 2017.
- [13] S. P. Chepuri, S. Liu, G. Leus, and A. O. Hero, "Learning sparse graphs under smoothness prior," in *IEEE International Conference on Acoustics*, *Speech and Signal Processing (ICASSP)*, 2017, pp. 6508–6512.
- [14] R. Shafipour, S. Segarra, A. G. Marques, and G. Mateos, "Network topology inference from non-stationary graph signals," in *IEEE International Conference on Acoustics, Speech and Signal Processing* (*ICASSP*), 2017, pp. 5870–5874.
- [15] S. Segarra, A. G. Marques, G. Mateos, and A. Ribeiro, "Network topology identification from spectral templates," in *IEEE Statistical Signal Processing Workshop (SSP)*, Palma de Mallorca, Spain, 2016, pp. 1–5.

- [16] I. Viola, H. P. Maretic, P. Frossard, and T. Ebrahimi, "A graph learning approach for light field image compression," in *Applications of Digital Image Processing XLI*, vol. 10752. International Society for Optics and Photonics, 2018, p. 107520E.
- [17] S. Sardellitti, S. Barbarossa, and P. Di Lorenzo, "Graph topology inference based on transform learning," in *IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, Greater Washington, D.C., USA, 2016, pp. 356–360.
- [18] H. P. Maretic, D. Thanou, and P. Frossard, "Graph learning under sparsity priors," in *IEEE International Conference on Acoustics, Speech* and Signal Processing (ICASSP), New Orleans, LA, USA, 2017, pp. 6523–6527.
- [19] H. Ma, H. Yang, M. R. Lyu, and I. King, "Mining social networks using heat diffusion processes for marketing candidates selection," in *Proceedings of the 17th ACM conference on Information and knowledge management*, 2008, pp. 233–242.
- [20] J. Friedman, T. Hastie, and R. Tibshirani, "Sparse inverse covariance estimation with the graphical lasso," *Biostatistics*, vol. 9, no. 3, pp. 432– 441, 2008.
- [21] V. Matta, A. Santos, and A. H. Sayed, "Graph learning with partial observations: Role of degree concentration," in *IEEE International Symposium on Information Theory (ISIT)*, Paris, France, 2019, pp. 1312– 1316.
- [22] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Nonbayesian social learning," *Games and Economic Behavior*, vol. 76, no. 1, pp. 210–225, 2012.
- [23] 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.
- [24] P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie, "Foundations of nonbayesian social learning," *Columbia Business School Research Paper*, no. 15-95, 2017.
- [25] —, "A theory of non-bayesian social learning," *Econometrica*, vol. 86, no. 2, pp. 445–490, 2018.
- [26] V. Bordignon, V. Matta, and A. H. Sayed, "Adaptive social learning," *IEEE Transactions on Information Theory*, vol. 67, no. 9, pp. 6053–6081, 2021.
- [27] —, "Social learning with partial information sharing," in IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Barcelona, Spain, 2020, pp. 5540–5544.
- [28] 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.
- [29] X. Zhao and A. H. Sayed, "Learning over social networks via diffusion adaptation," in 2012 Conference Record of the Forty Sixth Asilomar Conference on Signals, Systems and Computers (ASILOMAR). IEEE, 2012, pp. 709–713.
- [30] D. Gale and S. Kariv, "Bayesian learning in social networks," *Games and Economic Behavior*, vol. 45, no. 2, pp. 329–346, 2003.
- [31] 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.
- [32] 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.
- [33] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Foundations and Trends* in *Machine Learning*, vol. 7, no. 4-5, pp. 311–801, 2014. [Online]. Available: http://dx.doi.org/10.1561/ 2200000051
- [34] V. Matta, A. Santos, and A. H. Sayed, "Graph learning under partial observability," *Proceedings of the IEEE*, vol. 108, no. 11, pp. 2049– 2066, 2020.