

Echo Chambers and Segregation in Social Networks: Markov Bridge Models and Estimation

Rui Luo^{ID}, Buddhika Nettasinghe^{ID}, and Vikram Krishnamurthy^{ID}, *Fellow, IEEE*

Abstract—This article deals with the modeling and estimation of the sociological phenomena called echo chambers and segregation in social networks. Specifically, we present a novel community-based graph model that represents the emergence of segregated echo chambers as a Markov bridge (MB) process. An MB is a 1-D Markov random field that facilitates modeling the formation and disassociation of communities at deterministic times, which is important in social networks with known timed events. We justify the proposed model with real-world examples and examine its performance on a recent Twitter dataset. We provide a model parameter estimation algorithm based on maximum likelihood and a Bayesian filtering algorithm for recursively estimating the level of segregation using noisy samples obtained from the network. Numerical results indicate that the proposed filtering algorithm outperforms the conventional hidden Markov modeling in terms of the mean-squared error. The proposed filtering method is useful in computational social science where data-driven estimation of the level of segregation from noisy data is required.

Index Terms—Bayesian filtering, echo chamber, Markov bridge (MB), segregation, social network.

I. INTRODUCTION

ONLINE social networks (OSNs) lay the foundation for online community formation. Billions of users rely on OSNs to connect with friends, share information, and advertise products. Echo chambers, i.e., situations where one is exposed only to opinions that agree with their own, are an increasing concern for the usage of OSNs. According to the theory of preferential attachment or homophily [1], users tend to link with other users who share similar attributes (e.g., opinions and interests). Furthermore, social influence [2] also increases users' tendency of becoming more similar to somebody as a result of social interaction. These two factors lead to segregated and polarized clusters known as “echo chambers” on social networks. A vivid example of such echo chambers in




Fig. 1. It indicates how Twitter users' political opinions are polarized into two communities (echo chambers) before the 2020 presidential election. Nodes represent Twitter users and edges represent retweets during the one-month period before election (October 1–November 1). Multiple snapshots illustrating the graph evolution during this period are shown in Fig. 6. The graph is laid out using DrL (a force-directed graph layout) and the nodes are assigned different colors and shapes (blue circles and red triangles) according to the two communities detected by the Louvain method [3]. Details can be found in Section V-B1.

a directed graph is from Twitter users' retweeting behavior before the 2020 presidential election, as shown in Fig. 1.

Echo chambers are studied on various social networks from different modeling perspectives. However, the evolution of echo chambers is characterized by certain temporal patterns in many cases, which is neglected by many proposed models. Having an anticipatory model of segregation in social networks allows us to incorporate the effects of periodic (i.e., seasonal) events into the model. This enables the real-time statistical inference as well as tasks such as offering incentives to reduce the effects of segregation (control strategies for preventing segregation and echo chambers). For example, one can imagine a control strategy, which offers incentives to users at each time instant (subject to budget restrictions) to influence the link formation in order to hinder the segregation in social networks.

Toward this end, the aim of this article is to develop and analyze a model for the anticipatory nature of the segregation

Manuscript received December 25, 2020; revised March 29, 2021; accepted June 6, 2021. Date of publication July 12, 2021; date of current version May 30, 2022. This work was supported in part by the U.S. Army Research Office under Grant W911NF-19-1-0365 and in part by the National Science Foundation under Grant CCF-2112457 and Grant CCF-1714180. This article was presented in part at the IEEE International Conference on Acoustics, Speech and Signal Processing, 2021. (Corresponding author: Rui Luo.)

Rui Luo is with the Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14850 USA (e-mail: rl828@cornell.edu).

Buddhika Nettasinghe and Vikram Krishnamurthy are with the Department of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14850 USA (e-mail: dwn26@cornell.edu; vikramk@cornell.edu).

Digital Object Identifier 10.1109/TCSS.2021.3091168

2329-924X © 2021 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See <https://www.ieee.org/publications/rights/index.html> for more information.

process in a social network, i.e., one can assign probabilities to the event that the social network will be segregated at a certain fixed time instant.

Main Results:

- 1) We present a dynamic network formation model that captures the dynamics of how a social network segregates into disconnected communities, i.e., echo chambers (and then integrates back again). The key idea behind our model is to represent the strength of the ties between communities (in terms of a graph clustering metric) as a Markov bridge (MB) process, which is a special case of an anticipatory process.
- 2) Based on the proposed MB-based segregation model, we propose a time-inhomogeneous Bayesian filter (called hidden Markov bridge (HMB) filter) for recursively estimating the state of the graph clustering metric. The HMB filter uses only a few (compared to graph size) noisy samples from the social network at each time instant.
- 3) We numerically compare the performance of the proposed HMB filter with the conventional hidden Markov model (HMM) filter in terms of mean-squared error. Our results show that the proposed method outperforms the traditional HMM filter. This shows that the Bayesian filter yields useful real-time information about the disassociation and association of communities in a network.
- 4) We evaluate the performance on a publicly available dataset [4], which encompasses 7 million tweets related to the 2020 U.S. presidential election. Our results illustrate the proposed model and filter are useful in estimating Twitter users' state of political opinion polarization under real-world settings.

II. RELATED WORK AND MOTIVATION

A. Segregation and Echo Chambers

Previous works on segregation and echo chambers that are related to ours can be considered under two categories.

1) *Generation Mechanism of Echo Chambers*: Echo chambers are segregated communities that do not interact across communities but only form intracommunity communication. This phenomenon appears in many areas, including political discussion [5], [6], e-commerce [7], and urban planning [8]. Various mechanisms have been proposed in the literature to explain the emergence of echo chambers. One category of such mechanisms focuses on a self-reinforcement procedure, which polarizes a user's opinion by exposure to similar contents or interaction with similar users. Baumann *et al.* [9] proposed a radicalization mechanism, which reinforces extreme opinions from moderate initial conditions. Ge *et al.* [7] explored the effect of repeated exposure to similar contents on users' e-commerce shopping interest. The second category introduces a social feedback mechanism, which modifies a user's opinion by imposing social constraints, e.g., peer effects. Sasahara *et al.* [2] constructed a model of social influence and unfriending where users can change both their opinions and social connections based on the information they received. Banisch and Olbrich [10] considered the social feedback's

effect on users expressing alternative opinions and analyzed the sufficient conditions for stable bipolarization on a stochastic block model.

In this article, we associate users with their known and fixed labels (i.e., customer types in Section III-A and political ideologies in Section V-B). A similar approach has been considered in [10] and [11]. We model echo chambers via the community-level behavior, e.g., how dense are intracommunity and intercommunity connections. More precisely, we construct a segregation measurement using the graph conductance of the corresponding network and the ratio between the number of intracommunity connections and the number of total connections,¹ which represents how users' interactions are confined inside community.

2) *Dynamics of Segregation*: Another direction has explored the dynamics of segregation and opinion polarization. Many related works view segregation as a steady state of a multiagent system and build upon opinion dynamics models, such as the DeGroot model [13], the Friedkin–Johnsen model [14], and the voter model [15].² Dandekar *et al.* [16] complemented a biased assimilation term into DeGroot model, which strengthens the individual's self-opinion and ensures polarization at the steady state. Chitra and Musco [17] augmented the Friedkin–Johnsen model with an external network administrator, which reduces disagreement among interacting users and leads to echo chambers. Friedkin [18] studied the community cleavage problem by comparing different opinion dynamics models' results on different social structures. De *et al.* [19] modeled individual users' opinions over time by marked jump-diffusion stochastic differential equations and identify conditions under which opinions converge to a steady state.

Our work does not build upon agent-based opinion models. Instead, we model the segregation in the network as a temporal signal and propose expectation–maximization (EM)-based parameter estimation from data. This is more amenable to data-driven analysis that is typically used in real-world settings.

B. Why MB Dynamics?

In this article, we focus on modeling the dynamics of segregation and echo chambers by considering the information of future events. To this end, we propose an MB model and justify this model based on real-world examples and a Twitter dataset. The temporal dynamics of social networks give rise to states where the network is segregated into multiple echo chambers at certain known time instants and integrate back into a single community at other known time instants. In statistical signal processing (e.g., in target tracking), dynamical processes with long-range dependencies are typically modeled as MB processes [20]–[22]. An MB process can be viewed as a special case of an anticipatory process in which the distribution at a future known time instant is fixed. In the following,

¹Another option is modularity [12], which represents the strength of division of a network into modules (i.e., echo chambers).

²We remind readers that a user following the DeGroot model averages her opinion with the opinions of her neighbors, whereas a user following the voter model adopts a neighbor's opinion at random. Friedkin and Johnsen [14] extended the DeGroot model by associating a user with an innate opinion.

we discuss how the temporal dynamics of segregation in many social network scenarios can be modeled using MB dynamics.³

Example 1 (Schelling's Segregation Model): The segregation model developed by Schelling [23] is set in an $N \times N$ grid. Agents are split into two groups and occupy the spaces of the grid. Agents desire a fraction B_a of their neighborhood to be from the same group. This model shows how echo chambers and weak intercommunity connectivity might arise even with a moderate individual preference B_a . The physical grid space can be generalized to a social network as a grid graph. In case of external stimulus such as elections, the value of B_a or its distribution is anticipatory and can be modeled as an MB process.

Example 2 (Polarization of Political Opinion): OSNs (e.g., Twitter) users have different political leanings and tend to follow or retweet users of similar opinions. Such tendency can be moderate, i.e., users are open-minded to follow or hear someone from different ideology groups. However, during a politically polarization event, such as election or legislation, the tendency will become stronger and hinder users from connecting with others of different opinions, which results in highly segregated online echo chambers, as shown in Fig. 1. We can formulate the portion $x^{(t)}$ of interactions between users of different political leanings as an MB process. It is anticipatory to be at a low level during the polarization event.

Example 3 (Social Media Marketing): Consider a social media marketing scenario (e.g., Facebook Business page) where a company is connected with customers. Customers are classified into fans and utilitarian customers [24]. As shown in Fig. 2, while fans (bottom-left vertices) have a stable connection strength (i.e., fixed edge weight) with the company (center vertex), utilitarian customers (top-right vertices) have time-varying connection strength (i.e., time-varying edge weights) with the company due to reasons such as sales events. The variable connection leads to segregation and integration of the company–customer social network.

III. MB MODEL FOR DYNAMIC SOCIAL NETWORKS

This section presents a stochastic model to represent the evolution of a social network whose state is fixed at the beginning and at the end. The two fixed states correspond to a segregated social network (with multiple echo chambers) and a social network that has a single community (i.e., an integrated network). Thus, the model presented in this section is a useful, intuitive representation of the process of social network segregation. Furthermore, as we show later in Section IV, the proposed model is easily amenable to the Bayesian statistical inference, making it useful in data-driven contexts in computational social science.

A. Time-Varying Edge Weight Graph Model

This section explains the graph model with time-varying edge weights using a company–customer social network consisting of a company and two types of customers as the graph vertices. The binary classification of customers is motivated from the topological study of Facebook fans in [24].

³Also, see Appendix A for more motivating examples.




Fig. 2. Snapshots of a dynamic company–customer social network model in (1) at four time instants. The edges between utilitarian customers (top-right vertices) and the company (center vertex) grow weaker as the social network evolves into segregation with two communities (echo chambers).

The social network at discrete time instant t is modeled by an undirected, weighted graph $G^{(t)}(V, E, w^{(t)})$, with $|V|$ number of agents, $|E|$ number of undirected edges representing their connectivity, and $w^{(t)} : E^{(t)} \rightarrow \mathbb{R}^+$ representing the weights of the edges (i.e., the strengths of the connections).

Let $|V| = n$ be the number of vertices in the network, $M = \{v_1, \dots, v_m\} \subset V$ be the set of utilitarian customers who are all connected with each other (i.e., form a complete subgraph), $v_{m+1} \in V$ be the company, and $N = \{v_{m+2}, \dots, v_n\} \subset V$ be the set of other customers (fans) that form another complete subgraph. Then, edge weight function $w_{ij}^{(t)}$ between v_i and v_j , where $(v_i, v_j) \in E$, is as follows:

$$w_{ij}^{(t)} = \begin{cases} W_{ij}(t) & \text{if } v_i = v_{m+1} \ v_j \in M \\ & \forall v_j = v_{m+1} \ v_i \in M \\ 1 & \text{if } v_i = v_{m+1} \ v_j \in N \\ & \forall v_j = v_{m+1} \ v_i \in N \\ 1 & \text{if } (v_i, v_j) \in M \ v_i \neq v_j \\ & \forall (v_i, v_j) \in N \ v_i \neq v_j \end{cases} \quad (1)$$

where, $W_{ij}(t)$, $t = 1, 2, \dots$, is the MB process that we define in Section III-B.

Note that (1) classifies the edge weights into three groups: between company and utilitarian customers, between company and fans, and between two customers of the same type. Note that there is no edge between two customers of different types. The edge weights between company and utilitarian customers are subject to sales events and therefore described as a time-evolving random process $W(t)$ (specified in Section III-B). Other weights are simply set to be one. The simplification is reasonable because fans would be indifferent about sales events and have a more stable relationship with the company. Furthermore, we also assume that the customers of the same type are all connected with each other motivated by the concept of homophily [1].

B. MB Model of Edge Weights

We now propose an MB model for the evolution of the weight $W_{ij}(t)$ in the graph. Recall [25] that an MB is a 1-D Markov random field. It is clamped at the beginning and end time point and evolves in between with a three-point transition probability $p\{W_{ij}(t)|W_{ij}(t+1), W_{ij}(t-1)\}$. An MB for $W_{ij}(t)$ facilitates modeling a community that separates and then reintegrates with another community in a network. Unlike a Markov chain that enters a state at a geometrically distributed time, an MB enters a state at a fixed deterministic time [20].

We consider $(2T - 1)$ time steps as the period between two consecutive sales events. The edge weight $W_{ij}(t)$ between company and utilitarian customers reaches maximum at time 1 and time $2T - 1$ when sales event happens and decreases to minimum at time T in the middle of two sales events. The process can be described as two consecutive MBs as we explain next.

The Markov process $W_{ij}(t)$, $t = 1, \dots, 2T - 1$ takes value in some finite state space $S = \{0, (1/(N_S - 1)), \dots, ((N_S - 2)/(N_S - 1)), 1\}$, which is an arithmetic sequence with N_S elements. The transition matrix of the Markov process is chosen to be an $N_S \times N_S$ row-normalized Toeplitz matrix such that transitions from a given state to neighboring states (i.e., values in S that are closer to the given state) are more likely. Let the entries of the transition matrix be $P_{a,b} = P\{W_{ij}(t+1) = S[b] | W_{ij}(t) = S[a]\}$ for all edge weights W_{ij} in (1), where $S[a], S[b] \in S$ are two states of the social network with order $a, b \in \{1, \dots, N_S\}$, respectively.

In this setup, we fix the states at $t = 1, T$, and $2T - 1$ of a Markov process—this can be viewed as two sequential MBs: one that starts at time 1 and another one that starts at time T . Both MBs have their starting and end states fixed. The first MB's end state overlaps the second MB's starting state. The first MB (for each edge) is initialized as 1 and the state at time T is set to be 0, i.e., $W_{ij}^1 = 1$, $W_{ij}^T = 0$. Thus, the transition probability of the first MB going from state $S[a]$ to state $S[b]$ is obtained by applying the Bayes rule as follows [25]:

$$B_{a,b}^c(t) = P\{W_{ij}(t+1) = S[b] | W_{ij}(t) = S[a], W_{ij}(T) = 0\} \\ = \frac{P_{a,b}(P^{T-(t+1)})_{b,c}}{(P^{T-t})_{a,c}} \quad (2)$$

for $t = 1, \dots, T - 2$, where c is the order of 0 in the state space, i.e., $S[c] = 0$.

Likewise, the state of social network is fixed to be 1 at $2T - 1$ (i.e., the last time step). We can then formulate the transition probability of the second MB in a similar manner as follows:

$$B_{a,b}^{c'}(t) = \frac{P_{a,b}(P^{2T-1-(t+1)})_{b,c'}}{(P^{2T-1-t})_{a,c'}} \quad (3)$$

for $t = T - 1, \dots, 2T - 3$, where c' is the order of 1 in the state space, i.e., $S[c'] = 1$. Thus, the dynamics of edge weights is specified by two MBs with transition probability matrices given by (2) and (3) and the fixed initial state of the first MB.




Fig. 3. Two metrics (left: graph conductance $\phi(G)$ and right: algebraic connectivity λ_2) that indicate the strength of connectivity between communities in a graph. The figure shows that conductance (the state random variable in our model) resembles other metrics such as the algebraic connectivity.

C. Graph Clustering Metrics

The aim of this section is to discuss the graph metric called graph conductance that we use to express segregation and set as the state of our model. Graph conductance is a measurement of the level of clustering in a graph and is explained as follows.

We first define a cut $(S, \bar{S})$ as a partition of the vertices of a graph into two disjoint subsets S and $\bar{S}$. The conductance of a cut $(S, \bar{S})$ in a graph is defined as

$$\phi(S) = \frac{\sum_{i \in S} \sum_{j \notin S} w_{ij}}{\min\{a(S), a(V \setminus S)\}}, \quad S \subset V \quad (4)$$

where $a(S) = \sum_{i \in S} \sum_{j \in V} w_{ij}$ is the sum of the weights of all edges with at least one endpoint in S . Then, given a graph G , we define the graph conductance as the minimum conductance over all possible cuts

$$\phi(G) = \min_{S \subset V} \phi(S). \quad (5)$$

Graph conductance is also related to the algebraic connectivity, which is the second smallest eigenvalue of the Laplacian matrix of G . Algebraic connectivity is used in many results in spectral graph theory such as Cheeger's inequality [26]. The derivation of algebraic connectivity can be found in [27]. The weighted adjacency matrix $A^{(t)}$ of the graph is given by

$$A_{ij}^{(t)} = w_{ij}^{(t)}. \quad (6)$$

The degree matrix $D^{(t)}$ is given by

$$D_{ij}^{(t)} = \begin{cases} \sum_k A_{ik}^{(t)}, & \text{if } i = j \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

The Laplacian $L^{(t)}$ of the graph is given by

$$L^{(t)} = D^{(t)} - A^{(t)}. \quad (8)$$

Fig. 3 shows that the variations of both graph conductance and algebraic connectivity follow a similar dynamics. This implies that an estimate of the graph conductance also serves as a proxy for the algebraic connectivity under our model.

IV. BAYESIAN ESTIMATION OF GRAPH METRICS

Section III presented an MB for a social network segregation. A natural question is: assuming the MB with known parameters, how can one estimate the level of segregation in

a data-driven manner? An answer to this question is useful in computational social science and network science that deal with large-scale, partially observable (via noisy samples) social networks. As a solution, we propose a Bayesian filtering method based on the proposed segregation model.

A. Measuring Conductance via Sampled Edges

This section discusses our criteria for obtaining a noisy estimate of the conductance of the underlying dynamic graph (explained in Section III) using a sampled subgraph. We demonstrate that the sampling noise can be approximated as a Gaussian noise from the central limit theorem.

We assume that γN of the total N edges are uniformly sampled and observed at each time t (random sampling of edges has been used widely in the literature in statistical estimation tasks, see [28]–[30]). γ is a fixed ratio in $(0, 1]$. The observed graph conductance $\phi(\tilde{G}^{(t)})$ is computed from the partially sampled graph $\tilde{G}^{(t)}$ at time t . Graph conductance is a static function of edge weights

$$\phi(G^{(t)}) = f(w_1^{(t)}, \dots, w_N^{(t)}). \quad (9)$$

Also, the observed graph conductance is the same function of sampled edge weights

$$\phi(\tilde{G}^{(t)}) = f(w_{i_1}^{(t)}, \dots, w_{i_{\gamma N}}^{(t)}) \quad (10)$$

where $i_1, \dots, i_{\gamma N}$ are sampled from $1, \dots, N$ with equal probabilities. From (9), it follows straightforwardly that graph conductance as a static function of the edge weights follows the same MB dynamics. For the rest of this article, we denote $\phi(G^{(t)})$ and $\phi(\tilde{G}^{(t)})$ as $\phi^{(t)}$ and $\tilde{\phi}^{(t)}$ respectively.

To estimate the observation probabilities $p(\tilde{\phi}^{(t)}|\phi^{(t)} = j)$, we use a Monte Carlo simulation to obtain sample trajectories of the $(2T - 1)$ step graph evolution and compute the empirical cumulative distribution function (cdf) of noise of the conductance computed from the partial observation, i.e., the cdf of the difference between the estimated conductance $\gamma \tilde{\phi}^{(t)}$ and the true conductance $\phi^{(t)}$. Fig. 4 shows that the observation noise is approximately (in the sense of Kolmogorov–Smirnov (KS) test) a Gaussian distribution, i.e.,

$$p(\tilde{\phi}^{(t)}|\phi^{(t)} = i) \sim \mathcal{N}(\gamma \tilde{\phi}^{(t)} - \phi^{(t)}|\mu^{(t)}, (\sigma^{(t)})^2). \quad (11)$$

The normal distribution form of the observation noise (11) can also be viewed as a consequence of the central limit theorem; since we are sampling independent identically distributed (i.i.d) edge sequences from the social network and approximate the graph conductance using the average of their weights, it follows from the central limit theorem that the sample mean (scaled by the square root of the number of samples) converges in distribution to a Gaussian distribution centered around the true state.

B. HMB Filter

In this section, we aim to estimate the segregation level of a social network by computing the posterior probability of the graph conductance given its sampled observation (the conductance computed from sampled edges of the graph).




Fig. 4. Empirical cdf of the sampling noise for graph conductance can be fitted as a Gaussian distribution. It contains the empirical and fitted cumulative probability at four states in a simulation of six-state model. The p -value of KS test is above a 0.05 significance level, and therefore, the Gaussian distribution null hypothesis is unrejectable.

Section IV-A exploits the Gaussian approximation of measurement noise and we propose an HMB filter here for recursively tracking the state of the graph conductance. HMB filter is a generalization of the time-homogeneous HMM filter [31] and has been widely used in signal processing methods for target tracking [20]–[22].

Suppose that the MB process $\Phi = \{\phi^{(1)}, \dots, \phi^{(t)}\}$ is observed via the observation process $\tilde{\Phi} = \{\tilde{\phi}^{(1)}, \dots, \tilde{\phi}^{(t)}\}$. Assume that the observation at time t given the state $\phi^{(t)}$ is conditionally independent of $\phi^{(\tau)}$ and $\tilde{\phi}^{(\tau)}$, $\tau \neq t$. This conditional independence implies that

$$P(\tilde{\phi}^{(1)}, \dots, \tilde{\phi}^{(t)}|\phi^{(1)}, \dots, \phi^{(t)}) = \prod_{k=1}^t P(\tilde{\phi}^{(k)}|\phi^{(k)}). \quad (12)$$

The process $\tilde{\Phi}$ is called an HMB because the property (12) is analogous to the assumption made for HMM. Consider the HMB $\tilde{\Phi}$ with state Φ , known MB transition probability (2), and precomputed observation probability (11). The posterior probability can be evaluated recursively via Bayes' rule

$$q_j(t+1) = \frac{p(\tilde{\phi}^{(t+1)}|\phi^{(t+1)} = j) \sum_{i=1}^{\Phi} B_{i,j}^k(t) q_i(t)}{\sum_{l=1}^{\Phi} p(\tilde{\phi}^{(t+1)}|\phi^{(t+1)} = l) \sum_{i=1}^{\Phi} B_{i,j}^k(t) q_i(t)} \quad (13)$$

as shown in [32].

V. NUMERICAL EXAMPLES ON SOCIAL MEDIA MARKETING AND TWITTER POLITICAL RETWEETS

In this section, we numerically illustrate that the proposed HMB filter (Section IV) outperforms (in terms of mean-squared error) the widely used HMM filter for estimating the level of segregation on synthetic data. This highlights how the proposed model and filtering method can be useful in estimating the level of segregation with a better accuracy




Fig. 5. Mean-squared error of the proposed HMB filter compared with an HMM filter in the company–customer marketing network simulation. HMB filter outperforms HMM filter by approximately 20%, which indicates its better prediction ability of segregation in social networks.

compared to the baseline method of HMM filtering. We also evaluate the proposed model on a public Twitter election dataset.

A. Simulation on Weighted Customer-Merchant Graph

In this section, we consider a social media marketing scenario where the connection between company and customers is modeled as an MB. An HMB filter is implemented to estimate the intercommunity distance based on sampled observation of single edge weights and an additive Gaussian noise. It outperforms a hidden Markov chain filter regarding the mean-squared error.

1) *Simulation Setup*: We consider a company–customer network of ten utilitarian customers, 20 fans, and one company as discussed in Section III for $2T - 1$ time steps ($T = 20$). The state space of the weight of each edge between utilitarian customers and company is an arithmetic sequence $[1, ((N_S - 2)/(N_S - 1)), \dots, (1/(N_S - 1)), 0]$. The weight evolves according to a transition matrix that is a Toeplitz matrix. Each descending diagonal from left to right is constant: $[(1/4)^{N_S-1}, \dots, 1, \dots, (1/4)^{N_S-1}]$. Each row vector of the Toeplitz matrix is normalized so that the row elements add up to 1. We then implement the HMB filter in assuming that the measurement noise is Gaussian with the empirically estimated mean and covariance in Section IV-A. To assess the performance, we compare the mean-squared error of the HMB filter with an HMM filter that assumes the underlying process is a Markov chain (instead of an MB).

2) *Numerical Results of Filters*: Fig. 5 shows the results obtained using the above simulation setup. Results show that the proposed HMB filter outperforms the HMM filter for all considered numbers of states (N_S values). Thus, the numerical results indicate that the proposed Bayesian filter is capable of accurately estimating the level of segregation in a company–consumer network from noisy sampled edges.

B. HMB Model of Political Polarization During 2020 Election

In this section, we propose that the MB model is sociologically beneficial for predicting the emergence and segregation

TABLE I
SAMPLE OF ELECTION-RELATED TWITTER ACCOUNTS TRACKED IN THE DATASET

account name	political party
@realDonaldTrump	R
@GovBillWeld	R
@MarkSanford	R
@WalshFreedom	R
@JohnDelaney	R
@AmbassadorRice	R
@TrumpWarRoom	R
@TeamTrump	R
@JoeBiden	D
@CoryBooker	D
@GovernorBullock	D
@SenKamalaHarris	D
@BernieSanders	D
@SenWarren	D
@marwilliamson	D
@AndrewYang	D

level of echo chambers on a social network. We justify our conclusion on a real-world Twitter dataset where a polarization score is defined on the Twitter retweet network. We determine the model parameters (i.e., the transition matrix) from maximum-likelihood algorithm, which is derived in Appendix C. We apply the HMB filter to estimate the polarization score. The filter’s estimation accuracy outperforms an HMM filter regarding mean-squared error. We also verify that the observation noise can be approximated as a Gaussian distribution based on statistical hypothesis testing.

1) *Construction of Retweet Network and Polarization Score*: We leverage a publicly available dataset that encompasses 240 million tweets related to the 2020 U.S. presidential election. The dataset captured tweets with specific user mentions and accounts (57 in total) that are tied to president candidates and politicians. A sample of such accounts is shown in Table I. The column of political party denotes the party to which this account belongs to (D-democratic and R-republican). We select and sample tweets from October 1 to November 1, spanning a 30-day period before the election day (Nov. 3), which is about 7 million tweets in total (see Appendix B for data collection and sampling procedure).

From this subset of tweets, we constructed a dynamic retweet graph $G^{(t)}(V, E^{(t)})$, $t = 1, \dots, 30$, where the nodes represent $|V|$ Twitter accounts, directed and unweighted edges $E_{ij}^{(t)}$ from node i to node j if user j retweets a message⁴ originally posted by user i on day t . We filtered out nodes with out-degree fewer than 2 (which means that they only retweet once during the 30-day period). In summary, the retweet network has $|V| = 1\,399\,644$ vertices and $\sum_{t=1}^{30} |E^{(t)}| = 5\,047\,498$ edges during the specified period. In Fig. 6, we take a sample of the retweet graph and plot four snapshots of its largest weakly connected component. It clearly shows the pattern of partitioning into two polarized echo chambers.

To study the evolution of political opinion polarization during the 30-day period, we define a temporal variable named polarization score. We select users who have retweeted

⁴Here, we do not consider “quote tweets” (retweet with a comment added) to avoid the use of “quote tweets” for ironic or criticizing purposes.




Fig. 6. Snapshots of a dynamic retweet network with two communities at four different timestamps before the Nov. 3 presidential election. The ratio of intracommunity connections and total connections (sum of intracommunity and intercommunity connections) evolves and is modeled as an MB in (14). The graph is laid out using DrL (the same as in Fig. 1).

election-related accounts from both political parties. We estimate these users' political leaning (interchangeably referred to as ideology) as follows. Every retweet to accounts from either political party increases the count for that side by +1. The user's political leaning is classified as the side with more accumulated retweets. We dismiss the users whose retweets are equally sourced from two political parties. After this ideology classification, we denote $D = \{v_1, \dots, v_m\} \subset V$ be the set of m left-leaning users, and $R = \{v_{m+1}, \dots, v_{m+n}\} \subset V$ is the set of n right-leaning users. We then define the polarization score of the retweet network on day t as the ratio of the amount of intraideological retweets (e.g., user from D retweets an election-related account who is from the Democratic party) to the amount of intraideological retweets plus the amount of cross-ideological retweets (e.g., user from D retweets an election-related account who is from the Republican party) on day t . We collected 28 such 30-day sequences of polarization scores, which are used for training the transition matrix and computing the empirical distribution of the observation error

$$y^{(t)} = \frac{|E_{ij}^{(t)}|_{(v_i, v_j) \in D \vee (v_i, v_j) \in R}}{|E_{ij}^{(t)}|_{\forall (v_i, v_j)}}. \quad (14)$$

2) *HMB Filter Estimation Results:* We formulate $x^{(t)}$ as an HMB with $y^{(t)}$ as its observation. The hidden states form an arithmetic sequence with maximum and minimum




Fig. 7. Empirical cdf of the observation noise for polarization score on the Twitter dataset and the cdf of a Gaussian distribution of four hidden states (states 2–5) in the six-state model. The p -value of KS test indicates that the observation noise can be approximated by the Gaussian noise.




Fig. 8. Mean-squared estimation error of the polarization score on the Twitter dataset. It shows that the proposed HMB filter outperforms an HMM filter in estimating the level of polarization in a social network.

corresponding to those of the observation. We anticipate a high-level opinion polarization near or on election day since people are required to vote for one candidate from one party. Therefore, we set the destination, i.e., the final state to be the maximal state.

$x^{(t)}$ follows transition probabilities as depicted by (2) with c as the maximal state and transition matrix P derived from maximum-likelihood algorithm in Appendix C. The observation noise is hypothetically verified to follow a Gaussian distribution, as shown in Fig. 7.

To assess the performance of the proposed filter in (13), we compare its mean-squared error with an HMM filter. Results show that the proposed HMB filter outperforms the HMM filter for all considered number of hidden states and reduces the mean-squared error by 10% (see Fig. 8).

C. Discussion on the Experiment With Real-World Twitter Data

Section V explored modeling and estimation of segregation in social networks. For the first social media marketing example, we proposed a tractable model for segregation based on MB processes of the connection strength between company and customers. We also proposed a Bayesian filter (named HMB filter) to estimate the level of segregation (as measured by graph conductance) by polling some random pairs of neighbors (i.e., edges) in the network.

To validate the proposed model with real-world data, our second opinion polarization example studied Twitter users' retweet behavior before the 2020 presidential election. We anticipated a high level of polarization near the election and modeled the portion of retweets between users of the same political ideology as an MB process. To evaluate the computational complexity and the scalability to large networks of the proposed filter, we provide a discussion as follows.

1) *Computational Complexity of the Proposed Filter*: The computational complexity of the proposed HMB filter can be decomposed into two parts. The first part is the computation of the MB transition probability (2), which requires computing P^t , i.e., the power of the $n \times n$ Markov transition matrix.

1.1) If P is diagonalizable and $P = Q^{-1}DQ$, then computing P^t is equivalent to computing $Q^{-1}D^tQ$. It takes $O(n^3)$ to diagonalize P (using, e.g., QR iteration) and $O(n \log t)$ to take each diagonal element to the t th power. Therefore, the computational complexity is

$$O(n^3).$$

1.2) If P is not diagonalizable,⁵ then computing P^t can be done in time

$$O(n^3 \log t).$$

Above are cases where we do not consider memory. Suppose that we have memory for P^{t-1} , and then computing P^t will only require $O(n^3)$ time for both cases.

The second part is the computation cost of filtering, i.e., computing the posterior probability of the true polarization score according to (13). Writing (13) in matrix-vector notation

$$q(t+1) = \frac{O_{\tilde{\phi}(t+1)} B' q(t)}{\mathbf{1}' B' q(t)} \quad (15)$$

where

$$O_{\tilde{\phi}(t+1)} = \text{diag}[p(\tilde{\phi}^{(t+1)} | \phi^{(t+1)} = 1), \dots, p(\tilde{\phi}^{(t+1)} | \phi^{(t+1)} = n)] \quad (16)$$

$$q(t+1) = [q_1(t+1), \dots, q_n(t+1)]' \quad (17)$$

which requires $O(n^2)$ time for computing the posterior.

To summarize, the proposed filter imposes a natural tradeoff between estimation accuracy (i.e., how fine-grained is the state space) and computational complexity. Furthermore, the filter is scalable to large networks, provided that an efficient sampling (pooling) approach is implemented.

⁵An example is $P = \frac{1}{5} \begin{bmatrix} 2 & 2 & 1 \\ 1 & 2 & 1 \\ 2 & 1 & 3 \end{bmatrix}$, which has eigenvalues $1, \frac{1}{5}$, while the eigenspace of $\frac{1}{5}$ is 1-D.

2) *Future Directions*: One major limitation of the study is that we equally weight each retweet. The model and the approaches presented in this article can be made more practical via the following directions.

2.1) *Incorporating Tweet's Temporal Effect*: We could formulate the influence of tweets to decrease exponentially after posted so that newer tweets would be more informative of user's political opinion. While many tweets spread within a limited time window, some could achieve high virality. Therefore, we could assign different decay rates for each tweet as well.

2.2) *Defining User's Political Ideology as a Time-Evolving Variable*: Although a user is typically unchanged as to his favorable party, his political leaning might change due to personal experience and external information. We could adapt the correspondence analysis or the maximum-likelihood logit model in [33] to obtain users' political ideology at different times and assign political scores accordingly to the user-related tweets.

2.3) *Incorporating the Content and Sentiment of Tweets*: It is found in [34] that tweets with negative emotion tend to get reposted more rapidly and frequently than positive and neutral messages. Taking one step forward, we could decompose each tweet's popularity into its political effect and content effect and mitigate the unbalance of different sentiments or wording among different tweets.

2.4) *Combining Information From Multilevel Social Graphs*: In addition to retweet network, we could construct other types of social network from Twitter data, including mention, follow, and comment. We could certainly analyze the trend of polarization separately on these networks and average the results, and however, it is more promising to construct a multilevel heterogeneous graph to learn a structural representation of polarization.

VI. CONCLUSION

This article studied the sociological phenomena of segregation and echo chambers in social networks. We proposed an MB dynamics-based model for evaluating the interaction between customers and company in a social media marketing scenario. We then justified the model by looking at the evolution of political opinion polarization on a real-world Twitter dataset. We formulated an additive Gaussian measurement noise model for the MB, derived the EM algorithm for estimating parameters of the HMB model, and proposed an HMB filter to estimate the state of segregation and echo chambers based on samples of the social network. The numerical results indicated that our filter outperforms time-homogeneous filters, such as an HMM filter.

Future directions of this work include further improving the accuracy of the proposed method using different sampling methods based on friendship paradox (see [28]–[30]), enriching this framework to handle more sophisticated network topologies such as heterogeneous graphs, and incorporating the HMB model with generation models to forecast opinion dynamics (see [19], [35], [36]).

APPENDIX A EXAMPLES OF REAL-WORLD MB DYNAMICS

In the following, we illustrate several motivating examples (in addition to those provided in Section II-B) whose dynamics can be well captured by the MB model.

Example 4 (Spread of Coronavirus Fake News): On social media, people's perception of Coronavirus evolves during the outbreak of this pandemic and leaves room for related fake news. Many factors prevent the spread of fake news, including scientific reports from reliable news sources, government releases, and users' tendency to share health and prevention messaging. We can model the user's tendency $x^{(t)}$ to share Coronavirus fake news as an MB process. It will decrease with the elucidations from reliable sources and also bring down the probability that fake news echo chambers emerge. This example is related to Example 2 in Section II-B because it is found in [37] that partisanship correlates with sentiment toward government measures. Therefore, the evolution of Coronavirus fake news echo chambers is correlated with that of political ideology echo chambers.

Example 5 (Activity Level of Seasonal Sports League): Sports leagues typically have season and off-season. In different periods of a year, fans will have different involvement in the sports leagues on social media such as online sports forums. We can formulate the active level $x^{(t)}$ as a hidden state following an MB process. We can then use the data of posting, commenting, and time of stay on the forums as the observation. $x^{(t)}$ will evolve from a high level to a low level during the off-season and will evolve back into a high level in the next sports season and lead to the segregated fan's community. Online merchants may take advantage of this information to maximize their advertisement coverage and return on investment.

Example 6 (E-Commerce Sequential Recommendation): Relevance and diversity usually act as two competing objectives in recommender systems, where the former causes growing concern that it might lead to the self-reinforcing of user's interests due to narrowed exposure of similar items. The existence of echo chambers has been validated on user clicks, purchases, and browse logs from Alibaba Taobao in [7]. To examine and quantify the echo chambers in recommender systems, we can use a measure $x^{(t)}$ to represent the similarity of recommended items during the interaction with users. For conventional recommender systems that narrow down the contents provided to users, $x^{(t)}$ can be modeled as an MB process, which evolves from a low level to a high level. In this context, an extended recommendation framework can potentially avoid such echo chamber emergence by using collaborative filtering and sequential forecasting to recommend users items that they may find useful in the future, thus improving both user satisfaction and E-commerce platform's revenue.

APPENDIX B TWITTER DATASET DETAILS

The Twitter dataset used in Section V-B was provided in [4] and publicly available.⁶ We used Twitter's streaming API

⁶The dataset website: <https://github.com/eichen102/us-pres-elections-2020>

through Tweepy and kept track of tweets with specific keyword mentions and accounts related to the 2020 U.S. presidential election since May 2019. The data contain approximately 1% stream of all tweets in real time.

In obedience to Twitter's Developer Agreement & Policy, only tweet's Tweet ID is shared. The Tweet ID is preserved in text files in temporal order. We used DocNow's Hydrator⁷ to retrieve the tweet objects⁸ with full tweet payloads, including the tweet poster, content, timestamp, and the author who is retweeted from.

To reduce the amount of data, we use systematic sampling, which means that we pick every n th tweets ($n = 20$) to comprise the data used in the research. For each tweet, we keep its poster, author, and timestamp if it is a retweet (if the tweet object has a "retweeted_status" attribute).

APPENDIX C EM ALGORITHM FOR HMB PARAMETERS

For real-life application of the MB model proposed in Section III, we need to obtain the model parameters that are useful for filtering and forecasting the segregation state of the social network. This section presents the EM algorithm, which serves the purpose of finding the maximum-likelihood estimate of the parameters.

A. Forward-Backward Smoothing Algorithm for HMB

In this section, we derive the forward-backward (also named Baum-Welch) algorithm for smoothing the HMB model. Consider HMB model with parameter $\theta = (S, P, O)$, where S is an X -state Markov chain with transition matrix $P = (P_{a,b})$, $a, b \in S = \{s_1, \dots, s_X\}$, and O is the HMB emission probability function. The HMB model has unknown state sequence $X^{(T)} = (x^{(1)}, \dots, x^{(T)})$ and observation sequence $Y^{(T)} = (y^{(1)}, \dots, y^{(T)})$. We know the destination of the state sequence $X^{(T)}$ is $c \in S$.

We first go through the forward procedure by defining

$$\alpha_\theta^{(t)}(a) = P(x^{(t)} = a, y^{(1)}, \dots, y^{(t)} | \theta) = P(x^{(t)} = a, Y^{(t)} | \theta) \quad (18)$$

which is the probability of seeing the partial sequence $(y^{(1)}, \dots, y^{(t)})$ and ending up in state a at time t .

The backward procedure is similar by defining a backward variable

$$\beta_\theta^{(t|T)}(a) = P(y^{(t+1)}, \dots, y^{(T)} | x^{(t)} = a). \quad (19)$$

We now define

$$\gamma_\theta^{(t)}(a) = P(x^{(t)} = a, Y^{(T)} | \theta) \quad (20)$$

which is the probability of being in state a at time t for the observation sequence $(y^{(1)}, \dots, y^{(T)})$. It can be derived in terms of $\alpha_\theta^{(t)}$ and $\beta_\theta^{(t|T)}$

$$\gamma_\theta^{(t)}(a) = \frac{\alpha_\theta^{(t)}(a)\beta_\theta^{(t|T)}(a)}{\sum_{b=s_1}^{s_X} \alpha_\theta^{(t)}(b)\beta_\theta^{(t|T)}(b)}. \quad (21)$$

⁷<https://github.com/DocNow/hydrator>

⁸<https://developer.twitter.com/en/docs/twitter-api/v1/data-dictionary/object-model/tweet>

We also define

$$\gamma_{\theta}^{(t)}(a, b) = P(x^{(t)} = a, x^{(t+1)} = b, Y^{(T)} | \theta) \quad (22)$$

which is the probability of being in state a at time t and being in state b at time $t + 1$. This can be expressed in terms of $\alpha_{\theta}^{(t)}$ and $\beta_{\theta}^{(t|T)}$

$$\begin{aligned} \gamma_{\theta}^{(t)}(a, b) &= \frac{\alpha_{\theta}^{(t)}(a) B_{a,b}(t) O(y^{(t+1)} | x^{(t+1)} = b) \beta_{\theta}^{(t+1|T)}(b)}{\sum_{a=s_1}^{s_X} \sum_{b=s_1}^{s_X} \alpha_{\theta}^{(t)}(a) B_{a,b}(t) O(y^{(t+1)} | x^{(t+1)} = b) \beta_{\theta}^{(t+1|T)}(b)}. \end{aligned} \quad (23)$$

B. Maximum-Likelihood Estimation Algorithm

We assume that the HMB model is observed in Gaussian noise, i.e., the emission probability follows a zero-mean Gaussian distribution

$$y^{(t)} = x^{(t)} + v^{(t)}, \quad v^{(t)} \sim N(0, \sigma^2). \quad (24)$$

The E-step of the EM algorithm finds the expected value of the complete-data log likelihood with respect to the unknown state $X^{(T)} = (x^{(1)}, \dots, x^{(T)})$ given the observation $Y^{(T)} = (y^{(1)}, \dots, y^{(T)})$ and the current parameter estimates $\theta^I = (S, P, \sigma)$. This log likelihood is defined as the Q function

$$Q(\theta, \theta^I) = E\{\log P(Y^{(T)}, X^{(T)} | \theta) | Y^{(T)}, \theta^I\}. \quad (25)$$

The second step (the M-step) of the EM algorithm is to maximize the expectation we computed in the first step, that is, we find

$$\theta^{I+1} = \underset{\theta}{\operatorname{argmax}} Q(\theta, \theta^I). \quad (26)$$

The joint probability of states $X^{(T)}$ and observations $Y^{(T)}$ given the parameter is formulated as

$$\begin{aligned} \log P(Y^{(T)}, X^{(T)} | \theta) &= \log \prod_{t=1}^T P(y^{(t)} | x^{(t)}) P(x^{(t)} | x^{(t-1)}, x^{(T)} = c) \\ &= \sum_{t=1}^T [\log P(y^{(t)} | x^{(t)}) + \log P(x^{(t)} | x^{(t-1)}, x^{(T)} = c)] \\ &= \sum_{t=1}^T \sum_{a=s_1}^{s_X} [I(x^{(t)} = a) \log P(y^{(t)} | x^{(t)} = a)] \\ &\quad + \sum_{t=1}^T \sum_{a=s_1}^{s_X} \sum_{b=s_1}^{s_X} [I(x^{(t)} = a, x^{(t+1)} = b) \\ &\quad \log P(x^{(t+1)} = b | x^{(t)} = a, x^{(T)} = c)] \\ &= \sum_{t=1}^T \sum_{a=s_1}^{s_X} I(x^{(t)} = a) \left[\log \left(\frac{1}{\sqrt{2\pi}\sigma} \right) - \frac{(y^{(t)} - a)^2}{2\sigma^2} \right] \\ &\quad + \sum_{t=1}^T \sum_{a=s_1}^{s_X} \sum_{b=s_1}^{s_X} [I(x^{(t)} = a, x^{(t+1)} = b) \log B_{a,b}^c(t)]. \end{aligned} \quad (27)$$

The Q function can then be simplified as

$$\begin{aligned} Q(\theta, \theta^I) &= -\frac{T}{2} \log \sigma^2 - \frac{1}{2\sigma^2} \sum_{t=1}^T \sum_{a=s_1}^{s_X} (y^{(t)} - a)^2 \gamma_{\theta^I}^{(t)}(a) \\ &\quad + \sum_{t=1}^T \sum_{a=s_1}^{s_X} \sum_{b=s_1}^{s_X} \gamma_{\theta^I}^{(t)}(a, b) \log B_{a,b}^c(t) \\ &= \dots + \sum_{t=1}^T \sum_{a=s_1}^{s_X} \sum_{b=s_1}^{s_X} \gamma_{\theta^I}^{(t)}(a, b) \log P_{a,b} \frac{(P^{T-(t+1)})_{b,c}}{(P^{T-t})_{a,c}}. \end{aligned} \quad (28)$$

Recursively solving $(\partial Q(\theta, \theta^I) / \partial \theta) = 0$ for the model parameter θ^{I+1} . Each iteration is guaranteed to improve log likelihood, and the algorithm is guaranteed to converge to a local maximum.

REFERENCES

- [1] M. McPherson, L. Smith-Lovin, and J. M. Cook, "Birds of a feather: Homophily in social networks," *Annu. Rev. Sociol.*, vol. 27, no. 1, pp. 415–444, Aug. 2001.
- [2] K. Sasahara, W. Chen, H. Peng, G. L. Ciampaglia, A. Flammini, and F. Menczer, "Social influence and unfollowing accelerate the emergence of echo chambers," 2019, *arXiv:1905.03919*. [Online]. Available: <http://arxiv.org/abs/1905.03919>
- [3] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks," *J. Stat. Mech., Theory Exp.*, vol. 2008, no. 10, Oct. 2008, Art. no. P10008.
- [4] E. Chen, A. Deb, and E. Ferrara, "#Election2020: The first public Twitter dataset on the 2020 U.S. Presidential election," 2020, *arXiv:2010.00600*. [Online]. Available: <http://arxiv.org/abs/2010.00600>
- [5] J. Bright, "Explaining the emergence of political fragmentation on social media: The role of ideology and extremism," *J. Comput.-Mediated Commun.*, vol. 23, no. 1, pp. 17–33, Jan. 2018.
- [6] P. Barberá, J. T. Jost, J. Nagler, J. A. Tucker, and R. Bonneau, "Tweeting from left to right: Is online political communication more than an echo chamber?" *Psychol. Sci.*, vol. 26, no. 10, pp. 1531–1542, Oct. 2015.
- [7] Y. Ge *et al.*, "Understanding echo chambers in E-commerce recommender systems," in *Proc. 43rd Int. ACM SIGIR Conf. Res. Develop. Inf. Retr.*, Jul. 2020, pp. 2261–2270.
- [8] A. Mele, "Does school desegregation promote diverse interactions? An equilibrium model of segregation within schools," *Amer. Econ. J., Econ. Policy*, vol. 12, no. 2, pp. 228–257, 2020.
- [9] F. Baumann, P. Lorenz-Spreen, I. M. Sokolov, and M. Starnini, "Modeling echo chambers and polarization dynamics in social networks," *Phys. Rev. Lett.*, vol. 124, no. 4, Jan. 2020, Art. no. 048301.
- [10] S. Banisch and E. Olbrich, "Opinion polarization by learning from social feedback," *J. Math. Sociol.*, vol. 43, no. 2, pp. 76–103, Apr. 2019.
- [11] P. R. Neary, "Competing conventions," *Games Econ. Behav.*, vol. 76, no. 1, pp. 301–328, Sep. 2012.
- [12] M. E. J. Newman, "Modularity and community structure in networks," *Proc. Nat. Acad. Sci. USA*, vol. 103, no. 23, pp. 8577–8582, Jun. 2006.
- [13] M. H. DeGroot, "Reaching a consensus," *J. Amer. Statist. Assoc.*, vol. 69, no. 345, pp. 118–121, Mar. 1974.
- [14] N. E. Friedkin and E. C. Johnsen, "Social influence and opinions," *J. Math. Sociol.*, vol. 15, nos. 3–4, pp. 193–206, 1990.
- [15] R. A. Holley and T. M. Liggett, "Ergodic theorems for weakly interacting infinite systems and the voter model," *Ann. Probab.*, vol. 3, no. 4, pp. 643–663, Aug. 1975.
- [16] P. Dandekar, A. Goel, and D. T. Lee, "Biased assimilation, homophily, and the dynamics of polarization," *Proc. Nat. Acad. Sci. USA*, vol. 110, no. 15, pp. 5791–5796, Apr. 2013.
- [17] U. Chitra and C. Musco, "Analyzing the impact of filter bubbles on social network polarization," in *Proc. 13th Int. Conf. Web Search Data Mining*, Jan. 2020, pp. 115–123.
- [18] N. E. Friedkin, "The problem of social control and coordination of complex systems in sociology: A look at the community cleavage problem," *IEEE Control Syst. Mag.*, vol. 35, no. 3, pp. 40–51, Jun. 2015.
- [19] A. De, I. Valera, N. Ganguly, S. Bhattacharya, and M. G. Rodriguez, "Learning and forecasting opinion dynamics in social networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 397–405.

- [20] M. Fanaswala, V. Krishnamurthy, and L. White, "Destination-aware target tracking via syntactic signal processing," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2011, pp. 3692–3695.
- [21] M. Fanaswala and V. Krishnamurthy, "Detection of anomalous trajectory patterns in target tracking via stochastic context-free grammars and reciprocal process models," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 1, pp. 76–90, Feb. 2013.
- [22] M. Fanaswala and V. Krishnamurthy, "Spatiotemporal trajectory models for metalevel target tracking," *IEEE Aerosp. Electron. Syst. Mag.*, vol. 30, no. 1, pp. 16–31, Jan. 2015.
- [23] T. C. Schelling, "Dynamic models of segregation," *J. Math. Sociol.*, vol. 1, no. 2, pp. 143–186, 1971.
- [24] E. Wallace, I. Buil, L. de Chernatony, and M. Hogan, "Who 'likes' you... and why? A typology of Facebook fans: From 'fan'-atics and self-expressives to utilitarians and authentic," *J. Advertising Res.*, vol. 54, no. 1, pp. 92–109, 2014.
- [25] L. B. White and F. Carravetta, "Optimal smoothing for finite state hidden reciprocal processes," *IEEE Trans. Autom. Control*, vol. 56, no. 9, pp. 2156–2161, Sep. 2011.
- [26] F. R. Chung and F. C. Graham, *Spectral Graph Theory*, no. 92. Providence, RI, USA: American Mathematical Society, 1997.
- [27] N. M. M. de Abreu, "Old and new results on algebraic connectivity of graphs," *Linear Algebra Appl.*, vol. 423, no. 1, pp. 53–73, May 2007.
- [28] B. Nettasinghe and V. Krishnamurthy, "What do your friends think?: Efficient polling methods for networks using friendship paradox," *IEEE Trans. Knowl. Data Eng.*, vol. 33, no. 3, pp. 1291–1305, Mar. 2021.
- [29] B. Nettasinghe and V. Krishnamurthy, "The friendship paradox: Implications in statistical inference of social networks," in *Proc. IEEE 29th Int. Workshop Mach. Learn. Signal Process. (MLSP)*, Oct. 2019, pp. 1–6.
- [30] B. Nettasinghe and V. Krishnamurthy, "Maximum likelihood estimation of power-law degree distributions via friendship paradox based sampling," 2019, *arXiv:1908.00310*. [Online]. Available: <http://arxiv.org/abs/1908.00310>
- [31] V. Krishnamurthy, *Partially Observed Markov Decision Processes*. Cambridge, U.K.: Cambridge Univ. Press, 2016.
- [32] G. Stamatescu, L. B. White, and R. Bruce-Doust, "Track extraction with hidden reciprocal chains," *IEEE Trans. Autom. Control*, vol. 63, no. 4, pp. 1097–1104, Apr. 2018.
- [33] P. Barberá, "Birds of the same feather tweet together: Bayesian ideal point estimation using Twitter data," *Political Anal.*, vol. 23, no. 1, pp. 76–91, 2015.
- [34] S. Tsugawa and H. Ohsaki, "Negative messages spread rapidly and widely on social media," in *Proc. ACM Conf. Online Social Netw.*, Nov. 2015, pp. 151–160.
- [35] G. Amati, S. Angelini, F. Capri, G. Gambosi, G. Rossi, and P. Vocca, "Twitter temporal evolution analysis: Comparing event and topic driven retweet graphs," *IADIS Int. J. Comput. Sci. Inf. Syst.*, vol. 11, no. 2, pp. 155–162, 2016.
- [36] M. T. Thij, T. Ouboter, D. Worm, N. Litvak, H. van den Berg, and S. Bhulai, "Modelling of trends in Twitter using retweet graph dynamics," in *Proc. Int. Workshop Algorithms Models Web-Graph*. Cham, Switzerland: Springer, 2014, pp. 132–147.
- [37] J. Jiang, E. Chen, S. Yan, K. Lerman, and E. Ferrara, "Political polarization drives online conversations about COVID-19 in the United States," *Hum. Behav. Emerg. Technol.*, vol. 2, no. 3, pp. 200–211, 2020.



Rui Luo received the B.S. degree from the Department of Automotive Engineering, Tsinghua University, Beijing, China, in 2019. He is currently pursuing the Ph.D. degree with the Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY, USA.

His current research interests include network science and complex systems.



Buddhika Nettasinghe received the M.A.Sc. degree from the Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, BC, Canada, in 2016. He is currently a Ph.D. candidate with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY, USA.

His current research interests include statistical inference and learning, network science, and complex systems.



Vikram Krishnamurthy (Fellow, IEEE) received the Ph.D. degree from Australian National University Canberra, ACT, Australia, in 1992.

From 2002 to 2016, he was a Professor and the Canada Research Chair at The University of British Columbia, Vancouver, BC, Canada. He is currently a Professor with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY, USA. He is the author of the books *Partially Observed Markov Decision Processes* and *Dynamics of Engineered Artificial Membranes and Biosensors* (Cambridge University Press, 2016 and 2018, respectively).

His research interests include statistical signal processing and stochastic control in social networks and adaptive sensing.

Dr. Krishnamurthy served as a Distinguished Lecturer for the IEEE Signal Processing Society and the Editor-in-Chief for the IEEE JOURNAL ON SELECTED TOPICS IN SIGNAL PROCESSING. In 2013, he was awarded an Honorary Doctorate from the Royal Institute of Technology (KTH), Sweden.