



Cascade-based Echo Chamber Detection

Marco Minici
marco.minici@icar.cnr.it
University of Pisa, Pisa, Italy
ICAR-CNR, Rende(CS), Italy

Federico Cinus
cinus@diag.uniroma1.it
Sapienza University, Rome, Italy
ISI Foundation, Turin, Italy

Corrado Monti
corrado.monti@centai.eu
CENTAI, Turin, Italy

Francesco Bonchi
bonchi@centai.eu
CENTAI, Turin, Italy
Eurecat, Barcelona, Spain

Giuseppe Manco
giuseppe.manco@icar.cnr.it
ICAR-CNR, Rende(CS), Italy

ABSTRACT

Despite echo chambers in social media have been under considerable scrutiny, general models for their detection and analysis are missing. In this work, we aim to fill this gap by proposing a probabilistic generative model that explains social media footprints—i.e., social network structure and propagations of information—through a set of latent communities, characterized by a degree of echo-chamber behavior and by an opinion polarity. Specifically, echo chambers are modeled as communities that are permeable to pieces of information with similar ideological polarity, and impermeable to information of opposed leaning; this allows discriminating echo chambers from communities that lack a clear ideological alignment.

To learn the model parameters we propose a scalable, stochastic adaptation of the Generalized Expectation Maximization algorithm, that optimizes the joint likelihood of observing social connections and information propagation. Experiments on synthetic data show that our algorithm is able to correctly reconstruct ground-truth latent communities with their degree of echo-chamber behavior and opinion polarity. Experiments on real-world data about polarized social and political debates, such as the Brexit referendum or the COVID-19 vaccine campaign, confirm the effectiveness of our proposal in detecting echo chambers. Finally, we show how our model can improve accuracy in auxiliary predictive tasks, such as stance detection and prediction of future propagations.

CCS CONCEPTS

• **Computing methodologies** → **Learning in probabilistic graphical models**; • **Information systems** → *Social networking sites*;

KEYWORDS

echo chambers, information propagation, probabilistic modeling

ACM Reference Format:

Marco Minici, Federico Cinus, Corrado Monti, Francesco Bonchi, and Giuseppe Manco. 2022. Cascade-based Echo Chamber Detection. In *Proceedings of the*

31st ACM International Conference on Information and Knowledge Management (CIKM '22), October 17–21, 2022, Atlanta, GA, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3511808.3557253>

1 INTRODUCTION

Social-media platforms have substantially altered the landscape of societal debates. By delivering an extremely large amount of content to online users, they enable quick and easy access to information and facilitate participation in public debates. This positive effect is intertwined by the growing phenomenon that online political discourses, especially on socially relevant issues, tend to fragment and polarize opinions. As a result, the propagation of information is affected by users' propensity to select and promote claims that adhere to their beliefs and ignore or even contrast dissenting information.

The “echo chamber” effect in social media refers to groups of users that, by being exposed solely to like-minded individuals, tend to reinforce each other's pre-existing opinions. This effect has been put under scrutiny as a possible culprit of increased polarization and radicalization [22]. Thus, several studies have been devoted to providing empirical evidence of the existence of echo chambers [11, 16, 21], with variable results, depending on the specific platforms and contexts. For instance, on Reddit echo chambers seem to be less prominent [16], while on Twitter, it has been shown that information propagates in well-separated echo chambers [13, 14, 31]. However, this literature proposes ad hoc approaches for the detection of echo chambers in specific platforms and contexts, while a ground-up approach to detect echo chambers through a formal model of their behavior is still missing.

Prior studies [4–6, 9, 33, 39, 42] have explored the role of communities in information propagation. The underlying assumption of these studies is that a user's activities and social connections are the visible effects of a latent stochastic diffusion process governed by community-level causal factors. As a result, the proposed models successfully devise communities through the lenses of social contagion and are able to characterize community membership in terms of propensity to filter and/or promote community-relevant information. Unfortunately, these models do not take into account the latent relationships between polarization and information diffusion that justify the formation of ideological groups and ultimately characterize the echo-chamber effect.

Inspired by this family of approaches for community detection, in this paper we tackle the problem of detecting echo chambers by observing the way polarized information propagates in a social network. Similarly to communities, echo chambers are defined by groups

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.
CIKM '22, October 17–21, 2022, Atlanta, GA, USA
© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9236-5/22/10...\$15.00
<https://doi.org/10.1145/3511808.3557253>

of nodes who interact and exchange information in a social network. However, while communities in networks are simply defined as high-density clusters in the social graph, echo chambers only need enough structure to allow information to propagate, without the high density typical of a close-knit group of friends. Moreover, as put by Alatawi et al. [1], an echo chamber can be defined as a community that spontaneously emerges as the most effective for spreading polarized content, where conflicting opinions are ignored or even discredited. In other terms, we expect echo chambers to facilitate the flow, through its internals, of information that is ideologically aligned with its opinion, while preventing the flow of information with an opposed leaning. Based on this intuition of echo-chamber behavior, we introduce a generative probabilistic model that encodes, in probabilistic terms, the following core concepts: (i) Each community is characterized by a measurable degree of polarization and highly polarized communities represent echo chambers. (ii) Analogously, the participation of users to a specific community can be measured by means of a given engagement degree. (iii) A polarized cascade can only occur within an echo chamber, provided that the corresponding polarities are aligned. Furthermore, the likelihood of a user contributing to a cascade depends on their level of engagement in the corresponding community. (iv) Social connections are likely to occur between community members; however, it is possible to explain such connections differently according to whether the underlying latent community is considered an echo chamber or not.

Notably, the underlying latent parameters for such a model can be efficiently inferred by resorting to a suitable adaptation of the Generalized Expectation Maximization algorithm [8], which exploits samples of observable propagations and social connections to learn such parameters through an alternating gradient-based optimization strategy. As a result, the learning is scalable and likely to produce communities that can be amenable to coherent explanations in terms of echo-chamber behavior and opinion polarity.

Paper contributions and roadmap. Our technical contributions can be summarized as follows:

- We propose a community-aware information propagation model that explains the creation of social links and the diffusion of items in terms of homogeneity and alignment within a polarized ideological space (§3).
- We devise a scalable gradient-based optimization procedure to learn both the communities and their degree of polarization, maximizing an approximation of the likelihood of a set of information cascades (§4).
- We provide an extensive empirical evaluation of our proposal on synthetic and real-world datasets and show that our inference algorithm is effective, provides meaningful and interpretable communities, and can be used to predict auxiliary tasks such as activation of users on a given cascade, or their stance (§5).

In the next section, we discuss relevant related work.

2 RELATED WORK

Echo chambers: causes and traits. The mechanisms behind the formation of echo chambers are still subject of investigation, however, three main phenomena seem to play key roles in the formation process: algorithmic recommendation (both content [20], and people [12] recommenders), confirmation bias [17, 38], and homophily [28]. The literature also presents several characterizations of echo chambers highlighting their distinctive traits, such as distorted information patterns [27], users similarity growth [11], users psychological profiles [7], the presence of polarization effects [21, 25], and the capacity of spreading polarized content [1].

Echo chambers: detection. Some prior works have analysed social media information to detect echo chambers in a variety of different platforms, without proposing a general method [11, 38, 41]. Usually, a combination of different sources of information is considered, such as textual features (e.g., tweets and hashtags) and interaction network (e.g., retweets, mentions, follow). While some studies (e.g. [10]) focus mainly on the former, in this work we exploit, together with the social graph, the polarity of information pieces and their cascades, which have been shown to be more effective in detecting communities [37] and user-level stances [2]. Other researchers have approached the task of detecting echo chambers as a community detection task, thus exploiting the network structure, followed by an interpretation step to properly identify communities that have echo-chamber traits [14, 15, 18, 19, 24]. In our model, the community structure and the level of echo chamberness of each community are learnt *jointly*. A recent effort by Morini et al. [36] adopts a hybrid approach, employing network-based methods together with NLP tools for content analysis. In particular, the authors propose to first infer users' ideology on a controversial issue, thus constructing a debate network, then detect polarity-homogeneous communities. Their framework is the most similar to ours in terms of input and output: nonetheless, it differs from our proposal as it does not exploit the mutual influence of social bonds and cascades. Moreover, it lacks an explanation for the inferred echo chambers, as it is not based on a model, instead it is a concatenation of pre-existing techniques.

Community detection from cascades. Barbieri et al. [4] introduce a stochastic generative model that relies on a mixture of memberships to learn communities from observed cascades. Similarly to our approach, learning is performed by means of Expectation-Maximization: the derived model can both output the overlapping communities and the users' level of participation. Nevertheless, this approach does not take in consideration the polarity of content and the differences between normal social communities and echo chambers. In the context of dynamic networks, He et al. [26] use cascade diffusion models to discover overlapping communities. In the same context, Sattari and Zamanifar [40] propose a hybrid approach that relies on label propagation and cascades models to learn overlapping communities in dynamic networks. A related research line has attempted inferring the communities from the information cascades *only*, i.e., when the underlying network is not observable [6, 37]. Despite this major difference, the proposed approaches share important features with our work, such as inferring the individual membership by maximum-likelihood. Finally, Monti et al. [35] have recently proposed to exploit cascades models

to learn users' opinions/leanings in social networks. In particular, [35] introduces a stochastic model-based approach to learning, by gradient-based optimization, the ideological leaning of each user in a multidimensional ideological space. While their framework targets directly the user-level representation, in our work we focus on the mesoscale community structure.

3 MODELING ECHO CHAMBERS

We next introduce our probabilistic generative model of echo chambers, which has two main goals. First, to provide a generative description of the echo-chamber phenomenon. Second, to draw on this description in order to identify echo chambers in the wild. We start by identifying some variables as observables —representing the evidence of the phenomena we aim at modeling— and others as latent— i.e., the components that explain the observables.

Observables. We consider the following input:

- (1) A directed social graph $G = (V, E)$, where V represents a set of social media users, and E a set of links, where $(u, v) \in E$ represents the fact that user u is followed by user v .
- (2) A set of items $\mathbb{I}$, where each item $i \in \mathbb{I}$ is labeled with a polarity $p_i \in [-1, 1]$, that characterize its ideological content (e.g., with respect to a given political axis). As we are interested in modeling echo chambers, we assume that our input set $\mathbb{I}$ only contains polarized items, i.e., whose p_i value is not close to 0.
- (3) For each item $i \in \mathbb{I}$, its cascade $\mathbb{D}_i \subseteq V$ in the social graph, i.e., the set of nodes that propagated (or consumed) item i .

Latent variables. We assume that nodes in V can be grouped in latent communities. Some of these communities are *echo chambers*; that is, they facilitate the flow, through its internals, of information that is ideologically aligned with its opinion, while preventing the flow of information with an opposed leaning. By contrast, we call *social communities* those communities that are likely to incorporate ideologically heterogeneous nodes.

Given a set C of latent communities, the value $\eta_c \in [-1, 1]$ (with $c \in \{1, \dots, K\}$) indicates both the polarization and the degree of echo-chamber behavior of the community c . In particular, the value $|\eta_c| = 1$ indicates an ideal echo chamber community, while $|\eta_c| = 0$ indicate the perfect social community.

We assume that each observation (i.e., the links of the social graph and the cascades) is the result of a stochastic process where people act in the network according to their fuzzy membership to latent communities. More specifically, we assume two prior components for a given node $u \in V$ and community c :

- $\theta_{c,u} \in [0, 1]$ represents the level of *polarized engagement* of user u in an echo chamber c ;
- $\phi_{c,u} \in [0, 1]$ is the level of *social engagement* of user u in a social community c .

Both θ_c and ϕ_c represents categorical distributions and model how likely is that a user u contributes to c . Each phenomenon in the social network can be explained by the above latents: a link (u, v) can only be observed if u and v are part of the same close-knit community or if they are part of echo chambers with the same polarity. Similarly, each item i is produced by a community c and it propagates by flowing through the nodes of that community.

3.1 Modeling links and propagations

We use the above latent variables to devise a stochastic process generating the observables. As discussed previously, we assume that echo chambers facilitate the flow of information which is ideologically aligned with their opinion, while preventing the flow of information with an opposed leaning. Following this assumption, the propagation of an item i is generated by considering its polarity p_i and the polarity of the echo chamber η_c . Specifically, echo chambers only allow items with the same polarity; therefore, a propagation is allowed only if $\text{sign}(p_i) = \text{sign}(\eta_c)$, and is allowed with a probability depending on the degree of echo-chamber behavior of c ; i.e., a Bernoulli trial with probability $|\eta_c|$. Finally, it depends on how strong is the item polarity, i.e. a Bernoulli trial with probability $|p_i|$. If these conditions are respected, a node is chosen from the categorical distribution of the community, i.e. $u \sim \text{Cat}(\theta_c)$. This process gives rise to the following likelihood of observing a propagation $\mathbb{D}_i$ with polarity p_i :

$$P(\mathbb{D}_i | c) = \max(0, p_i \cdot \eta_c) \prod_{u \in \mathbb{D}_i} \theta_{c,u}. \quad (1)$$

In other words, item propagations can only be explained by an alignment between the item polarity and the sign of η_c . In fact, the term $\max(0, p_i \cdot \eta_c)$ is only positive when both p_i and η_c exhibit the same sign. If this is the case, users can contribute to the propagation according to their degree of echo-chamber involvement.

Similarly, social links in G are generated as follows. Each community c chooses whether it is an echo chamber with a Bernoulli random trial with probability $|\eta_c|$. If it is, it will create a link $(u, v) \in E$ by extracting two nodes u, v by using the θ_c , i.e. $u \sim \text{Cat}(\theta_c)$. Otherwise, it will do so by using the social engagement ϕ_c . This procedure defines the probability of each link (u, v) given that it was latently generated by a given community c as

$$P((u, v) \in E | c) = |\eta_c| \cdot \theta_{c,u} \theta_{c,v} + (1 - |\eta_c|) \cdot \phi_{c,u} \phi_{c,v}. \quad (2)$$

The above probability follows from this stochastic process where users contribute to the underlying community according to the polarity of the community itself. In fact, $|\eta_c|$ represents a characterization of the community either as an echo chamber (and consequently links forms with θ_c and propagations are possible) or as a social community (links occur thanks to ϕ_c).

Latent priors. Equations 1 and 2 model conditional probabilities for links and propagations, given c . In order to specify the unconditional likelihood, we introduce the categorical priors π_ℓ and π_f ; those define the probability of creating respectively links and propagations. The term $\pi_\ell(c)$ (resp. $\pi_f(c)$) represents the prior probability of a link (resp. a propagation) within c . However, according to our assumption, the probability $\pi_f(c)$ strongly depends on the polarity η_c : when $\eta_c \approx 0$, propagations cannot be explained through c . This constraint can be enforced by introducing the Dirichlet priors α^f and α^l defined as

$$\alpha_c^f = h \cdot |\eta_c| + \epsilon, \quad \alpha_c^l = s \cdot (1 - |\eta_c|) + h \cdot |\eta_c| \quad (3)$$

where hyperparameters $s > 0$, $h > 0$ represent the prior importance of social and echo chamber communities (respectively) in generating links, and ϵ is a regularization value (e.g. 10^{-5}). Therefore, we can generate π_f and π_ℓ through sampling from Dirichlet distributions parameterized by α^f and α^l .




Figure 1: Generated examples for the model with different values for s, h in order to have different balances of echo chambers and social communities in the network: Social with $(s = 16, h = 8)$, Balanced with $(s = 8, h = 8)$, and Polarized with $(s = 8, h = 16)$.

Likelihood. We can finally specify the likelihood for both links and propagations. Given the model parameters $\Theta = \{\theta, \phi, \eta\}$ and the hyperparameters s and h , we have:

$$P(\ell|\Theta; s, h) = \int \left\{ \sum_c P(\ell|c) \pi_\ell(c) \right\} \text{Dir}(\pi_\ell; \alpha^\ell) d\pi_\ell$$

$$P(\mathbb{D}_i|\Theta; s, h) = \int \left\{ \sum_c P(\mathbb{D}_i|c) \pi_f(c) \right\} \text{Dir}(\pi_f; \alpha^p) d\pi_f$$
(4)

For readers' convenience, we provide a notation reference in Table 1.

3.2 Generative process

We can summarize the aforementioned procedure as a simple generative stochastic process for data generation that adheres to the aforementioned modeling assumptions. The process assumes that V and I are given; then, based on the model parameters it generates both links and item propagations with the following processes.

Links. The generative process for a link ℓ is:

- (i) Pick a community $c_\ell \sim \text{Cat}(\pi_\ell)$.
- (ii) Pick $y_c \sim \text{Bernoulli}(|\eta_{c_\ell}|)$ (whether c_ℓ is an echo chamber).
- (iii) If $y_c > 0$, Pick two nodes $u, v \sim \text{Cat}(\theta_c)$.
- (iv) Else, pick two nodes $u, v \sim \text{Cat}(\phi_c)$.
- (v) Add the arc (u, v) to E .

Table 1: Notation reference.

Variable	Meaning
η_c	Polarity of community c
$\theta_{c,u}$	Polarized engagement of user u in community c
$\phi_{c,u}$	Social engagement of user u in community c
p_i	Polarity of item i
$\pi_\ell(c)$	Prior link probability in community c
$\pi_f(c)$	Prior propagation probability in community c
α_c^ℓ	Parameter of the Dirichlet distrib. that defines π_ℓ
α_c^f	Parameter of the Dirichlet distrib. that defines π_f
h	Link generation strength of echo chambers
s	Link generation strength of social communities
$\gamma_{\ell,c}$	Posterior to observe a link ℓ in community c
$\xi_{\delta,c}$	Posterior to observe a propagation δ in community c




Figure 2: Example community and propagation generated for the network represented in Figure 1a. Left (a) shows the membership for one social-type community with shades of green. Right (b) shows one propagation, spreading within an echo-chamber community with homogeneous polarity, where polarity is represented by node colors.

Items. The generative process for the propagation of an item i with polarity p_i is:

- (i) Repeat:
 - pick $c \sim \text{Cat}(\pi_f)$;
 - $y_i \sim \text{Bernoulli}(g_i)$ where $g_i(c) = \max(0, p_i \cdot \eta_c)$;
 - until $y_i > 0$.
- (ii) Pick a user $u \sim \text{Cat}(\theta_c)$ and let $\mathbb{D}_i = \{u\}$.
- (iii) Repeat:
 - let $F_i = \{u | (v, u) \in E, v \in D_i, u \notin D_i\}$;
 - pick the next user
$$u \sim \text{Cat}(\theta_c \cdot [u \in F_i]);$$
 - add u to the set of activated nodes D_i ;
 - until $F_i = \emptyset$ or D_i has reached a given size.

3.3 Generated networks

Here, we show and analyze some networks generated by the generative process we devised. We consider three different sets of parameters s, h (defined in Equation 3), in order to obtain a varying degree of echo-chamber behavior in the network. In all these networks, we generate 5 communities with a fixed $\eta = [-1, -0.5, 0.0, 0.5, 1]$: two opposing echo chambers, a purely social community, and two

cases in-between. Then, we randomly generate θ and ϕ for $N = 256$ nodes. Considering $s \in \{-1, 1\}$ we define two echo-chamber priors $\alpha_c^s = \max(0, s \cdot \eta_c) \cdot \sigma_s + \epsilon$ where σ_s are concentration parameters. Analogously, we define a social-type prior for ϕ with parameters $\alpha^0 = (1 - |\eta|) \cdot \sigma_0$. Then, for each node u we obtain the membership $\tilde{\theta}_u = U_p \cdot P + (1 - U_p) \cdot U_n \cdot N$ and $\tilde{\phi}_u = (1 - U_p \cdot U_n) \cdot S$, where $P \sim \text{Dir}(\alpha^P)$, $N \sim \text{Dir}(\alpha^N)$, $S \sim \text{Dir}(\alpha^0)$ and $U_p, U_n \sim \text{Bernoulli}(\delta)$ (we set $\delta = .3$ in experiments). Finally, we generate the network according to our generative procedure, producing 2048 links (an average of 8 links per node). Similarly, we generate 2048 propagations of items; to generate item polarities, we draw them as $p_i = 2X - 1$ with $X \sim \text{Beta}(\mu, \mu)$, where $\mu = 0.25$ is a parameter regulating the ideological strength of the generated items. Figure 1 shows the three graphs obtained with these different settings of (s, h) . In each visualization, the color gradient represents the polarities for each node, obtained as $\eta \cdot \theta$ (i.e., the weighted average polarity of each community for a given node). We observe that the first graph (Figure 1a) appears not to be shaped by polarized communities, but instead nodes with similar polarities are scattered across the network. In Figure 2 we further explore the first generated network from Figure 1a, where links are predominantly generated by social-type communities. We depict two features of this data set: a social community, and a propagation. We see that social communities, while embedded in a dense network (as in real-world data sets), are closely-knitted in the network. Nevertheless, propagations still happen inside echo-chamber communities, spreading across ideologically aligned nodes.

4 MODEL LEARNING

Given G and $\mathbb{I}$, the optimal Θ parameters can be learned by maximizing the total likelihood

$$P(E, \mathbb{I} | \Theta) = \prod_{\ell \in E} P(\ell | \Theta) \prod_{\substack{\mathbb{D}_i \\ i \in \mathbb{I}}} P(\mathbb{D}_i | \Theta).$$

First of all, we notice that Equation 4 can be simplified by exploiting the conjugacy of the Dirichlet Distribution [8]:

$$P(\ell | \Theta) = \sum_c P(\ell | c) \pi_\ell(c), \quad P(\mathbb{D}_i | \Theta) = \sum_c P(\mathbb{D}_i | c) \pi_f(c),$$

where

$$\pi_\ell(c) = \frac{\alpha_c^\ell}{\sum_{c'} \alpha_{c'}^\ell}, \quad \pi_f(c) = \frac{\alpha_c^f}{\sum_{c'} \alpha_{c'}^f}. \quad (5)$$

A potential problem with the resulting optimization problem is represented by the contribution of each propagation in the total likelihood.

By comparing Equations 1 and 2, we observe that the probability of a link embeds a product over two probabilities, whereas by contrast the probability of a propagation embeds the product over multiple probabilities. Thus, long propagations have very low probability and as a consequence the whole learning process is dominated by link probabilities. This issue makes it difficult to effectively learn the latent variable θ , and consequently the detection of echo chambers.

This problem can be addressed by resorting to a surrogate version of the above likelihood. In practice, we can consider the weighted multi-graph $G^\mathbb{I} = (V, E^\mathbb{I})$ induced by all propagations,

with $\mathcal{A}(u, v, p_i) \in E^\mathbb{I}$ if $u, v \in \mathbb{D}_i$ representing a *sharing link* (i.e., both u and v share an item i , characterized by polarity p_i). Then, the probability of observing such a link can be directly adapted from Equation 1:

$$P(\mathcal{A}(u, v, p) | c) = \max(0, p \cdot \eta_c) \theta_{c,u} \theta_{c,v}$$

Thus, the total likelihood can be rewritten into $P(E, E^\mathbb{I}) = \prod_{\ell \in E} P(\ell | \Theta) \prod_{\mathcal{A} \in E^\mathbb{I}} P(\mathcal{A} | \Theta)$ that allows a more balanced approach through stochastic backpropagation, where each batch can include a sample of both social connections and sharing links. We further simplify the optimization problem by resorting to a variational approximation. Let $X \subseteq E \cup E^\mathbb{I}$ be a batch of social connections and sharing links, Y be a set of corresponding binary variables representing the latent community assignment for both social connections and sharing links; that is, $y_{\ell,c} = 1$ (resp. $y_{\mathcal{A},c} = 1$) if ℓ (resp. $\mathcal{A}$) is associated to community c . Observe that

$$\log P(X, Y | \Theta) = \sum_{\ell, \mathcal{A} \in X} \sum_c \left\{ \log P(\ell | \Theta, c) + \log \pi_\ell(c) + \log P(\mathcal{A} | \Theta, c) + \log \pi_f(c) \right\},$$

and define

$$\begin{aligned} \mathcal{Q}(\Theta, \Theta' | X) &= \mathbb{E}_{Y|X, \Theta'} [\log P(X, Y | \Theta)] \\ &= \sum_{\ell, \mathcal{A} \in X} \sum_c \left\{ P(y_{\ell,c} | \ell, \Theta') (\log P(\ell | \Theta, c) + \log \pi_\ell(c)) \right. \\ &\quad \left. + P(y_{\mathcal{A},c} | \mathcal{A}, \Theta') (\log P(\mathcal{A} | \Theta, c) + \log \pi_f(c)) \right\}. \end{aligned} \quad (6)$$

Notably, whenever $\mathcal{Q}(\Theta, \Theta' | X) \geq \mathcal{Q}(\Theta', \Theta' | X)$, then $\log P(X | \Theta) \geq \log P(X | \Theta')$. In fact,

$$\begin{aligned} \log P(X | \Theta) &= \mathbb{E}_{Y|X, \Theta'} [\log P(X | \Theta)] \\ &= \mathbb{E}_{Y|X, \Theta'} \left[\log \frac{P(X, Y | \Theta)}{P(Y | X, \Theta)} \right] \\ &= \mathcal{Q}(\Theta, \Theta' | X) - \mathbb{E}_{Y|X, \Theta'} [\log P(Y | X, \Theta)] \\ &\geq \mathcal{Q}(\Theta', \Theta' | X) - \mathbb{E}_{Y|X, \Theta'} [\log P(Y | X, \Theta)] \quad (a) \\ &= \mathcal{Q}(\Theta', \Theta' | X) - \mathbb{E}_{Y|X, \Theta'} [\log P(Y | X, \Theta')] \\ &\quad - \mathbb{E}_{Y|X, \Theta'} \left[\log \frac{P(Y | X, \Theta)}{P(Y | X, \Theta')} \right] \\ &\geq \mathcal{Q}(\Theta', \Theta' | X) - \mathbb{E}_{Y|X, \Theta'} [\log P(Y | X, \Theta')] \quad (b) \\ &= \mathbb{E}_{Y|X, \Theta'} \left[\log \frac{P(X, Y | \Theta')}{P(Y | X, \Theta')} \right] \\ &= \log P(X | \Theta'), \end{aligned}$$

where (a) holds by hypothesis, and (b) by Jensen's inequality. This enables an iterative optimization strategy where, for each iteration t , we sample a batch X of social connections and sharing links, and then apply the following alternating steps:

- (*Expectation*) For each $\ell, \mathcal{A} \in X$ and community c , compute the posteriors

$$\begin{aligned} \gamma_{\ell,c} &\equiv P(y_{\ell,c} | \ell, \Theta^{(t)}) = \frac{P(\ell | c) \pi_\ell(c)}{\sum_{\hat{c}} P(\ell | \hat{c}) \pi_\ell(\hat{c})} \\ \xi_{\mathcal{A},c} &\equiv P(y_{\mathcal{A},c} | \mathcal{A}, \Theta^{(t)}) = \frac{P(\mathcal{A} | c) \pi_f(c)}{\sum_{\hat{c}} P(\mathcal{A} | \hat{c}) \pi_f(\hat{c})}, \end{aligned} \quad (7)$$

given the current parameter set $\Theta^{(t)} = \{\eta^{(t)}, \theta^{(t)} \phi^{(t)}\}$

Algorithm 1 ECD Inference

Input: Graph $G = (V, E)$; Sharing links E .
Hyper-parameters: number of communities C , social prior size s ,
echo-chamber prior size h , learning rate λ ,
number of optimization steps for each iteration H .
Output: polarities η , memberships θ and ϕ .

- 1: Randomly initialize $\Theta^{(0)} = \{\eta^{(0)}, \theta^{(0)}, \phi^{(0)}\}$ and set $t = 0$.
- 2: **repeat**
- 3: let $\Theta^{(*)} = \Theta^{(t)}$
- 4: **for** $w \in \{1, \dots, H\}$ **do**
- 5: Sample X from $E \cup E^{\perp}$.
- 6: **for each** $\ell, s \in X$ and $c \in \{1, \dots, C\}$ **do**
- 7: Compute posteriors $\gamma_{\ell,c}$ and $\xi_{s,c}$ according according to Eqs. 7, 5 and
the current parameters $\Theta^{(t)}$. ▷ E Step
- 8: **end for**
- 9: Compute the expected likelihood $\mathcal{Q}$ according to Eqs. 6 and 5 and the
posteriors γ and ξ .
- 10: Update the parameters: ▷ M Step
 $\Theta^{(*)} = \Theta^{(*)} + \lambda \nabla_{\Theta} \mathcal{Q}(\Theta^{(*)}, \Theta^{(t)} | X)$
- 11: **end for**
- 12: Set $\Theta^{(t+1)} = \Theta^{(*)}$ and increase t .
- 13: **until** convergence

- (Optimization) Ascend the gradient $\nabla_{\Theta} \mathcal{Q}(\Theta, \Theta^{(t)} | X)$ to obtain $\Theta^{(t+1)}$.

The whole procedure, dubbed ECD (Echo Chamber Detection), is described in Algorithm 1.

It converges to a local minimum, for a sufficiently small learning rate, since it preserves the general properties of stochastic backpropagation. In fact, although there is no guarantee that the improvement in the likelihood of the current batch corresponds to an improvement in the likelihood of the whole set of observables, this property occurs on average and can eventually be improved by, at each iteration, freezing $\Theta^{(t)}$ and applying the E and M steps on multiple batches.

Implementation details. The implementation follows the structure depicted in the previous sections. The vector η is fed into a $\tanh(\cdot)$ in order to constraint its values into $[-1, +1]$, θ and ϕ are modeled as a 2-layer GCN [29] using 1024 hidden units, the social graph, one-hot encoding attributes, and an output layer with $|C|$ components, that are then fed to a softmax and a sigmoid function, respectively for θ and ϕ . The latter decision is intuitive: echo-chamber communities compete to attract users, while each user could belong to multiple social communities. For modeling, the implementation of θ and ϕ is transparent since they are normalized w.r.t. communities before the likelihood computation as mentioned at the end of Section 3.2. We train the overall architecture through the stochastic algorithm described above, using Adam optimizer with default settings and one epoch. To balance the contribution of links and propagations, we randomly oversample the minority class between the two to achieve a balanced distribution.

5 EXPERIMENTS

In this section, we empirically asses our proposal and answer the following research questions:

- **RQ1.** Assuming that a data set is generated according to the generative process described in Section 3, is the ECD inference algorithm discussed in Section 4 able to estimate its original parameters? Under which conditions? (Section 5.1)

Table 2: Results from synthetic experiments with different configurations of parameters (s, h) used for generation. For each metric, we report its mean and its standard deviation across 10 experiments. Metrics indicate, respectively, the MAE between polarities for each community, between the original and estimated social interest of nodes in communities, between the original and estimated membership of nodes in communities, the correlation between original and estimated polarities for each node.

Input data set	$MAE(\eta, \eta^*) \downarrow$	$MAE(\phi, \phi^*) \downarrow$	$MAE(\theta, \theta^*) \downarrow$	$\rho(\eta\theta, \eta^*\theta^*) \uparrow$
Social ($s = 16, h = 8$)	0.27 ± 0.11	0.21 ± 0.00	0.24 ± 0.03	0.91 ± 0.03
Balanced ($s = 8, h = 8$)	0.27 ± 0.10	0.22 ± 0.00	0.22 ± 0.03	0.93 ± 0.01
Polarized ($s = 8, h = 16$)	0.27 ± 0.10	0.21 ± 0.00	0.19 ± 0.03	0.96 ± 0.03

- **RQ2.** Do polarized communities detected in real world data sets exhibit typical features associated with echo chambers? (Section 5.2)
- **RQ3.** Can our model be used to provide relevant information to auxiliary predictive tasks, such as predicting activations or individual stances? (Section 5.3)

5.1 Synthetic experiments

In order to answer our first research question, we generate an array of data sets according to the generative model described in Section 3, with different combinations of hyper-parameters.

Reconstruction experiment. We use data sets generated with the procedure defined in Section 3 to test whether the ECD algorithm is able to reliably infer the latent communities. The goal here is first, to present experimental evidence that our algorithm fits the intended purpose. Second, as with any inference procedure, it is necessary to check if a reasonably-sized amount of data is sufficient for a meaningful estimate of the latent communities with the presented algorithm, and under which conditions.

To do so, we consider the same three different settings of parameters s, h from Section 3.3 and visualized in Figure 1. For each parameter setting, we generate 10 data sets composed of a graph and a set of propagations each characterized by a polarity. On each data set, we run our estimation algorithm (initialized with the values for s, h used to generate the data set). From our algorithm, we obtain an estimate for the node-community membership θ^* , the community polarities η^* , and the social interest of nodes ϕ^* . We then measure the absolute error between the original value used to produce the data, and the estimated values obtained by our inference algorithm. Since there is not a natural ordering in the community space, the absolute error is computed as the best result achieved through an exhaustive search in the community indexes.

We present results in Table 2. We observe that our algorithm obtains a low error (between 0.19 and 0.27) for all parameters, and for ϕ and θ in particular. Moreover, the reconstruction of the individual nodes' polarities is very precise ($\rho > 0.9$). The estimate of the node memberships θ gains reliability as the echo-chamber behavior is more apparent: this is quite expected since propagations are more informative when the network is dominated by echo chambers.




Figure 3: Impact of the average number of items per user on the inference of node polarities (measured by Pearson’s Correlation, on the left) and community memberships (measured by MAE, on the right). In both cases, we observe an average number of items per user equal to 10 is sufficient to reach maximum performances, even if node polarities are recovered also with 4 items.

Efficiency Analysis. Then, we investigate the amount of data needed to reliably reconstruct the latent communities. We do so by generating data sets with a growing number of propagations. Specifically, we test from an average number of propagations per user of 1 to 16, with steps of 1. Using these data sets, we perform a grid of experiments by using the same setting described in Section 3.3 ($s = 8$, $h = 16$). The results are shown in Figure 3. We observe that both the estimation of node polarities and θ memberships are affected by different amounts of input data. When looking at the membership reconstruction error, an average of 10 items per user is sufficient for the model to reach its top performance, which then saturates. Individual node polarities, instead, are well estimated even with 4 items per user. This analysis gives a hint at the real-world applicability of our method.

5.2 Echo chamber assessment

In order to answer our second research question, we apply ECD to three real-world data sets extracted from Twitter.

Data sets. Each data set is focused on a different controversial topic:

- Brexit [43] regards the remain-leave discourse before the 2016 UK Referendum to exit the EU. Since this data set does not include retweets, we scrape all the retweets in the period May-July 2016 that contain at least one of the first 100 most used hashtags about Brexit, for a core set of users with at least 5 tweets.
- Referendum [30] is gathered during the Italian constitutional referendum in 2016.
- VaxNoVax [14] comprises polarized discussions related to the vaccine debate in Italy in 2018.

Summary statistics for each data set are reported in Table 3. Each data set exposes different features from the others, relative to the number of users and items, polarity distribution, and cascade size.

In all data sets, we use *follows* to construct the social graph G and *retweets* as propagations $\mathbb{D}$. In cases where the two diverge (a user retweets an item that has not been shared by any of the users they follow in the graph), we insert such missing links in the social graph. Our model also needs as input a polarity p_i for each item:

Table 3: Summary statistics of real-world data sets. The last column refers to the ratio between the number of items with positive $\mathcal{G}^+$ and negative $\mathcal{G}^-$ polarity.

data set	No. of users	No. of items	Cascade size	$\frac{ \mathcal{G}^+ }{ \mathcal{G}^- }$
Brexit	7589	19963	3.6 ± 5.6	1.37
Referendum	2879	40344	5.6 ± 9.9	0.18
VaxNoVax	14315	21312	7.6 ± 28.9	0.80

to compute it for the Brexit and Referendum data sets, we train a supervised text classifier on the labeled subset of tweets provided in the original works. Specifically, we first subselect polarized tweets (AUC ROC 0.78 and 0.75 in 10-fold cross-validation, respectively), and then assign a polarity to each polarized tweet (AUC ROC 0.81 and 0.86). For the VaxNoVax data set, we use the original labeled data to train a text-based classifier that separates tweets from the two ideological sides (F1 0.87) and then sub-select tweets that obtain a classification score larger than 0.75 in absolute value.

Experimental protocol. We next apply the ECD algorithm to these data sets. We chose social prior size $s = 8$ and echo-chamber prior size $h = 16$, since these data sets are collected around polarizing topics, where we expect to find a configuration similar to Figure 1c. We set $K = 8$ as the number of communities; if the data set can be explained by a fewer number of communities, our method simply assigns them a near-zero membership. Thus, in the following, we will consider only non-empty communities. We then adopt the evaluation method proposed by Morini et al. [36]: analyzing each community in terms of its *conductance*—i.e., how closely-knitted is the community with the rest of the graph—and its *purity*—i.e., the ratio of users with the same ideological alignment, measured as the average polarity of the tweets they reshare. Morini et al. identify a low conductance and a high purity as typical properties of echo chambers. ECD training time takes $\sim 2, 6$, and 120 minutes on the three dataset Brexit, Referendum, and VaxNoVax. In light of Table 3, cascade size is a determinant factor for scalability. This is intuitive since we model cascades through all pairs of users $\mathcal{A}(u, v, i)$ that interacts with item i .

Results. We report our results in Figure 4. We observe that all the echo chambers detected by our method (i.e., the communities with a high value of $|\eta|$) indeed display typical echo-chamber traits. Specifically, in the case of Brexit and Referendum (Figures 4a, 4b), we obtain two echo-chamber communities with high purity and low conductance. On VaxNoVax, besides the two echo-chamber communities we also obtain two social communities ($|\eta| \sim 0$). From an empirical analysis, one of these social communities contains all authoritative news sources (e.g., SkyTG24 and AdKronos), while the other contains users who are arguably pro-vax, but without significant pro-vax propagations. Indeed, both social communities have very high conductance, thus missing the segregation exhibited by typical echo chambers.

5.3 Predictive tasks

The latent communities and echo chambers discovered by the ECD model provide valuable information to describe social media users. Such information could be therefore useful for other applicative




Figure 4: Purity-Conductance plots of the communities detected by our method for three data sets: (a) Brexit, (b) Referendum, and (c) VaxNoVax. Each dot represents a community, whose coordinates are its level of conductance (x-axis) and purity (y-axis). The community assignment for a user u is derived from $\arg\max_c \theta_u$, and the colors are associated with the η values inferred by our model. For higher values of $|\eta|$ (echo chambers), we obtain communities well isolated (low conductance) and containing users of the same ideology (high purity), with a clear correspondence in the retweet networks (bottom). Lower absolute values of η correspond to neutral in-between communities which correspond to news media accounts. Force Atlas 2 with gravity= 100 is used in Gephi to define the layout of the network and the inferred community assignment to color the nodes.

predictive tasks. To assess the significance of the produced communities in such tasks, we study two typical prediction problems: *graph-based stance detection* and *next-activation prediction*.

Graph-based stance detection. In the first task, we wish to assign an individual polarity to each node in the network. To evaluate our performance, we manually label a set of ~ 100 users for each data set. Then, we apply our model, excluding their activations from the training set (our model sees their social links). Finally, we assign to each of them the polarity as the weighted average of the polarities of the communities they belong (i.e., $\eta^* \cdot \theta^*$).

We compare its results to the following baselines:

- **1-Hop Average:** given a user u we compute their stance as the average polarity of the propagations of the users that u follows. A similar method was proposed by Barberá [3].
- **node2vec** [23]: we embed the social graph G using the embedding dimension $K = 128$; then, we train a logistic regression using these embeddings. Since this method is supervised, we test it through a leave-one-out cross-validation procedure w.r.t. our set of manually-labeled users; we then report the average.
- **GCN** [29]: we adopt a 2-layer Graph Convolutional Neural Network, using as node features x_u a one-hot encoding of the $|I|$ propagations, i.e. $x_u[i] = 1$ if $u \in \mathbb{D}_i$, and 0 otherwise. Since this method is also supervised, we again adopt leave-one-out to test its performance.

Results are reported in Table 4. We use ROC-AUC as a standard metric to compare the superiority of different models. On all three data sets, our method significantly outperforms the 1-Hop Average baseline. Node2Vec and GCN perform substantially worse than our method, except on VaxNoVax, where all methods achieve good results. In practice, although propagations or the social graph are

valuable sources, ECD is the only model that can efficiently combine information coming from both.

Next-activation prediction. The second predictive task we test is predicting future propagations. In order to evaluate our approach, for each propagation, we split the set of nodes that activated on it into training and test. That is, a fraction (to be determined later) of the activated nodes is not visible during training. Then, we use only the training activations to estimate our model parameters. To approximate the probability of a node u activating on an item i , we consider the maximum probability of propagation from each of the other activated nodes $\mathbb{D}_i$; then, following our model, each probability is computed marginalizing over each community c , thus obtaining

$$P(u|\mathbb{D}_i) = \max \left\{ \sum_c \pi_f(c) \cdot P(\mathcal{J}(u, v, p_i)|c) \mid \forall v \in \mathbb{D}_i \right\}.$$

Using this probability as a prediction score, we evaluate the performance as a binary classification task where, given a pair (u, i) , the model predicts whether user u will activate on item i or not. Hence, we use ROC AUC to measure prediction quality. We apply this procedure for each of the three data sets introduced in the previous section. On each data set, we test different fractions for the train-test split, expressed as the percentage of masked activations during training. Since we treat the next-activation task as a binary classification problem, all pairs (u, i) s.t. $u \notin \mathbb{D}_i$ are attached to the test set as negative instances.

To benchmark the performance of our method, we compare its results with two heuristics. The first one, dubbed *MostPop*, gives higher probability to the most active users:

$$\text{MostPop}(u, \mathbb{D}_i) = \frac{\sum_{j \in \mathbb{I}} \mathbb{1}(u \in \mathbb{D}_j)}{\sum_{j \in \mathbb{I}} \sum_{v \in V} \mathbb{1}(v \in \mathbb{D}_j)}$$

Table 4: ROC-AUC on the stance-detection task for our approach (ECD) and different graph-based supervised and unsupervised baselines (see text). ROC-AUC scores of supervised baselines are measured as average on leave-one-out cross validation.

Supervised	Method	ROC-AUC		
		Brexit	VaxNoVax	Referendum
	ECD	0.98	0.97	0.91
	1-Hop Average	0.47	0.85	0.49
✓	Node2Vec+LR	0.85	1.00	0.75
✓	GCN	0.92	0.94	0.87

while the second, *MostPop**, takes into account the item polarity by assigning higher weight to those users activating on items with similar polarity:

$$\text{MostPop}^*(u, \mathbb{D}_i) = \frac{\sum_{j \in \mathbb{I}, \text{sign}(p_i) = \text{sign}(p_j)} \mathbb{1}(u \in \mathbb{D}_j)}{\sum_{j \in \mathbb{I}, \text{sign}(p_i) = \text{sign}(p_j)} \sum_{v \in V} \mathbb{1}(v \in \mathbb{D}_j)}$$

Results are shown in Figure 5. With a 90%-10% train-test split, we report an AUC ROC of around 0.9 for all three data sets—substantially better than the tested baselines. Moreover, the performance of our method degrades gracefully when the train set size decreases: we do not observe sudden variations in these curves. This result also suggests that the model is not impacted by different sampling choices for the training set.

Reproducibility: our code and data are available at <https://github.com/mminici/Echo-Chamber-Detection.git>

6 CONCLUSIONS AND FUTURE WORK

In this work, we fill the gap between modeling and data analysis approaches when studying echo chambers in social networks. We propose a gradient-based inference algorithm derived from a probabilistic model, which implements realistic assumptions on echo chambers, distinguishing them from other types of communities, and can be used to generate polarized networks and propagations.

Our solution inherits its explainability from this principled generative approach. This approach allows us to formalize the common intuition of a deep entanglement between the observed propagation patterns of polarized contents, and the latent association between users and communities.

The experimental analysis confirms that our algorithm successfully detects echo chambers exhibiting their typical traits of connectivity and opinion homogeneity. Comparisons against state-of-the-art baselines on auxiliary prediction tasks, such as stance detection and next-activation prediction, show the good performance of our algorithm for such tasks in cold-start settings. Also, experiments show that the algorithm is efficient in terms of the number of input propagations needed, and robust with respect to missing data.

Our approach relies on minimal and realistic assumptions that define the perimeter of its effectiveness, and allows for possible extensions. For instance, we consider one specific type of interaction that reflects endorsement, and neglects all the possible nuances in the users’ debates (e.g., *replies* on Twitter could be antagonizing). Furthermore, as presented in Section 5.1, the results of our method improve with the polarization of the input. However, it would be straightforward to extend our model by introducing a form of




Figure 5: Performance of our ECD model and baselines on the task of next-activation prediction, for our three data sets, as a function of the training set size.

hindered propagation for neutral content in non-echo-chamber communities; such an extension would leave our algorithm almost identical.

As in every experimental study, our empirical validation is limited by the available data. For instance, we use only one type of social network, i.e. Twitter. Nonetheless, we hypothesize that our framework would suit also other social media platforms, as long as they allow the existence of a social graph and propagations.

Our work focuses on a given single ideological axis; however, learning the interplay of different axes has been proved successful in the literature [35]: it would be worth devising an extension of our model able to deal with multiple ideological axes.

Another interesting direction for future investigation, would be to model the temporal aspects of propagations that, for simplicity, we do not consider here. On the one hand, propagations naturally happen over time, and their speed could provide further characterization of echo chambers [32]. On the other hand, also polarities do change over time: such an extension could offer improvements on the difficult task of learning opinion dynamics from data [34].

Finally, we consider the polarity of items (i.e., tweets) as part of the input, since they can be easily obtained from natural language processing techniques. However, it would be interesting to integrate such approaches with our model, by considering the items’ polarities as a latent variable, that can be estimated by looking at their propagations as well as their content.

ACKNOWLEDGMENTS

Marco Minici and Giuseppe Manco acknowledge partial support by the EU H2020 ICT48 project “HumanE-AI-Net” under contract #952026. Federico Cinus acknowledges partial support by SoBig-Data++ through the Transnational Access project.

REFERENCES

- [1] Faisal Alatawi, Lu Cheng, Anique Tahir, Mansoor Karami, Bohan Jiang, Tyler Black, and Huan Liu. 2021. A Survey on Echo Chambers on Social Media: Description, Detection and Mitigation. *arXiv preprint arXiv:2112.05084* (2021).
- [2] Abeer Aldayel and Walid Magdy. 2019. Your stance is exposed! analysing possible factors for stance detection on social media. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–20.
- [3] Pablo Barberá. 2015. Birds of the same feather tweet together: Bayesian ideal point estimation using Twitter data. *Political analysis* 23, 1 (2015), 76–91.
- [4] Nicola Barbieri, Francesco Bonchi, and Giuseppe Manco. 2013. Cascade-based community detection. In *Proceedings of the sixth ACM international conference on Web search and data mining*. 33–42.
- [5] Nicola Barbieri, Francesco Bonchi, and Giuseppe Manco. 2013. Influence-Based Network-Oblivious Community Detection. In *2013 IEEE 13th International Conference on Data Mining, Dallas, TX, USA, December 7–10, 2013*. 955–960.
- [6] Nicola Barbieri, Francesco Bonchi, and Giuseppe Manco. 2016. Efficient methods for influence-based network-oblivious community detection. *ACM Transactions on Intelligent Systems and Technology (TIST)* 8, 2 (2016), 1–31.
- [7] Alessandro Bessi. 2016. Personality traits and echo chambers on facebook. *Computers in Human Behavior* 65 (2016), 319–324.
- [8] Christopher M. Bishop. 2006. *Pattern Recognition and Machine Learning*. Springer.
- [9] Simon Bourigault, Sylvain Lamprier, and Patrick Gallinari. 2016. Representation Learning for Information Diffusion through Social Networks: An Embedded Cascade Model. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining (WSDM '16)*. 573–582.
- [10] Fernando H Calderón, Li-Kai Cheng, Ming-Jen Lin, Yen-Hao Huang, and Yi-Shin Chen. 2019. Content-based echo chamber detection on social media platforms. In *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. 597–600.
- [11] Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. 2021. The echo chamber effect on social media. *Proceedings of the National Academy of Sciences* 118, 9 (2021), e2023301118. <https://doi.org/10.1073/pnas.2023301118> [arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.2023301118](https://www.pnas.org/doi/pdf/10.1073/pnas.2023301118)
- [12] Federico Cinus, Marco Minici, Corrado Monti, and Francesco Bonchi. 2022. The effect of people recommenders on echo chambers and polarization. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 16. 90–101.
- [13] Michael Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. 2011. Political polarization on twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 5. 89–96.
- [14] Alessandro Cossard, Gianmarco De Francisci Morales, Kyriaki Kalimeri, Yelena Mejova, Daniela Paolotti, and Michele Starnini. 2020. Falling into the echo chamber: the Italian vaccination debate on Twitter. In *Proceedings of the International AAAI conference on web and social media*, Vol. 14. 130–140.
- [15] Wesley Cota, Silvio C Ferreira, Romualdo Pastor-Satorras, and Michele Starnini. 2019. Quantifying echo chamber effects in information spreading over political communication networks. *EPJ Data Science* 8, 1 (2019), 1–13.
- [16] Gianmarco De Francisci Morales, Corrado Monti, and Michele Starnini. 2021. No echo in the chambers of political interactions on Reddit. *Scientific Reports* 11, 1 (2021), 1–12.
- [17] Michela Del Vicario, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi. 2017. Modeling confirmation bias and polarization. *Scientific reports* 7, 1 (2017), 1–9.
- [18] Michela Del Vicario, Fabiana Zollo, Guido Caldarelli, Antonio Scala, and Walter Quattrociocchi. 2017. Mapping social dynamics on Facebook: The Brexit debate. *Social Networks* 50 (2017), 6–16.
- [19] Siying Du and Steve Gregory. 2016. The Echo Chamber Effect in Twitter: does community polarization increase?. In *International workshop on complex networks and their applications*. Springer, 373–378.
- [20] Francesco Fabbri, Yanhao Wang, Francesco Bonchi, Carlos Castillo, and Michael Mathioudakis. 2022. Rewiring What-to-Watch-Next Recommendations to Reduce Radicalization Pathways. In *Proceedings of the ACM Web Conference 2022*. 2719–2728.
- [21] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Quantifying controversy on social media. *ACM Transactions on Social Computing* 1, 1 (2018), 1–27.
- [22] R Kelly Garrett. 2009. Echo chambers online?: Politically motivated selective exposure among Internet news users. *Journal of computer-mediated communication* 14, 2 (2009), 265–285.
- [23] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. 855–864.
- [24] Stefano Guarino, Noemi Trino, Alessandro Chessa, and Gianni Riotta. 2019. Beyond fact-checking: Network analysis tools for monitoring disinformation in social media. In *International conference on complex networks and their applications*. Springer, 436–447.
- [25] Shahrzad Haddadan, Cristina Menghini, Matteo Riondato, and Eli Upfal. 2021. RePBubLik: Reducing Polarized Bubble Radius with Link Insertions. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (WSDM '21)*. Association for Computing Machinery, New York, NY, USA, 139–147. <https://doi.org/10.1145/3437963.3441825>
- [26] Ling He, Wenzhong Guo, Yuzhong Chen, Kun Guo, and Qifeng Zhuang. 2021. Discovering overlapping communities in dynamic networks based on cascade information diffusion. *IEEE Transactions on Computational Social Systems* (2021).
- [27] Kathleen Hall Jamieson and Joseph N Cappella. 2008. *Echo chamber: Rush Limbaugh and the conservative media establishment*. Oxford University Press.
- [28] Bohan Jiang, Mansoor Karami, Lu Cheng, Tyler Black, and Huan Liu. 2021. Mechanisms and Attributes of Echo Chambers in Social Media. *arXiv preprint arXiv:2106.05401* (2021).
- [29] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*.
- [30] Mirko Lai, Viviana Patti, Giancarlo Ruffo, and Paolo Rosso. 2018. Stance evolution and twitter interactions in an italian political debate. In *International Conference on Applications of Natural Language to Information Systems*. Springer, 15–27.
- [31] Mirko Lai, Marcella Tambuscio, Viviana Patti, Giancarlo Ruffo, and Paolo Rosso. 2019. Stance polarity in political debates: A diachronic perspective of network homophily and conversations on Twitter. *Data & Knowledge Engineering* 124 (2019), 101738.
- [32] Giuseppe Manco, Ettore Ritacco, and Nicola Barbieri. 2021. A Factorization Approach for Survival Analysis on Diffusion Networks. *IEEE Transactions on Knowledge and Data Engineering* 33, 1 (2021), 1–13.
- [33] Yasir Mehmood, Nicola Barbieri, Francesco Bonchi, and Antti Ukkonen. 2013. CSI: Community-Level Social Influence Analysis. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2013, Prague, Czech Republic, September 23–27, 2013, Proceedings, Part II (Lecture Notes in Computer Science)*, Vol. 8189. 48–63.
- [34] Corrado Monti, Gianmarco De Francisci Morales, and Francesco Bonchi. 2020. Learning opinion dynamics from social traces. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 764–773.
- [35] Corrado Monti, Giuseppe Manco, Cigdem Aslay, and Francesco Bonchi. 2021. Learning Ideological Embeddings from Information Cascades. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 1325–1334.
- [36] Virginia Morini, Laura Pollacci, and Giulio Rossetti. 2021. Toward a Standard Approach for Echo Chamber Detection: Reddit Case Study. *Applied Sciences* 11, 12 (2021), 5390.
- [37] Liudmila Prokhorenkova, Alexey Tikhonov, and Nelly Litvak. 2022. When Less is More: Systematic Analysis of Cascade-based Community Detection. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 16, 4 (2022), 1–22.
- [38] Walter Quattrociocchi, Antonio Scala, and Cass R Sunstein. 2016. Echo chambers on Facebook. *Available at SSRN 2795110* (2016).
- [39] Maryam Ramezani, Ali Khodadadi, and Hamid R. Rabiee. 2018. Community Detection Using Diffusion Information. *ACM Trans. Knowl. Discov. Data* 12, 2 (2018).
- [40] Mohammad Sattari and Kamran Zamanifar. 2018. A cascade information diffusion based label propagation algorithm for community detection in dynamic social networks. *Journal of Computational Science* 25 (2018), 122–133.
- [41] Giacomo Villa, Gabriella Pasi, and Marco Viviani. 2021. Echo chamber detection and analysis. *Social Network Analysis and Mining* 11, 1 (2021), 1–17.
- [42] Zheng Zhang, Jun Wan, Mingyang Zhou, Kezhong Lu, Guoliang Chen, and Hao Liao. 2022. Information diffusion-aware likelihood maximization optimization for community detection. *Information Sciences* 602 (2022), 86–105.
- [43] Lixing Zhu, Yulan He, and Deyu Zhou. 2020. Neural opinion dynamics model for the prediction of user-level stance dynamics. *Information Processing & Management* 57, 2 (2020), 102031.