

Users Are Known by the Company They Keep: Topic Models for Viewpoint Discovery in Social Networks

Thibaut Thonet, Guillaume Cabanac, Mohand Boughanem, Karen Pinel-Sauvagnat

IRIT, Université de Toulouse, CNRS

Toulouse, France

{thonet,cabanac,boughanem,sauvagnat}@irit.fr

ABSTRACT

Social media platforms such as weblogs and social networking sites provide Internet users with an unprecedented means to express their opinions and debate on a wide range of issues. Concurrently with their growing importance in public communication, social media platforms may foster *echo chambers* and *filter bubbles*: homophily and content personalization lead users to be increasingly exposed to conforming opinions. There is therefore a need for unbiased systems able to identify and provide access to varied viewpoints. To address this task, we propose in this paper a novel unsupervised topic model, the *Social Network Viewpoint Discovery Model* (SNVDM). Given a specific issue (e.g., U.S. policy) as well as the text and social interactions from the users discussing this issue on a social networking site, SNVDM jointly identifies the issue's topics, the users' viewpoints, and the discourse pertaining to the different topics and viewpoints. In order to overcome the potential sparsity of the social network (i.e., some users interact with only a few other users), we propose an extension to SNVDM based on the *Generalized Pólya Urn* sampling scheme (SNVDM-GPU) to leverage "acquaintances of acquaintances" relationships. We benchmark the different proposed models against three baselines, namely TAM, SN-LDA, and VODUM, on a viewpoint clustering task using two real-world datasets. We thereby provide evidence that our model SNVDM and its extension SNVDM-GPU significantly outperform state-of-the-art baselines, and we show that utilizing social interactions greatly improves viewpoint clustering performance.

KEYWORDS

Viewpoint Discovery, Topic Modeling, Social Networks

1 INTRODUCTION

The advent of social media platforms such as weblogs and social networking sites provided Internet users with an unprecedented means to express their opinions and debate on a wide range of issues, including society and policy. As a consequence, social media is increasingly impacting political life for lay users and policy makers alike, as a 2014 report by the Pew Research Center suggests [47]. Between 2010 and 2014, the share of American voters following political figures on social media has more than doubled. Additionally,

other works noted the influence of social media on key political factors such as democratic participation [53] and campaigning [24].

Concurrently with the growing importance of social media in public and political communication, a phenomenon known as the *echo chambers* [48] or *filter bubble* [34] effect was recently brought to light by scholars. This effect postulates that homophily (i.e., higher chance of bonding between like-minded people) and content personalization lead to an increased exposure to conforming opinions, along with the hiding of contrasting positions. This phenomenon has been for example analyzed in the context of online news consumption [14] and exposure to ideologically diverse content on Facebook [4] and Twitter [12, 16]. While it has been observed that Internet users are not necessarily interested in purely antagonistic content, the study of news consumption patterns suggests that balanced sources, which reflect varied viewpoints, are nonetheless appreciated [17]. Therefore, considering the instrumental role played by social media in politics nowadays, there is a need for unbiased systems that enable users to access different opinions and thus "burst" the filter bubble.

An initial and necessary step in the design of such unbiased systems is an algorithmic approach able to identify users' viewpoints on a specific issue (e.g., U.S. policy) and the topics pertaining to this issue. In order to address this task in a realistic context – i.e., where annotation data is scarce, we propose in the present paper a novel unsupervised topic model, the *Social Network Viewpoint Discovery Model* (SNVDM). Given a specific issue as well as the text and social interactions from the users discussing this issue on a social networking site, SNVDM jointly identifies the issue's topics, the users' viewpoints, and the discourse pertaining to the different topics and viewpoints. In addition to utilizing text data as in traditional topic modeling approaches, SNVDM exploits the homophily phenomenon and thus relies on the users' social interactions to infer their viewpoints. Indeed, users who connect together are more likely to have the same viewpoint, as prior work showed [10, 15].

While a user's social interactions provide rich insight on her opinions, social networks contain sparse components: some users interact with only a few other users. They are however weakly connected to other users: the online acquaintances of their online acquaintances (thereafter, for the sake of simplicity, "the acquaintances of their acquaintances"). Accounting for the virtual links between a user and the acquaintances of her acquaintances, we propose an extension to SNVDM, namely SNVDM-GPU, which is based on the *Generalized Pólya Urn* (GPU) scheme [28]. The GPU scheme has been applied in prior work to embed word similarity in topic models [8, 25, 30, 52]. To the best of our knowledge, our paper introduces the first application of GPU to social networking interactions.

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'17, November 6–10, 2017, Singapore.

© 2017 ACM. ISBN 978-1-4503-4918-5/17/11...\$15.00

DOI: <https://doi.org/10.1145/3132847.3132897>

The contributions of the present article are the following:

- We propose a topic model, the *Social Network Viewpoint Discovery Model* (SNVDM), that leverages text and social networking interactions to estimate users’ viewpoints and discussed topics in social media data;
- Accounting for social network sparsity and weaker connections between users, we extend SNVDM into the SNVDM-GPU model, which leverages the *Generalized Pólya Urn* (GPU) scheme;
- The two proposed models and variations thereof are evaluated against state-of-the-art baselines, namely TAM [35], SN-LDA [45], and VODUM [50], using two real-world datasets harvested from Twitter in [7] to assess the validity of the discovered viewpoints;
- Additionally, we publicly release the code of our models and baselines to foster reproducibility and facilitate future viewpoint discovery research.

Section 2 reviews the prior work related to viewpoint discovery. Then, in Section 3, we define our model SNVDM as well as its extension SNVDM-GPU and we detail their inference procedure. Section 4 describes the experiments conducted to validate these models and how they compete with state-of-the-art baselines. Finally, Section 5 concludes the article and provides future research directions to our work.

2 RELATED WORK

This section introduces the prior work related to viewpoint discovery in social media. We first review political polarization in social media. Then, we describe unsupervised topic models proposed in literature to address viewpoint discovery in formal text and social media. We finally provide a comparison of our models against prior approaches.

2.1 Political Polarization in Social Media

The growing interplay between social media and politics impeded substantial research on the polarization of online discussions. A seminal study reported on the interactions between political bloggers for the 2004 U.S. elections [1]: Democratic (resp. Republican) bloggers were primarily connected to Democratic (resp. Republican) bloggers. A similar trend was stressed from the analysis of political discussions on Twitter [10]. It was however noted that mentions across ideologically antagonistic users are frequent, as opposed to retweets across such users. Other works analyzed the topics and word usage of Twitter communities built on political or social polarizing issues [7, 18]. An interesting related body of work tackled the detection of online controversy, i.e., issues that polarize Internet users. In [11], the authors proposed a weakly-supervised approach to identify controversial Wikipedia articles based on edits and talk pages. Controversy detection was also studied on Twitter using supervised [38] and unsupervised methods [15].

The burst of interest for the study of polarization in social media led researchers to devise supervised machine learning models in order to automatically identify users’ political affiliation and viewpoint (e.g., Democratic or Republican) from their generated content and/or social interactions [2, 9, 13, 37, 43]. The proposed approaches were based on traditional machine learning classifiers

such as Naive Bayes [13, 19], Support Vector Machine [2, 9, 27], Decision Trees [37], or Neural Networks [22, 43]. While these classifiers usually reach high accuracy, they need to be fed labeled examples, which may be difficult to gather.

2.2 Unsupervised Viewpoint Discovery

A common method to overcome the lack of annotation data in text mining tasks is to explore unsupervised topic modeling approaches, e.g., based on Latent Dirichlet Allocation (LDA) [6]. Researchers therefore proposed Bayesian and topic models to discover viewpoints in formal documents such as essays [35, 36, 50] and legislation bills [33], or in social media data such as forum posts [40, 42] and tweets [5, 23, 44].

2.2.1 Application to Formal Text. In [35], the authors defined the Topic-Aspect Model (TAM) that jointly models topics and so-called aspects, which can be interpreted as viewpoints in the context of viewpoint discovery. A contrastive viewpoint summarization framework based on TAM was then described in [36], aiming to find phrases that best reflect the different viewpoints pertaining to a set of topics. Another topic model, the Viewpoint and Opinion Discovery Unification Model (VODUM), was proposed to identify viewpoints by leveraging parts of speech [50]. The proposed scheme helped discriminate between topic words and opinion words. The authors of [33] studied a slightly different yet related problem: the prediction and analysis of legislators’ ideal point given their votes on bills. Ideal point is a measure used in political science to assign policy makers on a one-dimension ideological spectrum ranging from liberal to conservative. The topic model devised in [33] integrates regression techniques to estimate real-valued ideal points and is able to extract ideological discourse vocabulary.

2.2.2 Models for Social Media Data. Topic modeling for viewpoint discovery has been applied as well to social media and particularly to forums, which naturally enable Internet users to debate. A more comprehensive literature review on the subject is provided in [39]. The topic models described in [40–42] utilize the threaded nature of forum posts. The authors assume that the viewpoint conveyed in a post depends on the post’s positive or negative polarity and the viewpoint of the parent post – thereby denoting agreement or disagreement between users. However, the proposed approach could not be applied to social networking sites such as Twitter, where threaded interactions are scarce.

Viewpoint discovery was also studied on Twitter. The authors of [44] designed a time-aware topic model to summarize contrastive opinions expressed in multilingual social text streams. However, contrast was based only on sentiment (positive, negative, or neutral) rather than on viewpoint – which goes beyond mere sentiment polarity and encompasses partisanship, e.g., pro-Palestine and pro-Israel, or Democratic and Republican. Conversely, the work described in [23] focused on detecting positions and political issues on Twitter. It proposed a topic model that uses politically affiliated users to estimate non-affiliated users’ position; yet it did not leverage interactions between users. On the other hand, a Bayesian ideal point model for Twitter users was devised in [5] but it was only based on follow interactions and it did not utilize text data.

Table 1: Comparison of our models SNVDM and SNVDM-GPU against related work approaches.

Reference	Application to social media	Use of social networking interactions	Discovery of viewpoints and related discourse	Unified approach
TAM [35, 36]; VODUM [50]; [33] [40]; [42]; [41]; [23] [44]	✗	✗	✓	✓
[5]; SN-LDA [45] [49]	✓	✓	✗	✓
SNVDM(GPU) [this paper]	✓	✓	✓	✓

Although addressing a different task, the topic model described in [45], namely the Social Network Latent Dirichlet Allocation (SN-LDA), combined both text and social networking interactions to detect communities on Twitter. The difference between traditional community detection and viewpoint discovery is that the former usually focuses on topical communities (e.g., communities discussing football or programming) while the latter studies the different viewpoints expressed on a common set of topics (e.g., the Democratic and Republican positions on abortion). Nonetheless, the homophily phenomenon is featured in both cases; the method used in [45] is therefore relevant to our viewpoint discovery goal. Community detection techniques were otherwise applied in conjunction with LDA to analyze advocating and conspiring discussions about vaccines and their underlying topics on Twitter [49]. However, the proposed approach did not unify the modeling of text and social interactions. Moreover, the community detection techniques used are non-parametric and thus lead to the discovery of many small communities which may be difficult to interpret and do not necessarily reflect different viewpoints.

We summarize in Table 1 the characteristics of our models SNVDM and SNVDM-GPU compared against related work from Section 2.2. Our topic models SNVDM and SNVDM-GPU use social networking interactions and text in a unified fashion to discover viewpoints, topics and related word usage in social media data. These characteristics are further detailed in Section 3.

3 METHOD

In this section, we describe the topic models proposed to discover viewpoints and topics in social networking data. We first establish preliminary definitions and notations. Secondly, we define our topic model, the Social Network Viewpoint Discovery Model (SNVDM), and give its inference procedure. Thirdly, we provide background on the Pólya Urn sampling scheme and we detail our extension to SNVDM, namely SNVDM-GPU, based on Generalized Pólya Urn.

3.1 Preliminaries

As the focus of this work is the modeling of *viewpoints* and *topics* in *social networks*, we provide definitions for related terms and phrases. Following the definitions of [50], a topic is “one of the subjects discussed in a document collection” and a viewpoint is “the standpoint of one or several authors on a set of topics.” We use the phrase “social network” to denote the directed graph of online social interactions (the edges) between users (the nodes) on a social media platform (e.g., Twitter). We thereafter simply use *social interactions* or *interactions* to denote online social interactions. Given the directed nature of a social network, a user may take part in two

Table 2: Definition of the notations used in the paper.

Symbol	Definition
$U, D_u, O_u, N_{ud}, I_{ud}$	Number of users, number of documents from u , number of outgoing interactions for u , number of tokens for u and d , number of incoming interactions for u and d , resp.
T, V, W	Number of topics, number of viewpoints, and number of words in the vocabulary, resp.
$w_{udn}, z_{udn}, l_{udn}, x_{udn}$	The n -th word token and its topic, level, and route assignments, resp., for document d from user u .
s_{udi}	The sender of incoming interaction i on document d from user u .
v_{ud}	The viewpoint assignment for document d from user u .
r_{uo}, v'_{uo}	The recipient and the viewpoint assignment, resp., for outgoing interaction o of user u .
$\phi_{00}, \phi_{01}, \phi_{10}, \phi_{11}, \beta$	$1 \times W$ matrix of background word distributions, $V \times W$ matrix of viewpoint word distributions, $T \times W$ matrix of topic word distributions, $V \times T \times W$ matrix of viewpoint-topic word distributions, and their concentration parameter, resp.
$\psi_0, \psi_1, \gamma_0, \gamma_1$	1×2 matrix of general route distribution, $T \times 2$ matrix of topic-specific route distributions, and their shape parameters, resp.
ξ, μ	$V \times U$ matrix of viewpoint-specific interacting user distributions, and their concentration parameter, resp.
$\sigma, \delta_0, \delta_1$	$V \times U$ matrix of user-specific level distributions, and their shape parameters, resp.
θ, α	$U \times T$ matrix of user-specific topic distributions and their concentration parameter, resp.
π, η	$U \times V$ matrix of user-specific viewpoint distributions and their concentration parameter, resp.
$n_{00w}, n_{01vw}, n_{10zw}, n_{11vzw}$	Number of background instances of word w , number of viewpoint instances of word w assigned to v , number of topic instances of word w assigned to z , number of viewpoint/topic instances of word w assigned to v and z , resp.
n_{0x}, n_{1zx}	Number of words assigned to level 0 and route x and number of words assigned to level 1, topic z , and route x , resp.
n_{vu}	Number of interactions with interacting user u assigned to viewpoint v .
n_{ul}	Number of words from user u assigned to level l .
n_{uz}	Number of words from user u assigned to topic z .
n_{uv}	Number of documents and outgoing interactions from user u assigned to viewpoint v .

distinct types of interactions: *incoming interactions* and *outgoing interactions*. A user u takes part in an incoming interaction with another user u' if u' initiated the interaction (e.g., u' retweeted or replied to u 's tweet). We then call u and u' the *recipient* and the *sender* of this interaction, respectively. Conversely, a user u takes part in an outgoing interaction with another user u' if u initiated the interaction (e.g., u retweeted or replied to u' 's tweet). Similarly, we call u and u' the *sender* and the *recipient* of this interaction, respectively. The notations for the variables used throughout our paper are provided in Table 2. Additionally, we use boldface symbols to represent a multiset or a vector (e.g., $\mathbf{v} = \{\{v_{ud}\}_{d=1}^{D_u}\}_{u=1}^U$). Count variables, i.e., variables expressing the number of specific assignments, are defined by n and subscripts (e.g., n_{uz} is the number of word tokens from user u assigned to topic z). A subscript “ \cdot ” means that the counting is marginalized over the corresponding variable (e.g., $n_{\cdot v} = \sum_{u=1}^U n_{vu}$). A superscript $-(y)$ on a count variable implies that the count variable excludes y from the counting (e.g., $n_{uz}^{-(udn)}$ is the number of word tokens from user u assigned to topic z , excluding word n from document d of user u).

3.2 Social Network Viewpoint Discovery Model

3.2.1 Model Definition. Before providing a more formal description of our proposed Social Network Viewpoint Discovery Model

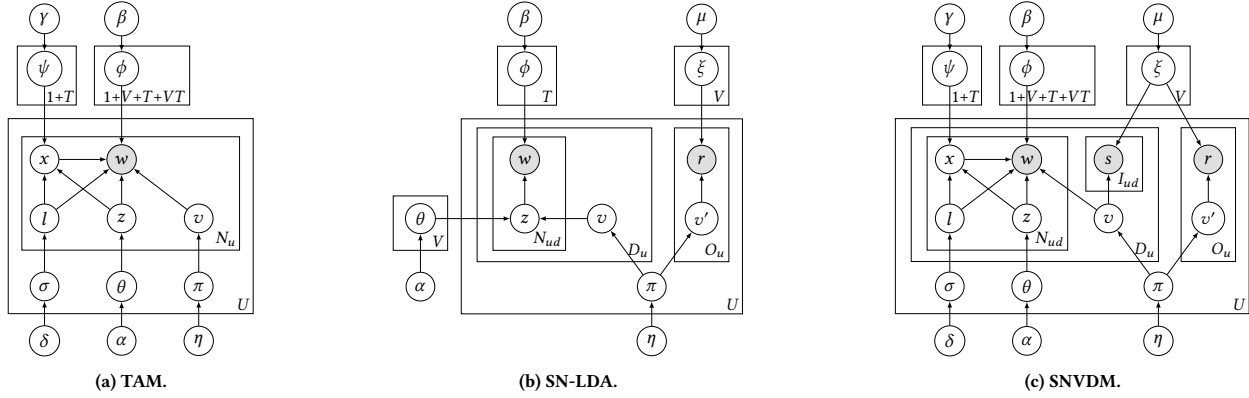


Figure 1: Graphical models for TAM (a), SN-LDA (b), and our proposed model SNVDM (c). The original variable names for TAM and SN-LDA were slightly modified from [35] and [45], respectively, to facilitate their comparison against SNVDM.

(SNVDM), we introduce its key characteristics. In traditional topic models such as LDA [6], words are assumed to be drawn from topic-specific word distributions. However, in SNVDM, we aim to model viewpoints in addition to topics. A naive approach to account for this would be to simply draw all words from viewpoint/topic-specific distributions. However, all words need not depend on both a viewpoint and a topic. A word can for example be a background word (e.g., *say*): it does not reflect any viewpoint or topic; it can otherwise be a topic word (e.g., *energy*), and depend only on a topic – and not on any viewpoint. To implement this observation, we follow the idea proposed in TAM [35]. The authors introduced Bernoulli latent variables, namely *levels* and *routes*, to account for the possible dependency of word tokens on topic and viewpoint, respectively. This yields four types of words: *background words* (level = 0, route = 0), *viewpoint words* (level = 0, route = 1), *topic words* (level = 1, route = 0), and *viewpoint-topic words* (level = 1, route = 1). These different types of words are very useful to analyze the viewpoint-specific and neutral word usage on a given topic.

Although word usage partially helps identify the expressed viewpoint in a document, this may not be enough as viewpoints can be expressed in a subtle way [50]. It is even more problematic in social media texts (e.g., tweets), which are often short, noisy, and contain numerous abbreviations. However, social media platforms enable social interactions between users. The underlying social network then provides precious insights on users’ viewpoint according to the *homophily* principle: like-minded users tend to connect together. This idea has been already used in prior topic models designed for community detection, e.g. in SN-LDA [45]. However, SN-LDA only utilized interactions and latent topics to identify communities – word usage did not directly depend on communities. Moreover, in SN-LDA, a given user is only aware of her outgoing interactions, i.e., the interactions she initiated: her incoming interactions are ignored. This may matter for a user that only or mostly receives incoming interactions (e.g., a user that is often retweeted but scarcely retweets). Under the modeling of SN-LDA, this user’s community will be difficult to identify as she has very few outgoing interactions. Therefore, in our model SNVDM, we propose to use both *incoming and outgoing interactions*. More precisely, in SNVDM, incoming interactions are used at document level to denote the retrospective

influence of interacting users (e.g., retweeting users) on the viewpoint expressed in a document. Note that, in this paper, we adopt a static view of the collection, we therefore assume in SNVDM’s generative process that incoming interactions on a document occur immediately after the document is written. In other words, we do not account for the temporality of incoming interactions.

We now give a more formal definition of our proposed model SNVDM by providing its graphical model (Figure 1c) compared against that of TAM (Figure 1a) and that of SN-LDA (Figure 1b). We describe as well SNVDM’s generative story:

- (1) Draw general and topic-specific route distributions ψ_0 and $\psi_{1z} \sim \text{Beta}(\gamma_0, \gamma_1)$ for $z = 1, \dots, T$;
- (2) Draw background, viewpoint, topic, and viewpoint-topic word distributions $\phi_{00}, \phi_{01v}, \phi_{10z},$ and $\phi_{11vz} \sim \text{Dirichlet}_W(\beta)$ for $v = 1, \dots, V$ and $z = 1, \dots, T$;
- (3) Draw viewpoint-specific interacting user distributions $\xi_v \sim \text{Dirichlet}_U(\mu \frac{1}{V})$ for $v = 1, \dots, V$;
- (4) For each user $u = 1, \dots, U$:
 - (a) Draw viewpoint distribution $\pi_u \sim \text{Dirichlet}_V(\eta \frac{1}{V})$;
 - (b) Draw topic distribution $\theta_u \sim \text{Dirichlet}_T(\alpha \frac{1}{T})$;
 - (c) Draw level distribution $\sigma_u \sim \text{Beta}(\delta_0, \delta_1)$;
- (5) For each document $d = 1, \dots, D_u$ from $u = 1, \dots, U$:
 - (a) Draw a document viewpoint $v_{ud} \sim \text{Discrete}(\pi_u)$;
 - (b) For each token $n = 1, \dots, N_{ud}$:
 - (i) Draw a topic $z_{udn} \sim \text{Discrete}(\theta_u)$;
 - (ii) Draw a level $l_{udn} \sim \text{Discrete}(\sigma_u)$;
 - (iii) If $l_{udn} = 0$, draw a general route $x_{udn} \sim \text{Discrete}(\psi_0)$;
 - Else if $l_{udn} = 1$, draw a topic-specific route $x_{udn} \sim \text{Discrete}(\psi_{1z_{udn}})$;
 - (iv) If $l_{udn} = 0$ and $x_{udn} = 0$, draw a background word $w_{udn} \sim \text{Discrete}(\phi_{00})$;
 - Else if $l_{udn} = 0$ and $x_{udn} = 1$, draw a viewpoint word $w_{udn} \sim \text{Discrete}(\phi_{01v_{ud}})$;
 - Else if $l_{udn} = 1$ and $x_{udn} = 0$, draw a topic word $w_{udn} \sim \text{Discrete}(\phi_{10z_{udn}})$;
 - Else if $l_{udn} = 1$ and $x_{udn} = 1$, draw a viewpoint-topic word $w_{udn} \sim \text{Discrete}(\phi_{11v_{ud}z_{udn}})$;
 - (c) For each incoming interaction $i = 1, \dots, I_{ud}$, draw sender $s_{udi} \sim \text{Discrete}(\xi_{v_{ud}})$;
- (6) For each outgoing interaction $o = 1, \dots, O_u$ from $u = 1, \dots, U$:
 - (a) Draw an outgoing interaction viewpoint $v'_{uo} \sim \text{Discrete}(\pi_u)$;
 - (b) Draw a recipient $r_{uo} \sim \text{Discrete}(\xi_{v'_{uo}})$.

In the next section, we detail the posterior inference procedure for SNVDM based on collapsed Gibbs sampling.

3.2.2 *Posterior Inference.* As for other probabilistic topic models, the exact posterior inference for SNVDM is intractable. We

thus rely on *collapsed Gibbs sampling* [26] to estimate the model distributions. In SNVDM, the latent variables of interest for a collapsed Gibbs sampler are topics \mathbf{z} , levels \mathbf{l} , routes \mathbf{x} , document viewpoints \mathbf{v} and outgoing interaction viewpoints \mathbf{v}' ; the observed variables are words \mathbf{w} , incoming interaction senders \mathbf{s} , and outgoing interaction recipients \mathbf{r} . The collapsed Gibbs sampler for SNVDM then successively samples the aforementioned latent variables from their conditional posterior probability, which we provide in the remainder of this section. Mathematical derivations are omitted due to space limitations; for a detailed description of Gibbs sampling derivations for topic models, the interested reader can refer to [20].

Sampling topics, levels, and routes. Similarly to TAM [35], assignments z_{udn} , l_{udn} and x_{udn} can be jointly sampled from the following joint conditional posterior probability:

$$p(z_{udn} = z, l_{udn} = l, x_{udn} = x | v_{ud} = v, \text{rest}) \propto \frac{n_{uz}^{-(udn)} + \alpha \frac{1}{T}}{n_{u\cdot}^{-(udn)} + \alpha} \quad (1)$$

$$\cdot \frac{n_{ul}^{-(udn)} + \delta_l}{n_{u\cdot}^{-(udn)} + \delta_0 + \delta_1} \cdot p(x_{udn} = x | \text{rest}) \cdot p(w_{udn} = w | \text{rest})$$

$$\text{where } p(x_{udn} = x | \text{rest}) = \begin{cases} \frac{n_{0x}^{-(udn)} + \gamma_x}{n_{0\cdot}^{-(udn)} + \gamma_0 + \gamma_1} & \text{if } l = 0, \\ \frac{n_{1zx}^{-(udn)} + \gamma_x}{n_{1z\cdot}^{-(udn)} + \gamma_0 + \gamma_1} & \text{if } l = 1; \end{cases}$$

$$\text{and } p(w_{udn} = w | \text{rest}) = \begin{cases} \frac{n_{0w}^{-(udn)} + \beta}{n_{00\cdot}^{-(udn)} + \beta W} & \text{if } l = 0, x = 0, \\ \frac{n_{01w}^{-(udn)} + \beta}{n_{01\cdot}^{-(udn)} + \beta W} & \text{if } l = 0, x = 1, \\ \frac{n_{10w}^{-(udn)} + \beta}{n_{10z\cdot}^{-(udn)} + \beta W} & \text{if } l = 1, x = 0, \\ \frac{n_{11w}^{-(udn)} + \beta}{n_{11vz\cdot}^{-(udn)} + \beta W} & \text{if } l = 1, x = 1. \end{cases}$$

Sampling document viewpoints. The conditional posterior probability for the viewpoint v_{ud} assigned to user u and her document d is given by:

$$p(v_{ud} = v | \text{rest}) \propto \frac{\prod_{w=1}^W \prod_{a=0}^{n_{01vzw}^{(ud)} - 1} (n_{01vzw}^{-(ud)} + a + \beta)}{\prod_{b=0}^{n_{01v\cdot}^{(ud)} - 1} (n_{01v\cdot}^{-(ud)} + b + \beta W)} \quad (2)$$

$$\cdot \prod_{z=1}^T \frac{\prod_{w=1}^W \prod_{a=0}^{n_{11vzw}^{(ud)} - 1} (n_{11vzw}^{-(ud)} + a + \beta)}{\prod_{b=0}^{n_{11vz\cdot}^{(ud)} - 1} (n_{11vz\cdot}^{-(ud)} + b + \beta W)}$$

$$\cdot \frac{n_{uv}^{-(ud)} + \eta \frac{1}{V}}{n_{u\cdot}^{-(ud)} + \eta} \cdot \frac{\prod_{u'=1}^U \prod_{a=0}^{n_{vu'}^{(ud)} - 1} (n_{vu'}^{-(ud)} + a + \mu \frac{1}{U})}{\prod_{b=0}^{n_{v\cdot}^{(ud)} - 1} (n_{v\cdot}^{-(ud)} + b + \mu)}$$

Sampling outgoing interaction viewpoints. For each outgoing interaction o of user u , viewpoint v'_{uo} is sampled as follows:

$$p(v'_{uo} = v | r_{uo} = u', \text{rest}) \propto \frac{n_{uv}^{-(uo)} + \eta \frac{1}{V}}{n_{u\cdot}^{-(uo)} + \eta} \cdot \frac{n_{vu'}^{-(uo)} + \mu \frac{1}{U}}{n_{v\cdot}^{-(uo)} + \mu} \quad (3)$$

Sampling hyperparameters. Prior work showed that sampling hyperparameters – as opposed to setting them to fixed values – impacts the model’s performance [51]. Therefore, we apply a full Bayesian treatment to our model by sampling the hyperparameters α , γ , δ , η , and μ from Gamma(1, 1) hyperpriors, following the auxiliary variable sampling technique as in [32]. We only fix $\beta = 0.01$ as

its sampling is more costly – since it depends on the usually large vocabulary size.

3.2.3 Discussion. Although SNVDM utilizes direct outgoing and incoming interactions between a user and her acquaintances, it does not leverage farther users in the network such as *acquaintances of acquaintances*, who can nevertheless bring precious insights on the user’s viewpoint. For example, consider a user u who exclusively interacts with a second user u' , who in turn interacts with a large number of users. Using only u ’s sparse direct interactions will not provide much information about her, whereas knowing that there exists a weaker link between u and u' ’s acquaintances can help describe u more accurately. We augment our model in the next section to implement this idea.

3.3 Generalized Pólya Urn Extension

We propose to extend SNVDM in order to account for virtual links between users and their “acquaintances of acquaintances”. We first introduce the simple Pólya urn scheme, then describe the extension of SNVDM based on generalized Pólya urn.

3.3.1 Simple Pólya Urn. Collapsed Gibbs sampling for Dirichlet-multinomial topic models such as LDA or SNVDM can be interpreted as an urn metaphor called the *simple Pólya urn* (SPU) scheme. SPU assumes that colored balls are successively drawn from an urn; if a ball of color c is drawn, then this ball is put back into the urn and one additional ball of color c is added. This over-replacement scheme enforces a property known as “the rich get richer”: the more balls of color c are drawn, the more likely balls of color c will be drawn in the future.

We illustrate the SPU metaphor assuming outgoing interactions are balls, the recipient user of the interaction is their color, and each urn is associated with a viewpoint. Initially, we assume that each urn contains $\mu \frac{1}{U}$ balls of each color. Now suppose that the interaction o for user u is the last ball we want to sample from urn v – i.e., we have already drawn and observed all other interactions/balls. Then the probability of obtaining color u' for ball (uo) is given by:

$$p(r_{uo} = u' | v'_{uo} = v, \text{rest}) = \frac{n_{vu'}^{-(uo)} + \mu \frac{1}{U}}{n_{v\cdot}^{-(uo)} + \mu} \quad (4)$$

where $n_{vu'}^{-(uo)}$ can be interpreted as the number of balls of color u' added in urn v , and $n_{v\cdot}^{-(uo)}$ as the total number of balls added in urn v , both counts excluding ball (uo) and the μ initial balls. Note that the rightmost term in Eq. (3) is the same as the one in Eq. (4) obtained using the SPU assumption, illustrating the equivalency between Dirichlet-multinomial and SPU models.

3.3.2 Generalized Pólya Urn. The *generalized Pólya urn* (GPU) scheme [28] extends SPU by altering the replacement rule: after a ball of color c is drawn, the ball is put back into the urn and a certain number of balls of colors *related* to c are added along. In prior topic models using GPU, balls are tokens and colors are word types [8, 25, 30, 52]; the “relatedness” of colors is therefore based on word similarity such as PMI [8, 52] or computed from word embeddings [25].

In our model, however, we utilize GPU to account for weaker links between users in a social network. As illustrated in 3.3.1, we consider interactions as balls and users as colors. Leveraging

the network structure, we assume that the users most related to a user u are the users who interact the most with u . Formally, we define the number of incoming interactions on u where u' is the sender as $i_{uu'} = \sum_{d=1}^{D_u} \sum_{i=1}^{I_{ud}} \mathbb{I}(s_{udi} = u')$, and the number of outgoing interactions from u where u' is the recipient as $o_{uu'} = \sum_{o=1}^{O_u} \mathbb{I}(r_{uo} = u')$. \mathbb{I} denotes the indicator function: $\mathbb{I}(\text{true}) = 1$ and $\mathbb{I}(\text{false}) = 0$. The set of the τ most interacting acquaintances of u (if such number of acquaintances exists for u) is then $\mathcal{R}_{u\tau}$ with $|\mathcal{R}_{u\tau}| = \min(\tau, \sum_{u'=1}^U \mathbb{I}(i_{uu'} + o_{uu'} > 0))$, and such that, for all $u' \in \mathcal{R}_{u\tau}$, $i_{uu'} + o_{uu'} > 0$ (i.e., u' is an acquaintance of u) and, for all $u'' \notin \mathcal{R}_{u\tau}$, $i_{uu'} + o_{uu'} \geq i_{uu''} + o_{uu''}$ (i.e., u' interacts more with u than u'' does). We now define a $U \times U$ real-valued matrix \mathbb{A} known as *schema* or *addition matrix* [28] which, given a withdrawn ball of color u , expresses the number of balls $\mathbb{A}_{uu'}$ for each color $u' = 1, \dots, U$ that is added along to the urn:

$$\mathbb{A}_{uu'} = \begin{cases} 1 & \text{if } u = u', \\ \lambda & \text{if } u \neq u' \text{ and } u' \in \mathcal{R}_{u\tau}, \\ 0 & \text{otherwise;} \end{cases} \quad (5)$$

where λ is a real-valued parameter in $[0, 1]$ denoting the number of balls added for colors related to u (i.e., in $\mathcal{R}_{u\tau}$).

Contrarily to the SPU scheme, balls sampled according to GPU are non-exchangeable, implying that the sampling order for a sequence of balls will impact the probability of the observed colors. Following [30], we approximate the exact conditional posterior probability used in Gibbs sampling by assuming that the ball or sequence of balls of interest is sampled as if we already sampled and observed all other balls in the urn. We thus ignore the current sampling's implications on the drawing of subsequent balls.

The application of the GPU scheme to SNVDM under this approximation leads to a slight modification of the collapsed Gibbs sampler described in Section 3.2.2. In the sampling of document viewpoint v_{ud} , the rightmost term in Eq. (2) under the GPU assumption can be shown to become:

$$\prod_{i=1}^{I_{ud}} \frac{\sum_{u''=1}^U \mathbb{A}_{u''s_{udi}} n_{vu''}^{-(ud)} + \sum_{j=1}^{i-1} \mathbb{A}_{s_{udj}s_{udi}} + \mu \frac{1}{U}}{\sum_{u''=1}^U \mathbb{A}_{u'' \cdot} n_{vu''}^{-(ud)} + \sum_{j=1}^{i-1} \mathbb{A}_{s_{udj} \cdot} + \mu}. \quad (6)$$

Similarly, the rightmost term in Eq. (3) for the sampling of outgoing interaction viewpoint v'_{uo} is changed into:

$$\frac{\sum_{u''=1}^U \mathbb{A}_{u''u'} n_{vu''}^{-(uo)} + \mu \frac{1}{U}}{\sum_{u''=1}^U \mathbb{A}_{u'' \cdot} n_{vu''}^{-(uo)} + \mu}. \quad (7)$$

4 EXPERIMENTS

In this section, we detail the experiments conducted on our models SNVDM and SNVDM-GPU, and their comparison against state-of-the-art baselines. We aim to validate the following hypotheses:

- **(H1)** The viewpoint clustering performance of the proposed models **(a)** outperforms that of state-of-the-art baselines; it is improved by **(b)** exploiting both incoming and outgoing interactions, and **(c)** implementing the generalized Pólya urn scheme.
- **(H2)** Generalized Pólya urn scheme increases models' robustness to social network sparsity, i.e., GPU-based models are less affected by limited number of social interactions.

- **(H3)** The proposed models are comparably efficient to state-of-the-art baselines, i.e., their execution times have the same order of magnitude.
- **(H4)** The viewpoints and topics discovered by our models are coherent.

In the remainder of this section, we first introduce the experimental setup. Secondly, we detail the viewpoint clustering performance of our models and baselines. Thirdly, we discuss the efficiency of the evaluated models. Finally, we illustrate the discovered viewpoints and topics by providing their most representative words.

4.1 Experimental Setup

4.1.1 Baselines and Evaluated Models. In order to test hypotheses (H1)-(H4), we compare prior topic models able to discover viewpoints against variations of our proposed models:

- **TAM** [35] was initially designed to discover topics and aspects – which represent viewpoints in our context. It was integrated into a viewpoint summarization framework in [36]. It does not account for interactions between users.
- **SN-LDA** [45] aims to jointly discover topics and communities – which we interpret as viewpoints – in a social network where users are associated with text data. The only interactions used are outgoing interactions.
- **VODUM** [50] models viewpoints and topics. It leverages part of speech to better discriminate opinion words (i.e., viewpoint/topic-specific words) from topic words. It does not use interactions between users.
- **SNVDM-WII** is a degenerate variation of our model SNVDM (described in Section 3.2) without incoming interactions. The only interactions used are outgoing interactions.
- **SNVDM** is the model we propose, as described in Section 3.2. It leverages incoming and outgoing interactions.
- **SNVDM-GPU** ($\tau = 10$) and **SNVDM-GPU** ($\tau = \infty$) extend SNVDM by integrating the generalized Pólya urn scheme into the inference procedure, as described in Section 3.3. The value of τ defines the maximum number of most interacting acquaintances used as related users. In SNVDM-GPU ($\tau = 10$), only the 10 most interacting acquaintances of a given user are used as her related users, whereas for SNVDM-GPU ($\tau = \infty$) all acquaintances are used as related users.

We release the Java code for our models and baselines at <https://github.com/tthonet/SNVDM>.

4.1.2 Datasets. There exists only few annotated datasets to evaluate models for viewpoint discovery in social networks. In 2016, the International Workshop on Semantic Evaluation (SemEval) introduced a benchmark for stance detection on Twitter [31]. However, the task proposed by SemEval was focused on tweet-level stance detection, whereas in this work we are interested in discovering user-level viewpoints. The methods and data relevant to these two problems differ as only the latter enables the exploitation of social interactions. Therefore, we chose to validate our approaches on two Twitter datasets¹ introduced in [7], which we refer the reader to for additional details. The first dataset, thereafter denoted as *Indyref*,

¹<http://dx.doi.org/10.6084/m9.figshare.1430449>

Table 3: Statistics of the datasets used in the experiments.

Dataset	#Users		#Tweets	#Tokens	Vocabulary	#Interactions
	Yes/Dem.	No/Rep.				
Indyref	589	575	270,075	2,043,204	38,942	696,654
Midterms	767	778	113,545	975,199	25,312	241,741

contains tweets about the 2014 Scottish Independence Referendum posted between 11/08/2014 and 20/10/2014. The two represented viewpoints are *Yes* and *No*, respectively expressing support and opposition to Scottish independence. The second dataset, thereafter denoted as *Midterms*, is constituted of the tweets written by the policy makers who were active during the 2014 U.S. Midterm Election. Note that, in order to extend the *Midterms* dataset, we used users’ full Twitter timelines (until 21/11/2014) instead of restricting to the tweets posted during the Midterm Election timespan. Only *Democratic* and *Republican* users appear in the dataset. For both datasets, users’ groundtruth viewpoints were originally obtained in [7] either from self-description in user profiles or from official and unofficial lists. Similarly to prior work on political polarization [10, 27], the Twitter social interactions that we considered to define the users’ social network are *retweets* and *replies*.

An issue we had to address with the datasets was the fact that they contained a large amount of tweets that did not concern the studied political subject (e.g., tweets about leisure or daily life). To remove these irrelevant tweets, we discarded all the tweets with no interactions from the users in the dataset, i.e., tweets with neither reply nor retweet from the dataset users. Although this process removed some relevant tweets, we observed that it provided us with considerably more focused datasets. Moreover, we only kept unique tweets in the datasets (removing the retweeted duplicates) and assigned them to their original authors. After the denoising step, users left with no tweets were discarded (32 for *Indyref* and 232 for *Midterms*). Note that this process is widely applicable to any dataset with social interactions and it does not require supervision. However, the number of interactions per tweet ratio is affected; therefore we further investigate in Section 4.2.2 the impact of the network sparsity on the models’ and baselines’ performance.

We then performed the following preprocessing steps on the datasets using *Lingpipe*² and *TweetNLP*³. We applied the part-of-speech tagger provided by *TweetNLP* to the tweets, as it is required by the baseline *VODUM* to discriminate opinion words from topic words. Following [50], nouns were used as topic words and verbs, adverbs, adjectives, and prepositions (and additionally hashtags) were used as opinion words in *VODUM*. Tokens not matching these parts of speech were discarded from the datasets used by all models – the rationale is that, for a fair comparison, we wanted to test all models on exactly the same text data. We then removed stopwords, user mentions, URLs and word tokens that appeared only once. Following [46], we did not perform stemming. Users with no tweets (or with only token-less tweets) were discarded. The statistics of the preprocessed datasets are detailed in Table 3. Reported interactions include both incoming and outgoing interactions.

²<http://alias-i.com/lingpipe/>

³<http://www.cs.cmu.edu/~ark/TweetNLP/>

4.1.3 Parameter Setting. For both datasets *Indyref* and *Midterms*, we set the number of viewpoints (for *VODUM* and all *SNVDM*-based models), aspects (for *TAM*), and communities (for *SN-LDA*) to 2. For all baselines and models, hyperparameters were initialized to 1 and sampled as described in Section 3.2.2, except for word distributions parameters (β in *SNVDM*-based models) which are set to 0.01. The parameter λ used in *GPU* models was fixed to 0.5. We observed (not shown due to space limitations) that this yielded better performance than setting $\lambda = 1$ (i.e., a uniform treatment of direct acquaintances and acquaintances of acquaintances). The Gibbs sampling inference for all baselines and models is performed on 5 Markov chains of 1,000 iterations each, with 500 iterations for burn-in. After burn-in, one sample was collected every 50 iterations, and models’ distributions were eventually estimated based on the 10 collected samples.

4.2 Viewpoint Clustering

In this section, we describe the results obtained by baselines and proposed models on the clustering of users’ viewpoints, evaluated in terms of *Purity* and *Normalized Mutual Information* (NMI) [29]. *Purity* measures the proportion of users who are assigned to the correct groundtruth class. NMI is an information theoretic clustering metric based on Mutual Information and Entropy. We also tested the *BCubed F* measure [3] but we found almost perfect Spearman’s rank correlations between *Purity* and *BCubed F* ($\rho = 0.998$) and between NMI and *BCubed F* ($\rho = 0.999$) – computed from the 280 measurements obtained in Section 4.2.1. Therefore, for the sake of brevity, we only report the clustering quality results in terms of *Purity* and NMI in the experiments.

The clusters are built as follows: In *SNVDM*-based models, each user u is assigned to the viewpoint (cluster) v that maximizes the user-specific viewpoint distribution π_{uv} . *TAM*’s aspect and *SN-LDA*’s community assignments for users are similarly obtained. In *VODUM*, a viewpoint is naturally attributed to each user, as viewpoints are defined at user level.

4.2.1 Clustering Performance. The viewpoint clustering results in terms of *Purity* and NMI for different number of topics $T \in \{5, 10, 15, 20\}$ on both the *Indyref* and *Midterms* datasets are reported in Figure 2. Note that the error bars denote the 95% confidence interval about the mean, which are computed from the 5 repeated executions (i.e., the 5 Markov chains). Confirming (H1a), we observe that for both datasets and for both metrics, our proposed models *SNVDM*, *SNVDM-GPU* ($\tau = 10$) and *SNVDM-GPU* ($\tau = \infty$) outperform all baselines with superior mean *Purity* and mean NMI on both *Indyref* and *Midterms*. The fact that their 95% confidence intervals do not intersect with those of baselines also confirm the statistical significance of this observation at significance level $\alpha = 0.05$. We also notice that our models are mostly insensitive to the number of topics.

Comparing our models *SNVDM*, *SNVDM-GPU* ($\tau = 10$) and *SNVDM-GPU* ($\tau = \infty$) against *SN-LDA* and the degenerate variation *SNVDM-WII*, which both do not exploit incoming interactions, support (H1b). Indeed, all models that use incoming interactions in addition to outgoing interactions significantly outperform those using only outgoing interactions. We nevertheless note that the baseline *SN-LDA* still achieved a strong clustering performance on

both datasets, especially on Indyref. This could be explained by the fact that Indyref features a larger number of interactions per user, upon which SN-LDA heavily relies. It also underlines the key importance of social interactions to viewpoint discovery. Other baselines TAM and VODUM both performed markedly worse than SN-LDA. They obtained similar performance on Midterms and TAM got better results on Indyref. The underperformance of VODUM may be due to its reliance on parts of speech, which are possibly less discriminative in identifying viewpoint-specific words in noisy, short texts like tweets than they are in long and formal documents. Both VODUM and TAM seem to be moderately sensitive to the number of topics, VODUM obtaining better performance for lower number of topics on both datasets.

Hypothesis (H1c) is also validated by the observation that GPU-based models SNVDM-GPU ($\tau = 10$) and SNVDM-GPU ($\tau = \infty$) attained slightly higher clustering performance than that of SNVDM. However, note that SNVDM-GPU ($\tau = 10$) only significantly outperforms SNVDM at significance level $\alpha = 0.05$ on the Midterms dataset: on Indyref their 95% confidence intervals intersect both for Purity and NMI and the difference is therefore not significant. We suspect the large number of interactions in the Indyref dataset to be responsible for the similar clustering results of SN-LDA, SNVDM, SNVDM-GPU ($\tau = 10$) and SNVDM-GPU ($\tau = \infty$). We therefore investigate in the next section to what extent these models are robust to data with less interactions, i.e., sparser social networks.

4.2.2 Robustness to Network Sparsity. In order to study the robustness of models to sparse interactions, we artificially included different percentages of the interactions (100%, 50%, 25%, and 10%) that are available in the Indyref dataset. The interactions to be removed – either incoming or outgoing ones – were randomly chosen. Then, for the obtained datasets derived from Indyref (identical for all evaluated models), we analyzed the viewpoint clustering performance of SN-LDA, SNVDM, SNVDM-GPU ($\tau = 10$) and SNVDM-GPU ($\tau = \infty$). Given that the number of topics only has a mild effect on the performance, as observed in Section 4.2.1, we set $T = 10$ in this experiment. The results are shown in Table 4. Here as well, we provide the 95% confidence interval about the mean. We observe a similar trend for all models: the clustering performance is substantially degraded for low percentage of interactions, especially 10%. This stresses again that interactions are key to the identification of viewpoints in a social network.

Overall, in this studied case, GPU-based models are only marginally more robust to sparsity than other models, with significant improvements only over SN-LDA for 100%, 50%, and 25%; this observation thus merely provides weak support to hypothesis (H2). However, interestingly, SNVDM-GPU ($\tau = 10$) seems to be more robust to interaction sparsity than SNVDM-GPU ($\tau = \infty$): from 50% and below, SNVDM-GPU ($\tau = 10$) performed slightly better. This may be explained by the fact that SNVDM-GPU ($\tau = 10$) is more selective than SNVDM-GPU ($\tau = \infty$) on interacting users: the former uses only the 10 most interacting acquaintances, while the latter leverages all available acquaintances.

4.3 Efficiency

We discuss in this section the efficiency in terms of execution time for the baselines and the proposed models. The machine used to

perform the experiments is a laptop with eight i7-4700MQ CPUs at 2.40 GHz (although our implementation is not parallelized) and 8GB RAM. We report in Table 5 the time (in seconds) taken by one Gibbs sampling iteration on Indyref (with $T = 10$) and Midterms (with $T = 15$). While we observe that the SN-LDA and TAM implementations are faster, the execution times for SNVDM and SNVDM-GPU ($\tau = 10$) are on the same order of magnitude: at most about 2 and 3 times slower, respectively. The execution times on Midterms are very similar for all models, which is explained by the lower number of interactions. With more interactions, the difference is greater on Indyref. Indeed, SNVDM-based and especially GPU-based models are highly dependent on interactions. This is confirmed by the very slow execution time of SNVDM-GPU ($\tau = \infty$) on Indyref. On the contrary, we see that SNVDM-GPU ($\tau = 10$) ran in reasonable time due to the more restrictive selection of related users; this model is therefore a good tradeoff between effectiveness and efficiency.

4.4 Discovered Viewpoints and Topics

In topic modeling literature, the traditional method to study the discovered topics’ quality is to analyze the coherence of the top words (i.e., most probable words) for the topic-specific distributions over words. Based on these top words, topic coherence metrics have been proposed (e.g., in [30]) to quantitatively compare the topics’ quality across different models. However, the models and baselines we wish to compare here have different distributions over words. For example, SN-LDA only has topic-specific distributions over words, while SNVDM-based models distinguish between four types of words, as described in Section 3.2.1: background words (drawn from ϕ_{00}), viewpoint words (drawn from ϕ_{01}), topic words (drawn from ϕ_{10}), and viewpoint-topic words (drawn from ϕ_{11}). Consequently, topic coherence metrics cannot be applied to compare the topics’ quality of our models and baselines. For that reason, we choose to provide in this section a qualitative analysis of the topics and viewpoints discovered by our models.

It is particularly interesting to study for a same topic the topic words as well as the viewpoint words for the different viewpoints: this enables the comparison between neutral word usage against subjective word usage. Therefore, we selected one topic discovered by our most effective model SNVDM-GPU ($\tau = \infty$) in each dataset, Indyref (with $T = 10$) and Midterms (with $T = 15$), and display its top 10 topic words as well as its top 10 viewpoint-topic words specific to “Yes”/“No” supporters in Table 6 and Democrats/Republicans in Table 7, respectively. Note that the topics were manually labeled.

Table 6 shows the top words for a central topic in the Indyref dataset: the question of the *Scottish independence*. As expected, discovered topic words focus only on neutral aspects such as the referendum (*#indyref, vote, campaign*). On the contrary, pro-Yes and pro-No viewpoints are clearly reflected in viewpoint-topic words. “Yes” supporters use specific hashtags such as *#voteyes* and *#yes*. On the other hand, “No” supporters use *#bettertogether*, which was the motto of the pro-No side. Moreover, “No” supporters seem to raise the issue of the *currency* in case of *separation*. In Table 7, we reported the top words about *energy and resources* discovered in the Midterms dataset. This topic showcases a striking difference in Democratic and Republican discourse. Democrats talk about environmental issues, with hashtags such as *#actonclimate* and




Figure 2: Viewpoint clustering results on Indyref (2a, 2b) and Midterms (2c, 2d) datasets in terms of Purity (2a, 2c) and NMI (2b, 2d) for different number of topics (5, 10, 15, and 20). Higher is better. Error bars denote 95% confidence intervals around the mean computed from 5 executions.

Table 4: Viewpoint clustering results in terms of Purity and NMI on Indyref for different percentages of social interactions (100%, 50%, 25%, and 10%). Higher is better, best scores are in boldface. Errors denote 95% confidence intervals around the mean computed from 5 executions.

Model	10% interactions used		25% interactions used		50% interactions used		100% interactions used	
	Purity	NMI	Purity	NMI	Purity	NMI	Purity	NMI
SN-LDA	0.644 ± 0.113	0.087 ± 0.089	0.795 ± 0.056	0.283 ± 0.083	0.863 ± 0.026	0.433 ± 0.045	0.917 ± 0.002	0.588 ± 0.008
SNVDM	0.621 ± 0.148	0.082 ± 0.136	0.912 ± 0.043	0.580 ± 0.135	0.946 ± 0.017	0.702 ± 0.066	0.957 ± 0.002	0.748 ± 0.011
SNVDM-GPU ($\tau = 10$)	0.704 ± 0.097	0.149 ± 0.121	0.938 ± 0.008	0.666 ± 0.031	0.951 ± 0.007	0.720 ± 0.031	0.958 ± 0.004	0.749 ± 0.019
SNVDM-GPU ($\tau = \infty$)	0.687 ± 0.125	0.139 ± 0.117	0.900 ± 0.041	0.546 ± 0.107	0.958 ± 0.003	0.752 ± 0.014	0.964 ± 0.002	0.776 ± 0.010

Table 5: Execution time (in seconds) of one Gibbs sampling iteration for models TAM, SN-LDA, VODUM, SNVDM-WII, SNVDM, SNVDM-GPU ($\tau = 10$) and SNVDM-GPU ($\tau = \infty$) on Indyref (with $T = 10$) and Midterms (with $T = 15$).

	Indyref	Midterms
TAM	1.45	0.87
SN-LDA	1.18	0.64
VODUM	2.78	1.85
SNVDM-WII	2.08	1.08
SNVDM	2.49	1.15
SNVDM-GPU ($\tau = 10$)	3.47	1.34
SNVDM-GPU ($\tau = \infty$)	14.67	2.56

Table 6: The top 10 topic words (1st column) and viewpoint-topic words for “Yes” supporters (2nd column) and “No” supporters (3rd column) about Scottish independence discovered by SNVDM-GPU ($\tau = \infty$) on Indyref.

	Topic: Scottish independence	
	Viewpoint: Yes	Viewpoint: No
#indyref	#voteyes	#indyref
scotland	yes	uk
independence	scotland	salmond
vote	independence	#bettertogether
campaign	westminster	#scotdecides
scottish	vote	separation
uk	independent	currency
people	country	thanks
future	#yes	today
independent	#scotland	say

Table 7: The top 10 topic words (1st column) and viewpoint-topic words for Democrats (2nd column) and Republicans (3rd column) about Energy and resources discovered by SNVDM-GPU ($\tau = \infty$) on Midterms.

	Topic: Energy and resources	
	Viewpoint: Dem.	Viewpoint: Rep.
energy	#actonclimate	#4jobs
house	climate	#obamacare
new	#p2	#jobs
gas	change	gop
natural	#climatechange	obama
#energy	clean	bills
#ff	oil	jobs
#kxl	energy	house
support	#gop	act
economic	sec	watch

#climatechange, while republicans focus on the economic impact such as jobs and bills. Overall, we observe that the discovered topics and viewpoints are reasonably coherent, which confirms (H4).

5 CONCLUSION AND FUTURE WORK

In this paper, we introduced a novel topic model for viewpoint discovery in social networks, namely SNVDM, and its extension based on the generalized Pólya urn (GPU) scheme. We showed throughout extensive experiments on two different Twitter datasets that our proposed models significantly outperform state-of-the-art models in a viewpoint clustering task. The experiments thereby underlined that leveraging both incoming and outgoing interactions as well as exploiting the GPU scheme lead to significant improvement in

terms of Purity and Normalized Mutual Information. This confirms previous findings which highlighted the key importance of social interactions to discovering viewpoints [27]. Additionally, the GPU scheme was found to slightly improve the robustness to social network sparsity. Moreover, proposed models SNVDM and SNVDM-GPU ($\tau = 10$) were shown to be comparably efficient to baselines, thus constituting good tradeoffs between efficiency and effectiveness. The discovered topics and viewpoints were also observed to be reasonably coherent.

Although our approach is widely applicable to study viewpoints in social networks, one of its limitations is that it is most effective for users that are strongly engaged (i.e., users who post and interact a lot). To study less active users’ viewpoints, post-level (e.g., tweet-level) techniques may complement our approach.

In future work, we plan to extend our proposed models and account for the temporal evolution of viewpoints, in order to illustrate the trends of topics and underlying views concurring with current events. Another interesting research direction would be to leverage users' geolocation (e.g., as in [21]) and explore region-level viewpoints. Eventually, we aim to devise an unbiased viewpoint summarization framework for social media to provide balanced viewpoints and thus mitigate the "filter bubble" effect.

ACKNOWLEDGMENTS

We would like to thank Éric Gaussier and Igor Brigadir for their valuable comments and helpful suggestions. This work was partially supported by the French National Research Agency under project LisTIC (ANR-16-CE26-0014-01).

REFERENCES

- [1] Lada A. Adamic and Natalie Glance. 2005. The Political Blogosphere and the 2004 U.S. Election: Divided They Blog. In *Link@KDD*. 36–43.
- [2] Faiyaz Al Zamal, Wendy Liu, and Derek Ruths. 2012. Homophily and Latent Attribute Inference: Inferring Latent Attributes of Twitter Users from Neighbors. In *ICWSM*. 387–390.
- [3] Enrique Amigó, Julio Gonzalo, Javier Artiles, and Felisa Verdejo. 2009. A Comparison of Extrinsic Clustering Evaluation Metrics based on Formal Constraints. *Inform. Retrieval* 12, 4 (2009), 461–486.
- [4] Eytan Bakshy, Solomon Messing, and Lada Adamic. 2015. Exposure to Ideologically Diverse News and Opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132.
- [5] Pablo Barberá. 2015. Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data. *Polit. Anal.* 23, 1 (2015), 76–91.
- [6] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* 3 (2003), 993–1022.
- [7] Igor Brigadir, Derek Greene, and Pádraig Cunningham. 2015. Analyzing Discourse Communities with Distributional Semantic Models. In *WebSci*. 27:1–10.
- [8] Zhiyuan Chen and Bing Liu. 2014. Mining Topics in Documents: Standing on the Shoulders of Big Data. In *KDD*. 1116–1125.
- [9] Raviv Cohen and Derek Ruths. 2013. Classifying Political Orientation on Twitter: It's Not Easy!. In *ICWSM*. 91–99.
- [10] Michael D. Conover, Jacob Ratkiewicz, Matthew R. Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. 2011. Political Polarization on Twitter. In *ICWSM*. 89–96.
- [11] Shiri Dori-Hacohen and James Allan. 2015. Automated Controversy Detection on the Web. In *ECIR*. 423–434.
- [12] Adam G. Dunn, Julie Leask, Xujuan Zhou, Kenneth D. Mandl, and Enrico Coiera. 2015. Associations Between Exposure to and Expression of Negative Opinions About Human Papillomavirus Vaccines on Social Media: An Observational Study. *J. Med. Internet Res.* 17, 6 (2015), e144.
- [13] Anjie Fang, Iadh Ounis, Philip Habel, Craig Macdonald, and Nut Limsopatham. 2015. Topic-centric Classification of Twitter User's Political Orientation. In *SIGIR*. 791–794.
- [14] Seth R. Flaxman, Sharad Goel, and Justin M. Rao. 2016. Filter Bubbles, Echo Chambers, and Online News Consumption. *Public Opin. Quart.* 80, S1 (2016), 298–320.
- [15] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2016. Quantifying Controversy in Social Media. In *WSDM*. 33–42.
- [16] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2017. Reducing Controversy by Connecting Opposing Views. In *WSDM*. 81–90.
- [17] R. Kelly Garrett and Natalie Jomini Stroud. 2014. Partisan Paths to Exposure Diversity: Differences in Pro- and Counterattitudinal News Consumption. *J. Comm.* 64, 4 (2014), 680–701.
- [18] Eduardo Graells-Garrido, Mounia Lalmas, and Ricardo Baeza-Yates. 2015. Finding Intermediary Topics Between People of Opposing Views: A Case Study. (2015). arXiv:1506.00963.
- [19] Eric A. Hardisty, Jordan Boyd-Graber, and Philip Resnik. 2010. Modeling Perspective using Adaptor Grammars. In *EMNLP*. 284–292.
- [20] Gregor Heinrich. 2008. *Parameter Estimation for Text Analysis*. Technical Report. Fraunhofer Institute for Computer Graphics. 1–31 pages.
- [21] Liangjie Hong, Amr Ahmed, Siva Gurumurthy, Alexander J. Smola, and Kostas Tsoutsoulouklis. 2012. Discovering Geographical Topics In The Twitter Stream. In *WWW*. 769–778.
- [22] Mohit Iyyer, Peter Enns, Jordan Boyd-Graber, and Philip Resnik. 2014. Political Ideology Detection Using Recursive Neural Networks. In *ACL*. 1113–1122.
- [23] Aditya Joshi, Pushpak Bhattacharyya, and Mark Carman. 2016. Political Issue Extraction Model: A Novel Hierarchical Topic Model That Uses Tweets By Political And Non-Political Authors. In *WASSA@NAACL HLT*. 82–90.
- [24] Andreas Jungherr. 2016. Twitter Use in Election Campaigns: A Systematic Literature Review. *J. Inform. Tech. Polit.* 13, 1 (2016), 72–91.
- [25] Chenliang Li, Haoran Wang, Zhiqian Zhang, Aixin Sun, and Zongyang Ma. 2016. Topic Modeling for Short Texts with Auxiliary Word Embeddings. In *SIGIR*. 165–174.
- [26] Jun S. Liu. 1994. The Collapsed Gibbs Sampler in Bayesian Computations with Applications to a Gene Regulation Problem. *J. Am. Stat. Assoc.* 89, 427 (1994), 958–966.
- [27] Walid Magdy, Kareem Darwish, Norah Abokhodair, Afshin Rahimi, and Timothy Baldwin. 2016. #ISISisNotIslam or #DeportAllMuslims? Predicting Unspoken Views. In *WebSci*. 95–106.
- [28] Hosam M. Mahmoud. 2008. *Pólya Urn Models*. Chapman & Hall/CRC.
- [29] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press.
- [30] David Mimno, Hanna M. Wallach, Edmund Talley, Miriam Leenders, and Andrew McCallum. 2011. Optimizing Semantic Coherence in Topic Models. In *EMNLP*. 262–272.
- [31] Saif M. Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. SemEval-2016 Task 6: Detecting Stance in Tweets. In *SemEval@NAACL HLT*. 31–41.
- [32] David Newman, Arthur Asuncion, Padhraic Smyth, and Max Welling. 2009. Distributed Algorithms for Topic Models. *J. Mach. Learn. Res.* 10 (2009), 1801–1828.
- [33] Viet-An Nguyen, Jordan Boyd-Graber, Philip Resnik, and Kristina Miler. 2015. Tea Party in the House: A Hierarchical Ideal Point Topic Model and Its Application to Republican Legislators in the 112th Congress. In *ACL/IJCNLP*. 1438–1448.
- [34] Eli Pariser. 2011. *The Filter Bubble: What the Internet Is Hiding from You*. The Penguin Press.
- [35] Michael J. Paul and Roxana Girju. 2010. A Two-Dimensional Topic-Aspect Model for Discovering Multi-Faceted Topics. In *AAAI*. 545–550.
- [36] Michael J. Paul, ChengXiang Zhai, and Roxana Girju. 2010. Summarizing Contrastive Viewpoints in Opinionated Text. In *EMNLP*. 66–76.
- [37] Marco Pennacchiotti and Ana-Maria Popescu. 2011. Democrats, Republicans and Starbucks Afficionados: User Classification in Twitter. In *KDD*. 430–438.
- [38] Ana-Maria Popescu and Marco Pennacchiotti. 2010. Detecting Controversial Events from Twitter. In *CIKM*. 1873–1876.
- [39] Minghui Qiu. 2015. *Mining User Viewpoints in Online Discussions*. Ph.D. Dissertation. Singapore Management University.
- [40] Minghui Qiu and Jing Jiang. 2013. A Latent Variable Model for Viewpoint Discovery from Threaded Forum Posts. In *NAACL HLT*. 1031–1040.
- [41] Minghui Qiu, Yanchuan Sim, Noah A. Smith, and Jing Jiang. 2015. Modeling User Arguments, Interactions, and Attributes for Stance Prediction in Online Debate Forums. In *SDM*. 855–863.
- [42] Minghui Qiu, Liu Yang, and Jing Jiang. 2013. Modeling Interaction Features for Debate Side Clustering. In *CIKM*. 873–878.
- [43] Adithya Rao and Nemanja Spasojevic. 2016. Actionable and Political Text Classification using Word Embeddings and LSTM. In *WISDOM@KDD*.
- [44] Zhaochun Ren, Oana Inel, Lora Aroyo, and Maarten de Rijke. 2016. Time-aware Multi-Viewpoint Summarization of Multilingual Social Text Streams. In *CIKM*. 387–396.
- [45] Mrinmaya Sachan, Avinava Dubey, Shashank Srivastava, Eric P. Xing, and Eduard Hovy. 2014. Spatial Compactness meets Topical Consistency: Jointly modeling Links and Content for Community Detection. In *WSDM*. 503–512.
- [46] Alexandra Schofield and David Mimno. 2016. Comparing Apples to Apples: The Effects of Stemmers on Topic Models. *T. Assoc. Comput. Ling.* 4, 1 (2016), 287–300.
- [47] Aaron Smith. 2014. *Cell Phones, Social Media and Campaign 2014*. Technical Report. Pew Research Center.
- [48] Cass R. Sunstein. 2009. *Republic.com 2.0*. Princeton University Press.
- [49] Didi Surian, Dat Quoc Nguyen, Georgina Kennedy, Mark Johnson, Enrico Coiera, and Adam G. Dunn. 2016. Characterizing Twitter Discussions About HPV Vaccines Using Topic Modeling and Community Detection. *J. Med. Internet Res.* 18, 8 (2016), e232.
- [50] Thibaut Thonet, Guillaume Cabanac, Mohand Boughanem, and Karen Pinel-Sauvagnat. 2016. VODUM: A Topic Model Unifying Viewpoint, Topic and Opinion Discovery. In *ECIR*. 533–545.
- [51] Hanna M. Wallach, David Mimno, and Andrew McCallum. 2009. Rethinking LDA: Why Priors Matter. In *NIPS*. 1973–1981.
- [52] Shuai Wang, Zhiyuan Chen, and Bing Liu. 2016. Mining Aspect-Specific Opinion using a Holistic Lifelong Topic Model. In *WWW*. 167–176.
- [53] Weiwu Zhang, Thomas J. Johnson, Trent Seltzer, and Shannon L. Bichard. 2010. The Revolution Will be Networked: The Influence of Social Networking Sites on Political Attitudes and Behavior. *Soc. Sci. Comput. Rev.* 28, 1 (2010), 75–92.