

DISCOVERING POLITICAL TOPICS IN FACEBOOK DISCUSSION THREADS WITH GRAPH CONTEXTUALIZATION

BY YILIN ZHANG^{*,1}, MARIE POUX-BERTHE^{†,2}, CHRIS WELLS[‡],
KAROLINA KOC-MICHALSKA^{†,2} AND KARL ROHE^{*,1}

University of Wisconsin-Madison^{}, Audencia Business School[†] and
Boston University[‡]*

We propose a graph contextualization method, `pairGraphText`, to study political engagement on Facebook during the 2012 French presidential election. It is a spectral algorithm that contextualizes graph data with text data for online discussion thread. In particular, we examine the Facebook posts of the eight leading candidates and the comments beneath these posts. We find evidence of both (i) candidate-centered structure, where citizens primarily comment on the wall of one candidate and (ii) issue-centered structure (i.e., on political topics), where citizens' attention and expression is primarily directed toward a specific set of issues (e.g., economics, immigration, etc). To identify issue-centered structure, we develop `pairGraphText`, to analyze a network with high-dimensional features on the interactions (i.e., text). This technique scales to hundreds of thousands of nodes and thousands of unique words. In the Facebook data, spectral clustering without the contextualizing text information finds a mixture of (i) candidate and (ii) issue clusters. The contextualized information with text data helps to separate these two structures. We conclude by showing that the novel methodology is consistent under a statistical model.

1. Introduction. Social networking sites (SNSs) such as Facebook and Twitter now make up a major part of Internet communications [Boyd and Ellison (2007), Kaplan and Haenlein (2010)], including political communication. By providing platforms for citizens to publicly communicate with each other and with politicians, SNSs may increase the accessibility of candidates and political dialog [Wellman et al. (2001)] and motivate political engagement within the public [Williams and Gulati (2009, 2013), Hebshi and O'Gara (2011), Kushin and Kitchener (2009)]. They also appear to facilitate the spread of false or offensive information, and a variety of forms of actors to reach micro-targeted publics with a high degree of efficiency [Kreiss and McGregor (2018)]. Since the 2008 US election, particularly [Wattal et al. (2010)], SNSs have been playing a significant role in advertising and interactions during the presidential elections.

Received August 2017; revised March 2018.

¹Supported by NSF Grant DMS-1612456 and ARO Grant W911NF-15-1-0423.

²Supported by Audencia Foundation Research grant.

Key words and phrases. Network, Facebook, topic, spectral clustering, node covariate, stochastic co-Blockmodel.

Drawing meaning from the massive text corpora of political discussion threads on SNSs has been a major project of scholars working in text mining [Pang and Lee (2008), Stieglitz and Dang-Xuan (2013), Grimmer and Stewart (2013)] and sentiment analysis in recent years. One popular text mining approach is the probabilistic topic models based on latent Dirichlet allocation (LDA) [Blei, Ng and Jordan (2003), Blei (2012), Chang and Blei (2009)], which have been extensively used in social science [Ramage, Connolly and Cox (2009)]. Sentiment analysis is another approach to analyze text. It focuses on understanding emotions in the text. Wang et al. (2012) provides a system for real-time sentiment analysis on Twitter during the 2012 US election. For instance, using sentiment analysis and regression, Stieglitz and Dang-Xuan (2012) finds that political tweets on Twitter that contain stronger emotions receive more public interactions. There are also studies of how political sentiment on SNSs reflect the offline political landscape [Tumasjan et al. (2011)], and how it can affect political elections [Choy et al. (2011)].

Apart from the topic or sentiment information, patterns of political discussion on SNSs are also of great theoretical and empirical interests to scholars of communication and political science. Such platforms have long been heralded for their potential to foster a “public sphere” in which ordinary citizens can recognize one another and hear reasons both for and against their own points of view [Papacharissi (2002)]. More recent analyses of online political discourse are less optimistic, identifying instead vitriol, “trolling,” and larger patterns of partisan polarization. As a result, a great deal of research investigates the extent to which online actors are connected to political opponents [Adamic and Glance (2005), Colleoni, Rozza and Arvidsson (2014), Bakshy, Messing and Adamic (2015)].

Another approach to understand structure of political discussions is social network analysis, which aims to identify influential political actors and communities in the discussions [Stieglitz and Dang-Xuan (2012)] and to study properties of the communities [Robertson, Vatrapu and Medina (2010) and Gonzalez-Bailon, Kaltenbrunner and Banchs (2010)]. One popular community detection approach is spectral clustering [von Luxburg (2007)], which is fast, easy to implement, and consistent in block models for network [Holland et al. (1983), Airoldi et al. (2008), Qin and Rohe (2013)].

In this paper, we combine text mining and community detection to investigate the multiple dimensions of citizens’ interactions with political content coming from political actors. In our data, which come from the 2012 French election, citizens commented on presidential candidate’s Facebook posts. This creates a communication network between two types of units: (i) citizens and (ii) candidate-posts, as the eight presidential campaigns each has posts on Facebook, and citizens comment on the posts. This paper studies the structure of the resulting discussion threads.

The activities of the citizens are characterized by (i) which of the candidate-posts they comment on and (ii) the text of their comments. We are interested in two broad types of patterns in these activities: (i) candidate-centered structure, where citizens primarily comment on the wall of one candidate; and (ii) issue-centered structure, in which citizens' attention and expression is directed toward a specific set of issues (e.g., economics, immigration, etc.). To search for such patterns, we cluster the citizens based on their activities. In each cluster, we examine whether the activities of the citizens focus on particular candidates (i.e., candidate-centered) (Section 2.2) or whether the activities focus on certain political issues (i.e., issue-centered) (Section 4). This distinction reflects the possibility that the Facebook conversation might be organized more along lines of partisanship (candidate-centered), as opposed to matters of concern to "issue publics" (issue-centered) [Kim (2009)].

There has been significant progress on both topic modeling for text [Blei (2012)] and community detection for social networks [Airoldi et al. (2008)]. Recently, there has been significant interest in clustering networks for which we have additional information on the citizens in networks [Chang and Blei (2010), Binkiewicz, Vogelstein and Rohe (2017)]. In this paper, we extend these ideas to the setting of discussion threads. Our network is two-way or bipartite, in which the two types of units, citizens and candidate-posts, are linked by commenting in a discussion thread. Below, we refer to the links showing which citizens commented on which candidate-posts as the *network* or the *graph*. We refer to both the text in candidate-posts and the text in citizen-comments as the *text*. The duality between citizens and candidate-posts also appears in the text; candidates say things differently from citizens.

A key difficulty in analyzing this process, and the key methodological innovation of this paper, is to combine these disparate sources of data, the graph information and the two types of text information (citizen-words and thread-words), in a meaningful way. We develop a graph contextualization technique, `pairGraphText`, to leverage high dimensional node covariates into spectral clustering. We extend and specialize the techniques of Binkiewicz, Vogelstein and Rohe (2017) to deal with both (i) the asymmetrical nature of the network between citizens and candidate-posts, and (ii) the high dimensional and sparse nature of the text. With noticeable themes, four subpopulations and four subgroups of the candidate-posts are uncovered by our method. We interpret the clusters by a word-content strategy: For each cluster, we (i) identify keywords, and then (ii) read through central conversations containing the keywords.

Our graph contextualization method, `pairGraphText`, is adaptable to symmetric or directed graphs, unipartite or bipartite, assortative or disassortative, weight or unweighted. It scales to hundreds of thousands of nodes and thousands of covariates (e.g., words). `pairGraphText` uses a sparsity penalty to select the key covariates that align with the graph. After combining the covariates with the graph, we use spectral clustering to compute a partition of the nodes. Finally, we

provide diagnostics to identify key covariates to interpret the different clusters. Theorem 5.2 shows that our method is consistent under the Node-Contextualized Stochastic co-Blockmodel. Section 4 uses `pairGraphText` to identify the issue centered structure in the Facebook discussion threads. In Section 6, we compare `pairGraphText` to a state-of-the-art topic modeling method, relational topic model (RTM) [Chang and Blei (2009)], by both the Facebook discussion threads and simulations. We show that RTM focuses more on the text data, while `pairGraphText` focuses more on the graph data.

This paper is organized as follows. In Section 2, we briefly describe the 2012 French presidential election, the discussion threads on Facebook, and the result of regularized spectral clustering without any contextualizing information. In Section 3, we introduce the graph contextualization technique, `pairGraphText`, which leverages node covariates in spectral clustering. In Section 4, we identify the issue-centered structure of the discussion threads using `pairGraphText`. The statistical consistency of our method is provided under the Node Contextualized Stochastic co-Blockmodel in Section 5. In Section 6, we discuss different choices for weights of words, and we compare `pairGraphText` with a state-of-the-art topic modeling method. Section 7 concludes with a discussion of our method.

2. Background and key summaries of the data. France’s presidential elections proceed in two stages. On April 22, 2012, the first round of voting narrowed the field of candidates from ten to two; the second round, between François Hollande and Nicolas Sarkozy, took place on May 6. In these analyses, we focus on the eight candidates who received at least 1% of the votes in the 1st round of the election. These eight candidates—François Hollande, Nicolas Sarkozy, Marine Le Pen, Jean-Luc Mélenchon, François Bayrou, Eva Joly, Nicolas Dupont-Aignan, and Philippe Poutou—made a total of 3239 posts on Facebook. In response, 92,226 Facebook users, which we call citizens, made 594,685 comments on the candidate-posts.³

There are two main structures that we aim to detect and study in the conversation: (i) candidate-centered structure, where citizens primarily comment on the wall of one candidate; and (ii) issue-centered structure, in which citizens’ attention and expression is directed toward a specific set of issues (e.g., economics, immigration, etc.).

³The data was gathered by sotrender.com. They collect all posts from the official Facebook profiles of the top eight candidates and all the comments beneath them. Citizens who commented on the candidate-posts are distinguished by identification numbers, which are corresponding to the urls of their Facebook profiles. sotrender.com does not control for citizens being human users (nonbots) or being unique users (e.g., without establishing artificial accounts in order to comment on candidate-posts).

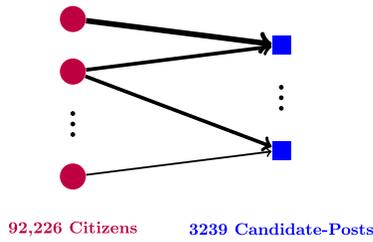


FIG. 1. The communication network is a bipartite graph between citizens and candidate-posts. Each edge weight corresponds to the number of times that a citizen comments on a candidate-post.

2.1. *The communication network.* To study the structure of the conversations, we construct a weighted bipartite network between citizens and candidate-posts (see Figure 1) from the discussion threads.

A citizen is linked to a candidate-post if and only if the citizen comments on the candidate-post. The weight of this link is the number of times the citizen comments on the candidate-post. To represent this network, we construct the weighted adjacency matrix $A \in \mathbb{R}^{92,226 \times 3239}$ with

$$(2.1) \quad A_{ij} = \# \text{ of times of citizen } i \text{ comments on candidate-post } j.$$

Denote the degree of a citizen i , $d_i = \sum_j A_{ij}$, as the number of comments by citizen i . Denote the degree of a candidate-post j , $d_j = \sum_i A_{ij}$, as the number of comments underneath the candidate-post. Figure 2(a) shows the proportion of citizens who have at least d comments, as a function of d . Figure 2(b) gives the same result for the post-degrees.

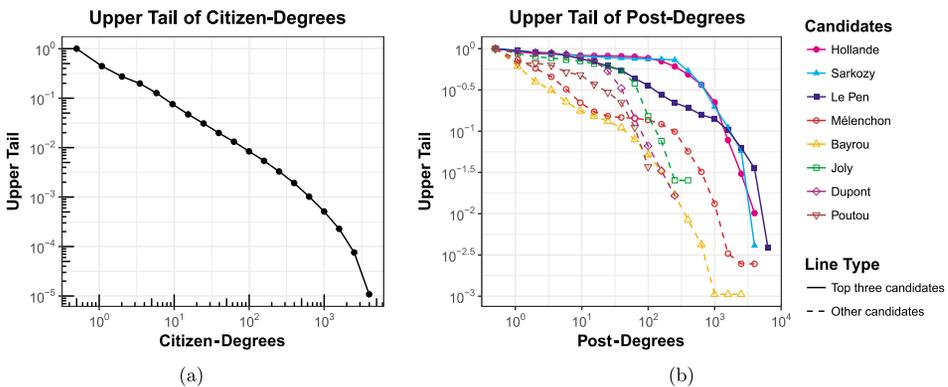


FIG. 2. Upper tail of degrees. Figure (a) shows the upper tail of citizen-degrees. 90% of the citizens write fewer than 10 comments, a small number of citizens write thousands of comments. Figure (b) shows the upper tail of post-degrees by candidate. The top three candidates: Hollande, Sarkozy, and Le Pen (on right), have the largest degrees.

2.2. *Citizens’ attention-ratio toward candidates.* Let ζ_{ij} be the number of times that citizen i comments under candidate j ’s wall. For each citizen i , we denote their attention-ratio as

$$\text{AttentionRatio}(i) = \frac{\max_{\ell} \zeta_{i\ell}}{d_i}.$$

When the attention-ratio is one, it indicates the citizen only comment on one candidate-wall, while smaller attention-ratio indicates the citizen comments across different candidate-walls. We say that citizen i focuses on candidate j if $\zeta_{ij} \geq \zeta_{i\ell}$ for any candidate ℓ . The citizens that have tied favorites are randomly assigned to one of their favorite candidates. Then the citizens are naturally partitioned into eight clusters based on the candidates they focus on. Figure 3 shows the histogram of attention-ratio for all citizens with $d_i \geq 10$. Most of the mass of this histogram is close to one, indicating that most citizens primarily comment on one candidate-wall. This gives the first impression of candidate-centered structure.

Categorizing the citizens based upon where they focus their attention produces a partition. For any partition of citizens, $\mathcal{P} : \{1, \dots, N_C\} \rightarrow \{1, \dots, K_C\}$ where $N_C = 92,226$ is the number of citizens and K_C is the number of citizen-clusters,

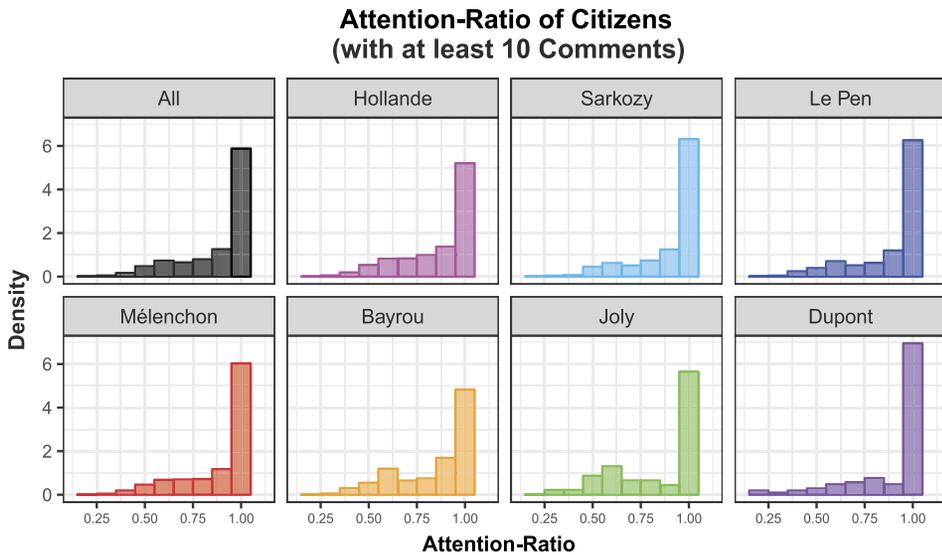


FIG. 3. *Distribution of citizens’ attention-ratio.* In this figure, we focus on citizens who have at least 10 comments. The first plot displays the histogram of attention-ratio for all citizens. The rest eight plots are for the eight citizen-clusters based on the candidates they focus on. We don’t display the citizens who focus on Poutou, because he attracts very few comments.

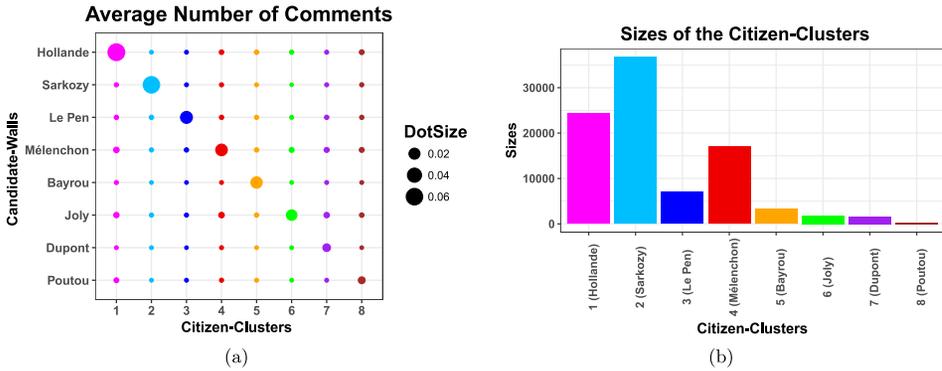


FIG. 4. Citizen-clusters. Figure (a) shows interactions between the citizen-clusters and candidate-walls. The sizes of the balloons are the elements of Ψ_C [defined in (2.2)]. Figure (b) shows the number of citizens in each cluster. In Figure (b), we label each citizen-cluster by the corresponding candidate. For example, the first citizen-cluster is Hollande-centered from Figure (a), so we label it as 1 (Hollande) in Figure (b).

define matrix $\Psi_C \in \mathbb{R}^{K_C \times 8}$ such that for any $a \in \{1, \dots, K_C\}$ and $b \in \{1, \dots, 8\}$,

$$\begin{aligned}
 (2.2) \quad [\Psi_C]_{a,b} = & \text{(# of comments from citizens in cluster } a \text{ under posts} \\
 & \text{on } b\text{th candidate-wall)} / (\text{(# of citizens in cluster } a) \\
 & \times \text{(# of posts on } b\text{th candidate-wall)}).
 \end{aligned}$$

Figure 4 gives a balloon plot of Ψ_C for the partition created by where citizens focus their attention. It also shows a clear candidate-centered structure: Each candidate has a corresponding citizen-cluster that mainly comment on their posts. Combined with the size of each citizen-cluster, it shows leading candidates attract larger clusters of citizens. See the supplementary material on Zhang et al. (2018) for more evidence for candidate-centered structure.

However, such strong candidate-centered structure, where citizens primarily comment on the wall of one candidate, does not lead to the conclusion that citizens devote their attention to candidates rather than issues. It might be an “illusion” from the “magnifying” effect of Facebook [Webster (2014)]. One possibility is many citizens may only follow one candidate on Facebook, so they can only see posts from one candidate. Even if they are interested in topics that are discussed by many candidates, they are likely to comment only on the candidate’s posts that they follow. In this case, even a slight more interest in one candidate can be magnified by Facebook to a strong candidate-centered structure. To understand whether the citizens’ attention is only directed by candidates, we dig more deeply into the discussion threads in the following sections.

Importantly, the partition of citizens in Figure 4, which is created by where citizens focus their attention, uses the additional information of *which of the eight candidates writes each post*. In other words, this partition of the rows of

$A \in \mathbb{R}^{92,226 \times 3239}$ uses a partition of the 3239 columns of A which is defined by which candidate writes the post. The next sections will define two additional partitions of the citizens. Neither of these partitions will use the information of which candidate writes the post. The summary Ψ_C will be computed with these new partitions to help interpret whether they are discovering candidate-centered structure.

2.3. *Studying the graph using DI-SIM.* Despite the overwhelming evidence for strong candidate-centered clusters in Figure 4, the spectral algorithm DI-SIM [Rohe, Qin and Yu (2016)] finds a different partition of the citizens. DI-SIM partitions both citizens and candidate-posts by applying a spectral clustering algorithm. It applies the singular value decomposition to a normalized version of the adjacency matrix A [defined in (2.1)].⁴ By applying k-means to the top left and right singular vectors, DI-SIM partitions the citizens and posts to different clusters.⁵ Figure 5 displays the matrix Ψ_C [defined in (2.2)] for the partition of citizens created by DI-SIM. Only the top three candidates have clusters that focus on them: Hollande and Sarkozy each has two clusters and Le Pen has one cluster that focuses on her. Other citizen-clusters (6, 7, 8) spread across multiple candidates.

One possible reason for the discrepancy between the attention-based partition and the partition from DI-SIM is that there may be some additional structure and DI-SIM is finding a mixture of the candidate-centered structure with that additional structure. `pairGraphText`, which we will introduce in the following sections, confirms that there is also an issue-centered structure in the network by incorporating text information.

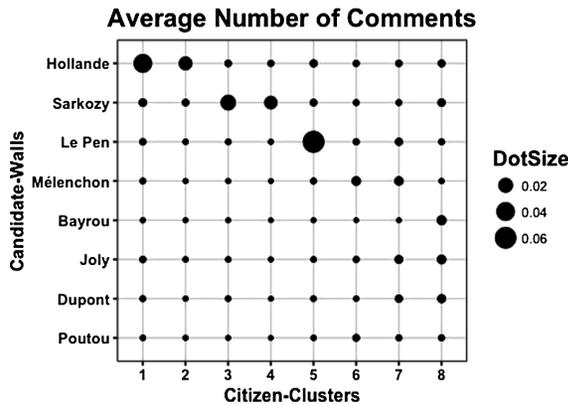


FIG. 5. The citizen-clusters by DI-SIM. Similar to Figure 4(a), this figure shows the balloon plot of Ψ_C corresponding to the citizen-clusters by DI-SIM.

⁴This normalized version of the adjacency matrix A is the regularized graph Laplacian which we will define in details in (3.1)

⁵DI-SIM is similar to the steps 4–7 in Algorithm 1. It applies the singular value decomposition on the graph Laplacian instead.

Algorithm 1 pairGraphText

Input: adjacency matrix $A \in \mathbb{R}^{N_P \times N_C}$, node covariate matrices $X \in \mathbb{R}^{N_P \times M_P}$ and $Y \in \mathbb{R}^{N_C \times M_C}$, number of citizen-clusters K_C , number of post-clusters K_P , weight h , and the significance level α .

1. Compute the regularized graph Laplacian L from A as in (3.1). Center X and Y by column. (In practice, scaling X and Y by rows and columns or using weighted X and Y might also be beneficial. See more details in Section 6.2.)
2. Compute $W = X^T L Y$. Set ω to be the $1 - \alpha$ quantile of $|W_{ij}|$'s.
3. Compute the similarity matrix for pairGraphText as

$$S = L + h X T_\omega(W) Y^T.$$

4. Compute the top K left and right singular vectors $U_C \in \mathbb{R}^{N_C \times K}$, $U_P \in \mathbb{R}^{N_P \times K}$ corresponding to the K largest singular values of S , where $K = \min\{K_C, K_P\}$.
5. Form matrices $U_C^* \in \mathbb{R}^{N_C \times K}$ and $U_P^* \in \mathbb{R}^{N_P \times K}$ such that for any $i \in \{1, \dots, N_C(N_P)\}$,

$$(2.3) \quad [U_C^*]_i = \frac{[U_C]_i}{\|[U_C]_i\|_2} \quad \text{and} \quad [U_P^*]_i = \frac{[U_P]_i}{\|[U_P]_i\|_2}.$$

6. Cluster the rows of U_C^* into K_C clusters with k-means. If the i th row of U_C^* falls in the k th cluster, assign citizen i to citizen-cluster k .
7. Cluster the candidate-posts by performing step 6 on the matrix U_P^* with K_P clusters.

3. Graph contextualization with pairGraphText. As shown in Section 2, there are at least two good clusterings of the nodes (by attention-ratio or by DI-SIM). Given the potentially large number of plausible clusterings of the nodes, the overarching aim of graph contextualization is to find a co-clustering of A (i.e., clustering both its rows and columns) such that these clusters align with a partition in the contextualizing information.

To quantify and utilize the contextualizing information, Section 3.1 describes how we preprocess the text in the discussion threads. Section 3.2 defines the document-term matrices to represent the text used by citizens and candidate-posts. Section 3.3 introduces the pairGraphText algorithm.

3.1. Preprocessing the text. To preprocess the text, we represent the text in document-terms, remove numbers, symbols (e.g., %, @, etc.), and stop words (e.g., le, la, en, au, etc.) and transfer words into their roots by stemming. For example, maintenaiement, maintenait, maintenant, maintenir are transferred into their root mainten.

3.2. *Document-term matrices (node covariate matrices).* From the cleaned text, we retain two different sets of words: “citizen-words” which are contained by at least 0.1% of the comments, and “thread-words” which are contained in at least 0.1% of the contents in threads (i.e., posts and comments). In this data, over 99% of citizen-words and thread-words are overlapped, such as `franc`, `vot`, `plus`, etc. There are also thread-words that are not in citizen-words, such as `confrontaient`, `relancait`, etc.

To contextualize the citizens with the words that they write, define $X \in \mathbb{R}^{N_C \times M_C}$, where N_C is the number of citizens and $M_C = 2020$ is the number of citizen-words. For citizen i and citizen-word j ,

$$X_{ij} = \# \text{ of comments from citizen } i \text{ that contain citizen-word } j.$$

Representing the candidate-posts is not as simple. Candidate-posts provide platforms for conversations, but usually it is the comments underneath it that generate conversations. This phenomenon is colloquially referred to as “thread highjacking,” where the discussion thread (beneath a candidate-post) is used to discuss something other than what is discussed in the candidate-post. In particular, many of the candidate-posts direct their followers to interviews that happen in traditional media. Thus, to properly contextualize the thread, one must include the text that citizens are responding to, which is not necessarily the candidate-post. To represent the text that citizens are responding to when they post a comment in a thread, we use matrix $Y \in \mathbb{R}^{N_P \times M_P}$, where $N_P = 3239$ is the number of candidate-posts and $M_P = 2021$ is the number of thread-words. For candidate-post i and thread-word j ,

$$Y_{ij} = \mathbf{1}\{\text{candidate-post } i \text{ contains thread-word } j\} \\ + \# \text{ of comments underneath candidate-post } i \text{ that contain thread-word } j.$$

We refer to X and Y document-term matrices and consider them as node covariate matrices that contain the text information about both types of nodes (citizens and candidate-posts). The rows index the nodes (citizens or candidate-posts) and columns index the dictionaries (citizen-words or thread-words). Our setting allows citizen-covariates and post-covariates to differ in both type and number. In general, there could be various types of covariates. Note that categorical covariates should be re-expressed with dummy variables. In practice, node covariate matrices X and Y should be centered and scaled by column before analysis.

3.3. `pairGraphText`. `pairGraphText` is a refinement of Covariate Assisted Spectral Clustering (CASC) [Binkiewicz, Vogelstein and Rohe (2017)]. In CASC, the graph is uni-partite. Denote $X \in \mathbb{R}^{N \times M}$ as the node covariate matrix and $L \in \mathbb{R}^{N \times N}$ as the regularized graph Laplacian

$$(3.1) \quad L = D_C^{-1/2} A D_P^{-1/2},$$

where D_C and D_P are diagonal matrices with $[D_C]_{ii} = \sum_j A_{ij} + \tau_c$ and $[D_P]_{jj} = \sum_i A_{ij} + \tau_p$, where $\tau_c(\tau_p)$ is set to be the average row (column) degree. When the uni-partite graph is undirected, $D_C = D_P$. CASC adds the covariate assisted part $C = XX^T$ to the regularized graph Laplacian and performs spectral clustering on the following similarity matrix:

$$S_{\text{casc}}(h) = L + hC.$$

To generalize CASC, `pairGraphText` refines the matrix C in several ways. This refinement will first be expressed in terms of a uni-partite graph where $X = Y$. Replace $C = XX^T$ with

$$C_W = XWX^T$$

for some matrix W . Note that when W is identity matrix, $C_W = C$. By imposing matrix W , `pairGraphText` addresses the following limitations of CASC:

- For any matrix H , denote its i th row as H_i . and its j th column as H_j . Note that $C_W = \sum_{ij} W_{ij} X_i X_j^T$. So, when W_{ij} is nonzero for $i \neq j$, it creates an “interaction” between X_i and X_j , that is, i th and j th covariates. Such interactions are not included in $C = XX^T = \sum_j X_j X_j^T$.
- In C , there is not a natural way of excluding covariates, that is, discarding columns of X . However, in many settings, several covariates could be unaligned with the graph and they should be excluded from the similarity matrix. C_W can select covariates by setting some elements (or rows/columns) of W to zero.
- C presumes that two nodes are more likely to be connected when they have similar covariates. But in some situations, this is not true. For example, in a dating network, relationships are more prevalent among men and women than two people of the same gender. In C_W , if W_{ii} is negative, then two nodes are closer in the similarity matrix C_W if they have different values for the i th covariate.
- The symmetric matrix C only allows for symmetric contributions of covariates, which may not be the case for directed graphs. This can be addressed by allowing W to be asymmetric.
- Finally, CASC was not designed for bipartite networks. In a bipartite graph, the rows of A might have different contextualizing measurements than the columns of A . In the Facebook data, these measurements correspond to the matrices X and Y . Because they have a different number of measurements, the multiplication XY^T is not defined for the Facebook data. However, the multiplication XWY^T is well-defined for a rectangular W . This removes the need for a one-to-one correspondence between the columns of X and Y ; they could contain entirely different types of measurements.

We propose estimating a matrix W to address the issues above. Define the *call-response matrix*

$$(3.2) \quad W = X^T LY,$$

which measures the correlation between thread-words and citizen-words *along the graph*. For example, if discussion threads containing the word `franc` have comments from citizens that are likely to say `vot`, then citizen-word `vot` is highly correlated with a thread-word `franc` along the graph.

To illustrate $W = X^T L Y$, examine a single element $x^T L y$, where $x \in \mathbb{R}^{92,226}$ is a column of X corresponding to word `vot` and $y \in \mathbb{R}^{3239}$ is a column of Y corresponding to word `franc`. So, x_i is the number of times that citizen i uses `vot` and y_j is the number of times that `franc` appears in the thread for candidate-post j . If x is centered and independent of L and y , then x is an uninformative covariate, and $\mathbb{E}[x^T L y] = \mathbb{E}(\mathbb{E}(x^T | L, y) L y) = 0$. Conversely, if for centered x and y ,

$$x^T L y = \sum_{i,j:A_{ij}=1} \frac{x_i y_j}{\sqrt{[DC]_{ii} [DP]_{jj}}}$$

is large (positive or negative), it suggests that linked nodes in L have (positively or negatively) correlated values of x and y . Figure 6 gives a small part of the call-response matrix.

There are thousands of words in the discussion threads. To select the highly correlated words along the graph, we define a hard-threshold function on W ,

$$(3.3) \quad [T_\omega(W)]_{sr} = \begin{cases} W_{sr} & \text{if } W_{sr} > \omega, \\ 0 & \text{o.w.} \end{cases}$$

In practice, we can set the threshold ω as the $1 - \alpha$ quantile of $|W_{ij}|$'s.

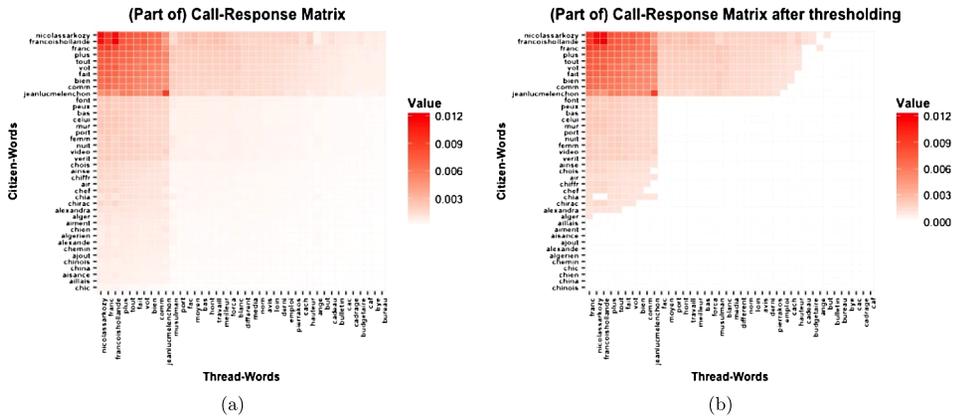
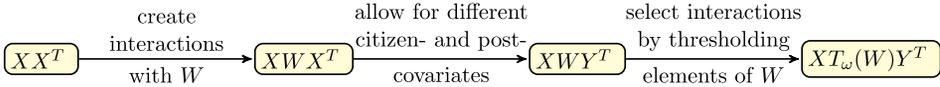


FIG. 6. Part of the call-response matrix before and after thresholding. Some pairs of words are relatively more highly correlated, like `nicolassarkozy` and `francoishollande`, `jeanlucmelenchon` and `jeanlucmelenchon`, `vot` and `franc`, etc. After thresholding, only the relatively highly correlated pairs of words are left, making the call-response matrix much more sparse.

Thus, we finally define the matrix that replaces C from CASC. For `pairGraphText`, define

$$(3.4) \quad C_T = XT_\omega(W)Y^T.$$

The following diagram reviews how `pairGraphText` refines the matrix C from CASC:



Note that

$$C_T = \sum_{ij} [T_\omega(W)]_{ij} X_{:i} Y_{:j}^T$$

shows closeness of citizens and candidate-posts based on their usage of words in the network. $[C_T]_{ij}$ is large when citizen i and candidate-post j use many highly correlated pairs of words. The threshold function $T_\omega(\cdot)$ helps select pairs of words, and imposes sparsity when W is high-dimensional.

Therefore, `pairGraphText` applies DI-SIM to the similarity matrix:

$$(3.5) \quad S = L + hC_T.$$

This similarity matrix combines both the graph information, represented by L , and the text information, represented by $C_T = XT_\omega(W)Y^T$, with a tuning parameter h to balance between these two parts.

4. Issue-centered structure. We identify topics that attract public’s attention in the Facebook discussion threads using `pairGraphText`. We scale the document-term matrices by both rows and columns.⁶ From the scree plot of the singular values of S [see Figure 3 in supplementary material Zhang et al. (2018)], we decide to study the top $K = 4$ clusters due to the large gap after the fourth singular value. To study how the text in discussion threads affects the partition of citizens and candidate-posts, we show the clustering results in three cases: (i) when we use no text, that is, the tuning parameter h in equation (3.5) is $h = 0$,⁷ (ii) when we incorporate text, that is, $h = 0.035$,⁸ and (iii) when we only use the text assisted part [defined in (3.4)], that is, $h = \infty$.

⁶We replace X_{ij} and Y_{ij} by $X_{ij}/\sqrt{\sum_i X_{ij} \sum_j X_{ij}}$ and $Y_{ij}/\sqrt{\sum_i Y_{ij} \sum_j Y_{ij}}$.

⁷When $h = 0$, `pairGraphText` is equivalent to DI-SIM.

⁸In case (ii), h can be any real positive value. We choose $h = 0.035$ since it shows clusters with major differences from both cases when $h = 0$ and when $h = \infty$. Recall the similarity matrix $S = L + hC_T$ [see (3.5)]. For identification of $h = 0.035$, we scale the text-assisted part C_T to have the same second singular value with L . Then h means how much we weigh the text-assisted part in `pairGraphText`. $h = 0.035$ means that we weigh the text-assisted part 0.035 times of the graph information.

Section 4.1 shows that with more text incorporated (i.e., with larger h), the clusters become less candidate-centered. Section 4.2 introduces a word-content strategy to extract topics of clusters. Section 4.3 describes the cluster topics and supports Section 4.1 by showing that clusters with larger h are more heavily focused on the contextualizing information.

4.1. *The clusters from pairGraphText with larger h are less candidate-centered.* For each partition of candidate-posts, $\mathcal{P} : \{1, \dots, N_P\} \rightarrow \{1, \dots, 4\}$, we define the matrix $\Psi_P \in \mathbb{R}^{4 \times 8}$ such that for any $a \in \{1, \dots, 4\}$ and $b \in \{1, \dots, 8\}$,

$$(4.1) \quad [\Psi_P]_{ab} = \frac{\text{\# of posts in cluster } a \text{ from candidate } b\text{'s wall}}{(\text{\# of posts in cluster } a) \times (\text{\# of posts from candidate } b\text{'s wall})}.$$

Ψ_P shows how post-clusters distribute on candidate-walls. This is similar to Ψ_C defined in (2.2), which shows how citizen-clusters interact with candidate-walls. Figure 7 displays Ψ_P and Ψ_C in balloon plots in the three cases. When we use no text, that is, $h = 0$, there appears some candidate-centered structure in both citizen-clusters and post-clusters. As we incorporate text, in the case when $h = 0.035$, each post-cluster spreads across multiple candidates. With even more text incorporated, in the case $h = \infty$, neither of the post-clusters nor citizen-clusters are candidate-centered. In the following subsections, we identify the cluster topics using key words, comments, and posts.

4.2. *A word-content strategy to identify cluster topics.* To identify the cluster topics, we first identify keywords for each cluster, which we will define in Section 4.2.1. These keywords give the first impression of the cluster topics.

However, it is insufficient to examine the words in isolation, because the same word is often used differently by different subsets of the population. For example, religion is often used by citizens talking about the religion of peace and it is also often used by atheists criticizing its appearance in the public sphere. Thus, to identify the cluster topics, besides identifying keywords, we also need to read through the conversations that contain these keywords. We focus on the central conversations in each cluster, which we will define in Section 4.2.2.

We call this strategy word-content strategy, where for each cluster, we (i) identify the keywords and (ii) read through the central conversations that contain the keywords in the cluster.

4.2.1. *Identify the keywords.* We identify the keywords in each cluster by setting “scores.” For any $k \in \{1, \dots, 4\}$ and $j \in \{1, \dots, M_C\}$, define the score of

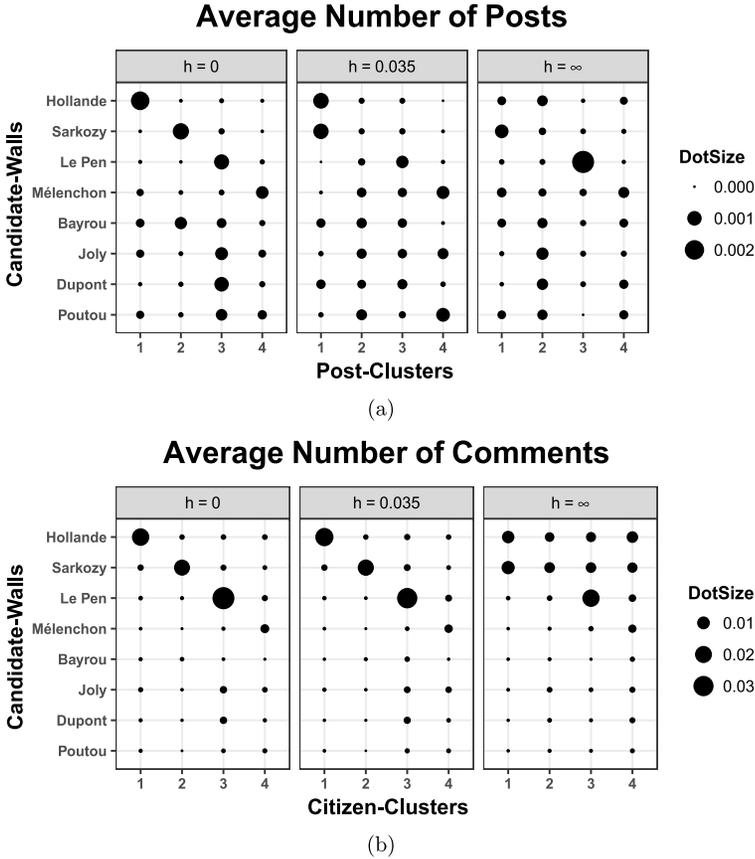


FIG. 7. Clusters and candidate-walls. Figure (a) and (b) display Ψ_P and Ψ_C in balloon plots for the three cases.

citizen-word j in citizen-cluster k as

$$\Phi_{kj} = \frac{\sum_{i \in k} X_{ij}}{\sum_{i \in k} \hat{X}_{ij}} \quad \text{where } \hat{X}_{ij} = \frac{\sum_j X_{ij} \sum_i X_{ij}}{\sum X_{ij}},$$

and $i \in k$ denotes the citizen i belongs to cluster k . We similarly define the scores of thread-words in post-clusters based on the document-term matrix of candidate-posts Y . These scores are also discussed in [Witten \(2011\)](#), where they are derived by maximum likelihood on a Poisson model. We define the keywords in a cluster to be the words with the largest scores in the cluster. We show keywords of each cluster in Section 4.3.

4.2.2. *Identifying central conversations.* We identify the central conversations by diagnostics from k-means clustering. Recall the `pairGraphText` algorithm partitions citizens by applying k-means on the N_C rows of matrix $U_C^* \in \mathbb{R}^{N_C \times 4}$

[defined in (2.3)] which correspond to the N_C citizens. For any citizen i , we denote their cluster-centrality as

$$\rho_i = [U_C^*]_i^T [\mu_C^*]_i,$$

where $[\mu_C^*]_i$ is the cluster centroid of citizen i from k-means on rows of U_C^* . There are four different cluster centroids. For each cluster, the central citizens are the citizens in the cluster with the largest cluster-centrality, that is, those that align best with the cluster centroid. We similarly define the central posts for post-clusters. For a citizen-cluster, the central conversations are the comments from the central citizens; for a post-cluster, the central conversations are the discussion threads (including posts and comments) initiated by the central posts.

We read through the central conversations that contain the keywords in each cluster. This word-content strategy helps us identify topics that attract citizens' attention. We will show these topics in Section 4.3.

4.3. *Topics of clusters.* We extract topics of the clusters by the word-content strategy in three cases, $h = 0$, $h = 0.035$, and $h = \infty$. Figure 8, 9 and 10 show the cluster topics with the keywords and a brief description of the central conversations in each cluster. In these figures, the links indicate major interactions⁹ between citizen-clusters and post-clusters, with the link widths proportional to elements of

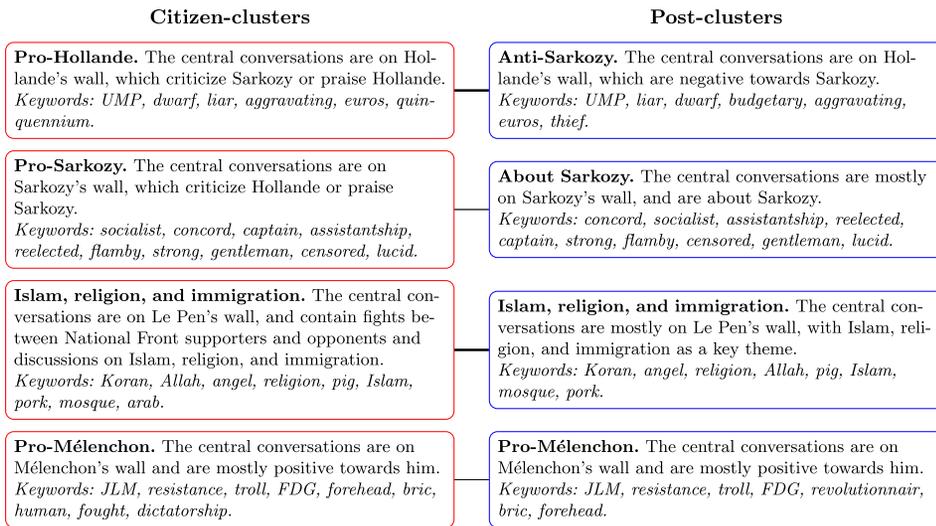


FIG. 8. Cluster topics when $h = 0$.

⁹We only display the links that correspond to the three or four largest elements of Ψ in each case.

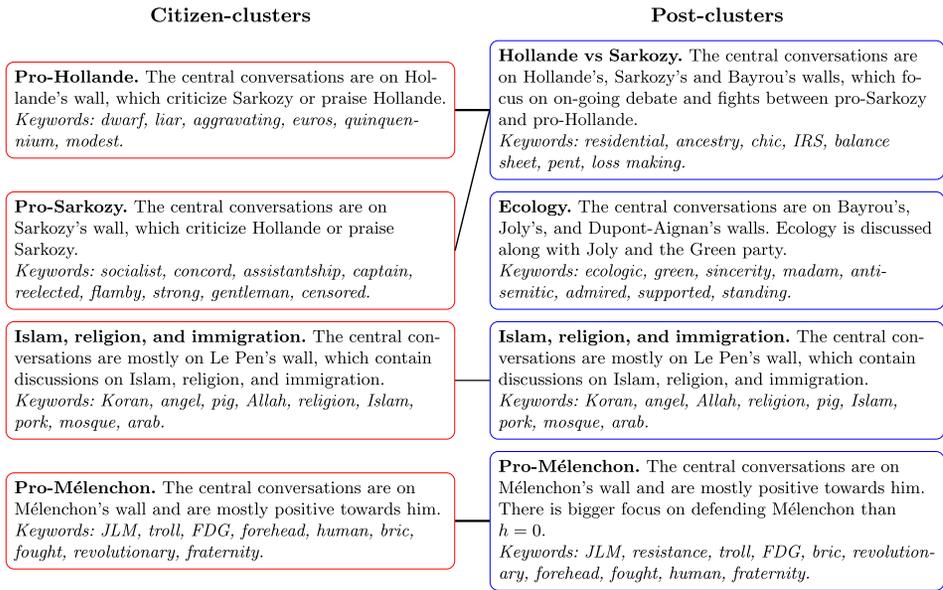


FIG. 9. Cluster patterns when $h = 0.035$.

matrix $\Psi \in \mathbb{R}^{4 \times 4}$, where for any $a, b \in \{1, \dots, 4\}$,

$$\Psi_{ab} = (\# \text{ of comments from citizens in citizen-cluster } a \text{ under candidate-posts from post-cluster } b) / ((\# \text{ of citizens in citizen-cluster } a) \times (\# \text{ of candidate-posts in post-cluster } b)).$$

This is similar to matrices Ψ_C defined in (2.2) and Ψ_P defined in (4.1), which show how clusters (for citizens or candidate-posts) distribute on the eight candidate-walls. Ψ shows how the citizen-clusters interact with the post-clusters.

When $h = 0$ (see Figure 8), clusters focus on candidates or the radical discussions. As we incorporate the text, in the case when $h = 0.035$ (see Figure 9), the citizen-clusters are similar to those when $h = 0$, but there appears a post-cluster about ecology. As we incorporate more text, in the case when $h = \infty$ (see Figure 10), we identify more topics, such as economic and crises. There also appear a cluster for both citizens and candidate-posts with many copy-paste comments. More data analysis results are in a Shiny App available at <https://yilinzhang.shinyapps.io/FrenchElection>.

Incorporating the text makes the central conversations more vivid representations of the clusters, allowing for a more precise interpretation of the topic. During the 2012 French election, the citizens devoted their attention and expression in (i) the debates and fights among different candidates, (ii) radical discussions on Islam,

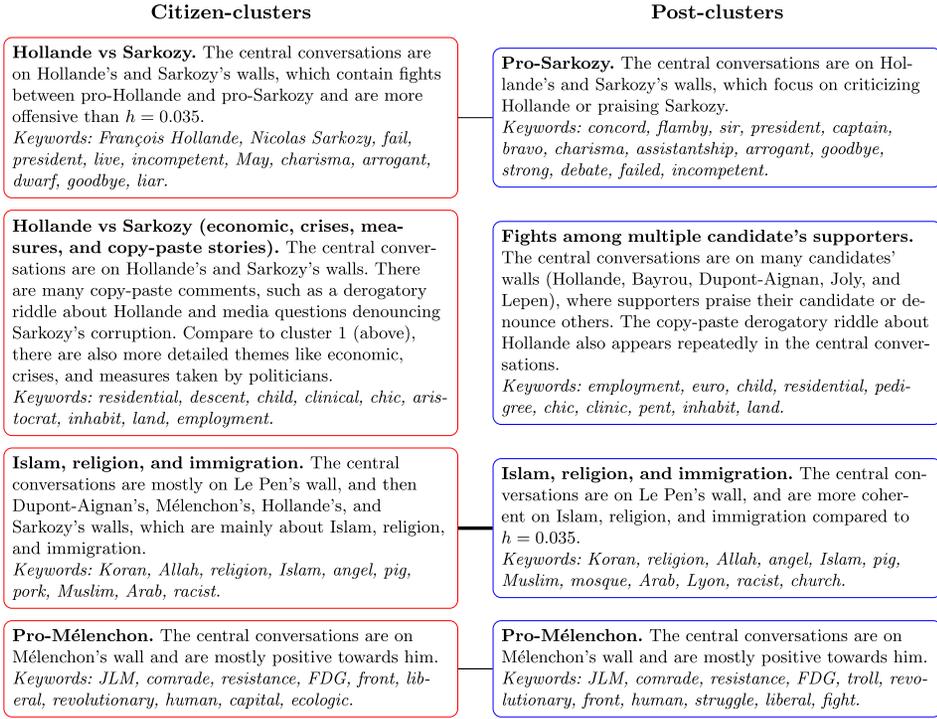


FIG. 10. Cluster patterns when $h = \infty$.

religion, and immigration, and (iii) other topics including ecology, economy, and crises.

5. Statistical consistency of pairGraphText. This section shows that our graph contextualization method, pairGraphText, is statistically consistent under the Node Contextualized Stochastic co-Blockmodel (NC-ScBM), which is a fusion of the NC-SBM [Binkiewicz, Vogelstein and Rohe (2017)] and ScBM [Rohe, Qin and Yu (2016)].

DEFINITION 5.1. Let $Z_C \in \{0, 1\}^{N_C \times K_C}$ and $Z_P \in \{0, 1\}^{N_P \times K_P}$, such that there is only one 1 in each row and at least one 1 in each column. Let $B \in [0, 1]^{K_C \times K_P}$ be of rank $K = \min\{K_C, K_P\}$. Let $E_C \in \mathbb{R}^{K_C \times M_C}$ and $E_P \in \mathbb{R}^{K_P \times M_P}$. Under the NC-ScBM, the adjacency matrix $A \in \{0, 1\}^{N_C \times N_P}$ contains independent Bernoulli random variables with

$$(1) \quad \mathcal{A} = \mathbb{E}[A] = Z_C B Z_P,$$

and the node covariate matrices $X \in \mathbb{R}^{N_C \times M_C}$ and $Y \in \mathbb{R}^{N_P \times M_P}$ contain independent sub-Gaussian elements with

$$(2) \quad \mathcal{X} = \mathbb{E}[X] = Z_C E_C \quad \text{and} \quad \mathcal{Y} = \mathbb{E}[Y] = Z_P E_P.$$

Recall the similarity matrix for pairGraphText defined in equation (3.5), $S = L + hXT_\omega(W)Y^T$. We define the population similarity matrix as

$$(5.1) \quad \mathcal{S} = \mathcal{L} + h\mathcal{X}\mathcal{W}\mathcal{Y}^T,$$

where $\mathcal{L} = \mathcal{D}_C^{-1/2} \mathcal{A} \mathcal{D}_P^{-1/2}$ and $\mathcal{W} = \mathcal{X}^T \mathcal{L} \mathcal{Y}$, where diagonal matrices $[\mathcal{D}_C]_{ii} = \sum_j \mathcal{A}_{ij} + \tau_c$ and $[\mathcal{D}_P]_{jj} = \sum_i \mathcal{A}_{ij} + \tau_p$. Let U_C and $U_P \in \mathbb{R}^{N_C \times K}$ (U_P and $U_P \in \mathbb{R}^{N_P \times K}$) contain the top K left (right) singular vectors of S and \mathcal{S} .

The basic outline of the proof for statistical consistency is: Under some conditions:

1. the element-wise difference between $T_\omega(W)$ and \mathcal{W} is bounded by ω in probability;
2. the similarity matrix S converges to \mathcal{S} in probability;
3. the singular vectors U_C and U_P converge to \mathcal{U}_C and \mathcal{U}_P within some rotations in probability;
4. the misclustering rates for citizens and candidate-posts goes to zero in probability.

The definition of misclustered is the same as in Rohe, Qin and Yu (2016) and is given in Section 3.2 in supplementary material Zhang et al. (2018). The complete proof is given in Section 3.3 in supplementary material Zhang et al. (2018).

Denote $\|\cdot\|$ as the spectral norm and $\|\cdot\|_F$ as the Frobenius norm. For any matrix H , we define $\text{sym}(H) = \begin{pmatrix} 0 & H \\ H^T & 0 \end{pmatrix}$ and $\|H\|_2 = \max(\|\max_i H_i\|_2, \max_j \|H_{\cdot j}\|_2)$. Denote $\|\cdot\|_{\phi_2}$ as the sub-Gaussian norm, such that for any random variable ξ , there is $\|\xi\|_{\phi_2} = \sup_{t \geq 1} t^{-1/2} (\mathbb{E}|\xi|^t)^{1/t}$. To simplify notation, we denote N as the number of nodes and M as the number of covariates, though N_C and N_P , M_C , and M_P can be different.

THEOREM 5.2. *Suppose A , X , and Y , are the adjacency matrix and the node covariate matrices sampled from the NC-ScBM. Let $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_K > 0$ be the K nonzero singular values of \mathcal{S} . Let \mathcal{M}_C and \mathcal{M}_P be the misclustered citizens and the misclustered candidate-posts. Denote q_c and q_p as the largest sizes of citizen-clusters and post-clusters. Define $\delta = \min(\min_i [\mathcal{D}_C]_{ii}, \min_j [\mathcal{D}_P]_{jj})$ and $\gamma = \max(\|X\|_2, \|Y\|_2, \|\mathcal{X}\|_2, \|\mathcal{Y}\|_2)$. Define $\xi = \max(\sigma^2 \|L\|_F \sqrt{\ln M}, \sigma^2 \|L\| \ln M, \frac{\gamma^2}{\delta} \sqrt{\ln M})$, where L is the regularized graph Laplacian defined in equation (3.1) and $\sigma = \max(\max_{ij} \|X_{ij} - \mathcal{X}_{ij}\|_{\phi_2}, \max_{ij} \|Y_{ij} - \mathcal{Y}_{ij}\|_{\phi_2})$. For any $\epsilon \in (0, 1)$, assume:*

- (1) $\delta > 3 \ln(2N) + 3 \ln(8/\epsilon)$,
- (2) $\xi = o(\omega)$, and
- (3) $h \leq \min(\frac{a}{\gamma^2 \|\text{sym}(\mathcal{W})\|}, \frac{a}{\gamma^2 \omega})$, where $a = \sqrt{\frac{3 \ln(16N/\epsilon)}{\delta}}$.

Then, with probability at least $1 - \epsilon$, for large enough N , the misclustering rates

$$\frac{|\mathcal{M}_C|}{N} \leq \frac{c_0 q_c K \ln(16N/\epsilon)}{N \lambda_K^2 \delta} \quad \text{and} \quad \frac{|\mathcal{M}_P|}{N} \leq \frac{c_0 q_p K \ln(16N/\epsilon)}{N \lambda_K^2 \delta},$$

for some constant c_0 .

REMARK. Assumption (1) indicates the sparsity of the graph. Assumption (2) and (3) are conditions on parameters ω and h for consistency. Note the largest sizes of clusters q_c and q_p are $O(N)$. Suppose λ_K is lower bounded by some constant $c_1 > 0$, which indicates the “signal” of each of the K blocks is strong enough to be detected. Then, when δ grows faster than $\ln N$, we have misclustering rates goes to zero as $N \rightarrow \infty$.

6. Comparison analysis. Section 6.1 shows the importance of the call-response matrix W . Section 6.2 discusses different scaling and weighting choices for document-term matrices. In Section 6.3, we compare `pairGraphText` with state-of-the-art topic modeling approach, relational topic model (RTM) [Chang and Blei (2009)], on both the Facebook discussion threads (Section 6.3.1) and on the simulated data (Section 6.3.2).

6.1. *Importance of the call-response matrix W .* Recall the call-response matrix W [defined in (3.2)], which shows the correlation between citizen-words and thread-words on the communication network. It induces weights on different pairs of citizen-words and thread-words; word-pairs with higher correlation on the network are weighted more. In this section, we show the importance of the weights induced by the matrix W . We compare `pairGraphText` with the all-one `pairGraphText`, which replaces matrix W by the “all-one” matrix, $J \in \mathbb{R}^{M_C \times M_P}$, where $J_{ij} = 1$, for all i, j . For comparison, we set the tuning parameter in (3.5) as $h = \infty$. Tables 1 and 2 show the keywords of each cluster by `pairGraphText` and by all-one `pairGraphText`.

Without weights on the word-pairs, the all-one `pairGraphText` fails to extract some topics in the citizen-clusters, such as Islam and the debates among top candidates, which are clear in the citizen-clusters by `pairGraphText`. Moreover, some words appear in multiple clusters by the all-one `pairGraphText`. For example, the word *Allah* appears in both post-clusters 3 and 4 in Table 2. This makes it harder to distinguish different topics between different clusters.

6.2. *Different choices for document-term matrices.* Recall the document-term matrices X and Y (defined in Section 3.2). These matrices don’t consider lengths of comments and posts or popularities of words. We can address this issue by either (1) scaling the document-term matrices by rows and columns as in Section 4, that is, replacing X_{ij} and Y_{ij} by

$$X_{ij} / \sqrt{\sum_i X_{ij} \sum_j X_{ij}} \quad \text{and} \quad Y_{ij} / \sqrt{\sum_i Y_{ij} \sum_j Y_{ij}},$$

TABLE 1
Keywords in clusters by pairGraphText

	Citizen-Clusters	Post-Clusters
Cluster 1	François Hollande, Nicolas Sarkozy, fail, president, live, incompetent, May, charisma, arrogant, dwarf, goodbye, liar	Concord, flamby, sir, president, captain, bravo, charisma, assistantship, arrogant, goodbye, strong, debate, failed, incompetent
Cluster 2	Residential, descent, child, clinical, chic, aristocrat, inhabit, land, employment	Employment, euro, child, residential, pedigree, chic, clinic, pent, inhabit, land
Cluster 3	Koran, Allah, religion, Islam, angel, pig, pork, Muslim, Arab, racist	Koran, religion, Allah, angel, Islam, pig, Muslim, mosque, Arab, Lyon, racist, church
Cluster 4	JLM, comrade, resistance, FDG, front, Jean-Luc Mélenchon, liberal, revolutionary, human, capital, ecologic	JLM, comrade, resistance, FDG, troll, Jean-Luc Mélenchon, revolutionary, front, human, struggle, liberal, fight

or (2) using the weighted document-term matrices. One standard weighting method is TF-IDF (term frequency-inverse document frequency), which is commonly used in information retrieval and text mining [Salton, Wong and Yang (1975), Joachims (1996), Sivic and Zisserman (2003), Ramos (2003)]. TF-IDF weights words based on both the document length and the word popularity. For each word i and document j , the TF-IDF is

$$\frac{\text{\# of occurrences of word } i \text{ in document } j}{\text{\# of words in document } j} \times \log_2 \frac{\text{\# of documents}}{\text{\# of documents that contain word } i}.$$

TABLE 2
Keywords in clusters by all-one pairGraphText

	Citizen-Clusters	Post-Clusters
Cluster 1	Identity, trick, fascist, opposite, reducer, Allah, flamby, top, continuous, incarnate, commercial, mission	Continued, resistance, great, passion, channel, bravo, debate, fight, difficult, goodbye, great, beat, stand, hope
Cluster 2	Baptist, professor, suburb, king, happiness, aristocrat, sincerity, school, regime, residential, exist, erasable, place	Troll, military, comrade, raid, concord, max, Philippe Poutou, tomorrow, soldier, killer, victim, hateful, bulletin, Jewish, fraud
Cluster 3	Dismissal, fraud, multiple, lump, aggravating, unfair, review, gift, parliamentary, budget, referendum	Allah, Israel, altarpiece, foul, angel, dozen, list, cuckoo, municipal
Cluster 4	Parisian, Russian, discriminant, defense, land, vineyard, flag, revel, pedigree, captain, conceivable	Lucid, Allah, African, boat, clandestine, sister, successful, realist, old, movie, angel, tear, promise

In our data, documents are the posts and comments in the discussion threads. For the weighted document-term matrix of citizens Y , we first calculate the TF-IDF matrix of comments, and then add up those comments from the same citizen. The weighted document-term matrix X is the TF-IDF matrix of posts. We compare plain (unscaled and unweighted), scaled, and TF-IDF weighted document-term matrices on the Facebook discussion threads in Section 3 in the supplementary material Zhang et al. (2018).

6.3. *Comparison with relational topic model.* This section compares pair-GraphText and Relational Topic Model(RTM) on both Facebook discussion threads (Section 6.3.1) and simulated data (Section 6.3.2).

6.3.1. *Comparison with relational topic model on the Facebook discussion threads.* Relational topic model (RTM) [Chang and Blei (2009)] is a popular approach to extract topics from documents with a network structure (e.g., citation network). RTM is designed for unipartite networks, where there is only one type of nodes. To apply RTM on the bipartite network with both candidate-posts and citizens, we consider two approaches, (1) symmetrized network (Table 3) and (2) co-occurrence network (Table 4).

We define the symmetrized network as a network with posts and citizens, disregarding the different types of nodes. Recall the adjacency matrix $A \in \mathbb{R}^{92,226 \times 3239}$ (defined in 2.1), the adjacency matrix for the symmetrized network is $\text{sym}(A) = \begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix}$. Table 3 shows the keywords of the four topics by RTM on the symmetrized network. Words such as *Nicolas Sarkozy* appears in both Clusters 1 and 3, making it hard to distinguish between different topics. The racial topic (Islam, religion, and immigration), which is clear by pairGraphText, is not that clear in Table 3.

TABLE 3
Keywords in topics by RTM on symmetrized network

Topics	
Cluster 1	Nicolas Sarkozy, president, live, François Hollande, bravo, France courage, all, good, strong, debate
Cluster 2	Residential, land, pent, pedigree, clinical, inhabit, chic, functional, school, conceivable, childhood
Cluster 3	Nicolas Sarkozy, fair, good, other, must, polish, can, say, nothing, François Bayrou, generation
Cluster 4	Fair, good, Jean-Luc Mélenchon, nothing, other, speak, share, yes, say, generation, front, racist, Muslim

TABLE 4
Keywords in topics by RTM on co-occurrence networks

	Citizen-Topics	Post-Topics
Cluster 1	François Hollande, Nicolas Sarkozy, France, Jean-Luc Mélenchon, good, all, fair, nothing, must, president	Nicolas Sarkozy, François Hollande, social, debate, may, president, victory, change, rich, augment, tax, poor, million
Cluster 2	Residential, clinic, stock market, live, childhood, build, school, free, assembly, departure, functional, bourgeois, depend	President, franc, sir, live, bravo, all, aim, pay, win, good, want, FDG
Cluster 3	Muslim, religion, Islam, speak, racist, insult, Arab, Koran, evil, fear, angel	Islam, religion, racist, evil, immigrant, Arab, insult, Koran, know, from, Jewish, racism
Cluster 4	European, financial, public, undertaken, billion, public, budget, advice, bank, balance sheet, service, euros, jobs	More, good, France, Jean-Luc Mélenchon, all, fair, Nicolas Sarkozy, François Hollande, speak, can, politic, generation

We define the co-occurrence network of posts as a network of posts, where the link width is large when the two posts share many citizens who comment frequently on both posts. Similarly, we define the co-occurrence network of citizens as a network of citizens, where the link width is large when the two citizens comment a lot on many same posts. We define the adjacency matrix of the co-occurrence network for posts as $A^T A$ and the adjacency matrix of the co-occurrence network for citizens as AA^T .

Table 4 shows the keywords of the four topics by RTM on the co-occurrence networks. Similar to `pairGraphText`, RTM also extracts topics like Islam and religion, debates among top candidates, and economic issues.

6.3.2. Comparison with relational topic model on simulated data. In this section, we use simulation examples to compare `pairGraphText` and RTM based upon both statistical accuracy and computational running time.

We simulate documents with links and text, then use `pairGraphText` and RTM to cluster these documents. There are two sources of data, (1) links between documents, that is, graph, and (2) text in the documents, that is, text. We compare `pairGraphText` and RTM in three cases: (1) when both the graph and text contain block information (both signals), (2) when only the graph contains block information (graph signals), and (3) when only text contains block information (text signals). For each of the three cases, we simulate varying levels of signal strength. (See more details on how we define signals in the next paragraph.) For each signal level, we simulate 100 random data sets. Each data set consists of 1000 documents and 1000 words in total, with around 200 words and 20 links per document. In this step, we simulate the documents with links and words under a block model with two blocks, each with around 500 documents and around 500

words. (See more details for the block model in the next paragraph.) On each data set, we run `pairGraphText` and RTM to partition the 1000 documents into two clusters. For RTM, we define its estimated cluster label for each document i as $\max_k \#$ of words in document i belongs to block k .

We simulate all the adjacency matrices (graphs) and the document-term matrices (text) under a Degree Corrected Stochastic Blockmodel [Karrer and Newman (2011)]. Denote $z(i)$ as the block label of any document i and $z_{\text{text}}(w)$ as the block label of any word w . Under this model, two documents i and j are linked with each other with probability $\theta_i \theta_j B_{z(i)z(j)}$, and document i contains word w with probability $\theta_i \theta_w^{\text{text}} B_{z(i)z_{\text{text}}(w)}^{\text{text}}$, where $\theta_i, \theta_j, \theta_w^{\text{text}}$ are degree parameters. The element B_{uv} shows the expected number of links between blocks u, v , and the element B_{uv}^{text} shows the expected number of appearances for words in block v in documents in block u . We define $B \propto \begin{pmatrix} 0.1 & 0.1 \\ 0.1 & 0.1 \end{pmatrix} + \text{sig}_g \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ and $B^{\text{text}} \propto \begin{pmatrix} 0.1 & 0.1 \\ 0.1 & 0.1 \end{pmatrix} + \text{sig}_t \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, where the graph signal sig_g and the text signal $\text{sig}_t \in \{1e-1.8, 1e-1.6, \dots, 1e3\}$ separately show graph links and words contain how much block information.

For `pairGraphText`, we set weight h so that the first singular values of the graph Laplacian L and the text assisted part hC_T are equal. We choose the threshold ω to be the 95% quantile of nonzero $|W_{ij}|$'s. We set the number of random starts in the k-means steps (Step 6 and 7 in Algorithm 1) as 10^4 . For RTM, we use the function `rtm.collapsed.gibbs.sampler` in the R package `lda`. We set the scalar value of the Dirichlet hyperparameter for topic proportions $\alpha = 0.001$, the scalar value of the Dirichlet hyperparameter for topic multinomials $\eta = 0.1$, the numeric of regression coefficients expressing the relationship between each topic and the probability of link $\beta = (0.5, 0.5)$, and the number of sweeps of Gibbs sampling over the entire corpus to make as `num.iterations = 1e4`. We set α to be small since we aim to cluster each document to one topic instead of multiple topics. We set the `num.iterations` large enough so that the likelihood from each document converges.

Figure 11(a) compares the mis-clustering rate of `pairGraphText` and RTM. Without text signals (the middle plot), RTM fails to recover block labels even with large graph signals, but `pairGraphText` recovers block labels with the graph signal over 10. Without graph signals (the right plot), `pairGraphText` can only recover 90% of block labels, but RTM recovers all block labels, when the text signal is over 0.4. With both graph signals and text signals (the left plot), both methods perform better than the two cases when only one type of signals exists, and both methods can recover block labels with large enough signals.

RTM generalizes the text-based topic modeling method, LDA [Blei, Ng and Jordan (2003)], to integrate links (graph); it depends more on text and uses links to improve. On the other hand, `pairGraphText` generalizes the link-based spectral clustering to integrate text; it depends more on links and uses text to improve. From the Figure 11(a), RTM fails to recover block labels without text signals, but `pairGraphText` can still recover most block labels without graph signals. Figure 11(b) also shows that `pairGraphText` is much faster than RTM.

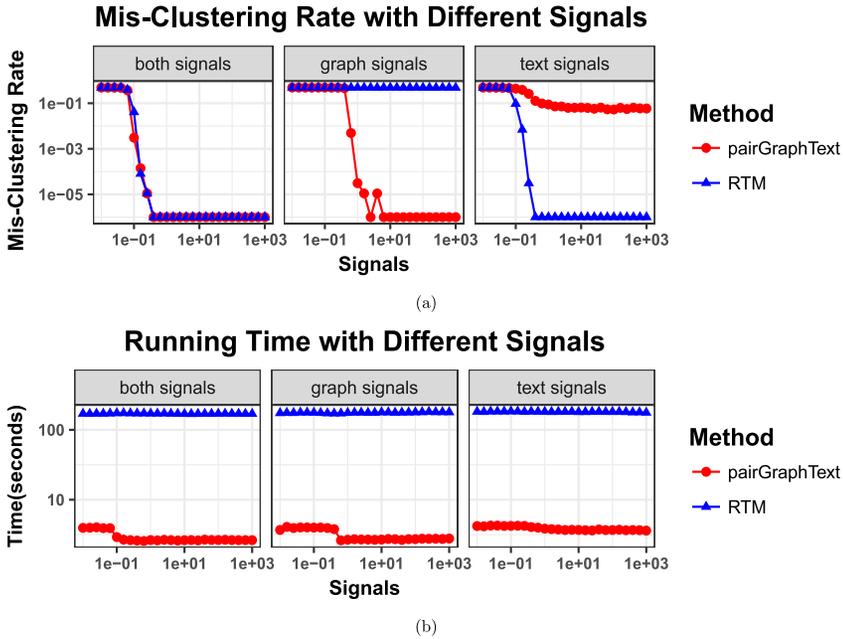


FIG. 11. Comparison between *pairGraphText* and *RTM*.

RTM enables us to predict keywords and citations for new documents [Chang and Blei (2009)]. However, to cluster massive documents into different topics, *pairGraphText* is a better choice.

See Section 4 in the supplementary material for more simulations comparing *pairGraphText* with multiple methods including *CASC* and spectral clustering.

7. Discussion. This paper searches for (i) candidate-centered structure and (ii) issue-centered structure in the political discussions on Facebook surrounding the 2012 French election. The candidate-centered structure is relatively easy to detect since we have the labels of each post belongs to which candidate. But the search for issue-centered structure is more challenging, because we have no such labels of citizens or any labels of issues. To identify topics in the discussions, we use both the graph and the text. *pairGraphText* synthesizes the graph and the text, and it addresses the noisy and high-dimensional problem for text by thresholding. Using *pairGraphText*, we identify topics that attract people’s attention, including Islam, religion, immigration, ecology, economy, and crises. During the interpretation of clusters, we propose the word-content strategy to extract the cluster topics, and our Shiny App <https://yilinzhang.shinyapps.io/FrenchElection> plays a significant role in the interdisciplinary collaboration between statisticians and social scientists. Our codes and data sets are available on Github <https://github>.

com/yzhang672/AOAS. We also provide an R package `pairGraphText` to implement our method on Github <https://github.com/yzhang672/pairGraphText>.

Chang and Blei (2010) proposed the relational topic model (RTM), a hierarchical probabilistic model for networks with node covariates. They modeled topic assignments for documents using latent Dirichlet allocation (LDA) [Blei, Ng and Jordan (2003)]. Instead of studying networks of documents or posts, we study the bipartite network between candidate-posts and citizens. Also, our method is unsupervised, more computationally efficient, and generally more accurate compared with RTM. RTM enables us to predict keywords and citations for new documents. However, to cluster documents into different topics, `pairGraphText` is a better choice than RTM.

`pairGraphText` is useful for applications outside of discussion threads. It is applicable to any network with node covariates. `pairGraphText` enhances the homogeneity of covariates within clusters. This boosts the signal of the clusters and helps with interpretation.

Acknowledgments. We thank Jonathan Chang from Physera and David Blei from Columbia University for their advice in implementing the Relational Topic Models. We thank Emma Krauska and Fan Chen from University of Wisconsin-Madison for their discussions to name `pairGraphText`. We thank the Editor, Associate Editor, and reviewer who provide helpful comments on the manuscript.

SUPPLEMENTARY MATERIAL

Supplementary Materials for “Discovering political topics in Facebook discussion threads with graph contextualization” (DOI: [10.1214/18-AOAS1191 SUPP](https://doi.org/10.1214/18-AOAS1191SUPP); .pdf). This supplementary consists of three parts. Part 1 provides more evidence for the candidate-centered structure. Part 2 explains our choice of the number of clusters K when searching for the issue-centered structure. Part 3 discusses different choices for document-term matrices. Part 4 provides more simulations comparing `pairGraphText` with RTM and other methods including CASC and spectral clustering. Part 5 provides theoretical justifications for `pairGraphText`.

REFERENCES

- ADAMIC, L. A. and GLANCE, N. (2005). The political blogosphere and the 2004 us election: Divided they blog. In *Proceedings of the 3rd International Workshop on Link Discovery* 36–43. ACM.
- AIROLDI, E. M., BLEI, D. M., FIENBERG, S. E. and XING, E. P. (2008). Mixed membership stochastic blockmodels. *J. Mach. Learn. Res.* **9** 1981–2014.
- BAKSHY, E., MESSING, S. and ADAMIC, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science* **348** 1130–1132. [MR3380630](https://doi.org/10.1126/science.1262516)
- BINKIEWICZ, N., VOGELSTEIN, J. T. and ROHE, K. (2017). Covariate-assisted spectral clustering. *Biometrika* **104** 361–377. [MR3698259](https://doi.org/10.1093/biomet/asw025)

- BLEI, D. M. (2012). Probabilistic topic models. *Commun. ACM* **55** 77–84.
- BLEI, D. M., NG, A. Y. and JORDAN, M. I. (2003). Latent Dirichlet allocation. *J. Mach. Learn. Res.* **3** 993–1022.
- BOYD, D. N. and ELLISON, N. B. (2007). Social network sites: Definition, history, and scholarship. *J. Comput.-Mediat. Commun.* **13** 210–230.
- CHANG, J. and BLEI, D. (2009). Relational topic models for document networks. In *Artificial Intelligence and Statistics* 81–88.
- CHANG, J. and BLEI, D. M. (2010). Hierarchical relational models for document networks. *Ann. Appl. Stat.* **4** 124–150. MR2758167
- CHOY, M., CHEONG, M. L., LAIK, M. N. and SHUNG, K. P. (2011). A sentiment analysis of Singapore presidential election 2011 using Twitter data with census correction. Preprint. Available at arXiv:1108.5520.
- COLLEONI, E., ROZZA, A. and ARVIDSSON, A. (2014). Echo chamber or public sphere? Predicting political orientation and measuring political homophily in Twitter using big data. *J. Commun.* **64** 317–332.
- GONZALEZ-BAILON, S., KALTENBRUNNER, A. and BANCHS, R. E. (2010). The structure of political discussion networks: A model for the analysis of online deliberation. *J. Inf. Technol.* **25** 230–243.
- GRIMMER, J. and STEWART, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Polit. Anal.* **21** 267–297.
- HEBSHI, S. and O’GARA (2011). The rohe of online social networking in the 2008 democratic presidential primary camps. Preprint. Available at <http://www.shoshanahebshi.com/wp-content/uploads/2011/08/Social-Medias-role-in-primary-campaigns.pdf>.
- HOLLAND, P. W., LASKEY, K. B. and LEINHARDT, S. (1983). Stochastic blockmodels: First steps. *Soc. Netw.* **5** 109–137. MR0718088
- JOACHIMS, T. (1996). A probabilistic analysis of the Rocchio algorithm with TFIDF for text categorization. Technical report, Carnegie-Mellon Univ., Pittsburgh, PA, Dept. of Computer Science.
- KAPLAN, A. M. and HAENLEIN, M. (2010). Users of the world, unite! The challenges and opportunities of social media. *Bus. Horiz.* **53** 59–68.
- KARRER, B. and NEWMAN, M. E. J. (2011). Stochastic blockmodels and community structure in networks. *Phys. Rev. E* **83** 016107. MR2788206
- KIM, Y. M. (2009). Issue publics in the new information environment: Selectivity, domain specificity, and extremity. *Communic. Res.* **36** 254–284.
- KREISS, D. and MCGREGOR, S. C. (2018). Technology firms shape political communication: The work of Microsoft, Facebook, Twitter, and Google with campaigns during the 2016 US presidential cycle. *Polit. Commun.* **35** 155–177.
- KUSHIN, M. J. and KITCHENER, K. (2009). Getting political on social network sites: Exploring online political discourse on Facebook. *First Monday* **14** 11-2.
- PANG, B. and LEE, L. (2008). Opinion mining and sentiment analysis. *Found. Trends Inf. Retr.* **2** 1–135.
- PAPACHARISSI, Z. (2002). The virtual sphere: The Internet as a public sphere. *New Media Soc.* **4** 9–27.
- QIN, T. and ROHE, K. (2013). Regularized spectral clustering under the degree-corrected stochastic blockmodel. In *Advances in Neural Information Processing Systems* 3120–3128.
- RAMAGE, H. R., CONNOLLY, L. E. and COX, J. S. (2009). Comprehensive functional analysis of Mycobacterium tuberculosis toxin-antitoxin systems: Implications for pathogenesis, stress responses, and evolution. *PLoS Genet.* **5** e1000767.
- RAMOS, J. (2003). Using TF-IDF to determine word relevance in document queries. In *Proceedings of the First Instructional Conference on Machine Learning*.
- ROBERTSON, S. P., VATRAPU, R. K. and MEDINA, R. (2010). Off the wall political discourse: Facebook use in the 2008 US presidential election. *Information Polity* **15** 11–31.

- ROHE, K., QIN, T. and YU, B. (2016). Co-clustering directed graphs to discover asymmetries and directional communities. *Proc. Natl. Acad. Sci. USA* **113** 12679–12684. MR3576189
- SALTON, G., WONG, A. and YANG, C.-S. (1975). A vector space model for automatic indexing. *Commun. ACM* **18** 613–620.
- SIVIC, J. and ZISSERMAN, A. (2003). Video Google: A text retrieval approach to object matching in videos. In *Proceedings Ninth IEEE International Conference on Computer Vision*. IEEE.
- STIEGLITZ, S. and DANG-XUAN, L. (2012). Political communication and influence through microblogging—an empirical analysis of sentiment in Twitter messages and retweet behavior. In *45th Hawaii International Conference on System Science*. IEEE.
- STIEGLITZ, S. and DANG-XUAN, L. (2013). Social media and political communication: A social media analytics framework. *Soc. Netw. Anal. Min.* **3** 1277–1291.
- TUMASJAN, A., SPRENGER, T. O., SANDNER, P. G. and WELPE, I. M. (2011). Election forecasts with Twitter: How 140 characters reflect the political landscape. *Soc. Sci. Comput. Rev.* **29** 402–418.
- VON LUXBURG, U. (2007). A tutorial on spectral clustering. *Stat. Comput.* **17** 395–416. MR2409803
- WANG, H., CAN, D., KAZEMZADEH, A., BAR, F. and NARAYANAN, S. (2012). A system for real-time Twitter sentiment analysis of 2012 US presidential election cycle. In *Proceedings of the ACL 2012 System Demonstrations* 115–120. Association for Computational Linguistics.
- WATTAL, S., SCHUFF, D., MANDVIWALLA, M. and WILLIAMS, C. B. (2010). Web 2.0 and politics: The 2008 US presidential election and an e-politics research agenda. *MIS Q.* 669–688.
- WEBSTER, J. G. (2014). *The Marketplace of Attention: How Audiences Take Shape in a Digital Age*. MIT Press.
- WELLMAN, B., HAASE, A. Q., WITTE, J. and HAMPTON, K. (2001). Does the Internet increase, decrease, or supplement social capital? Social networks, participation, and community commitment. *Am. Behav. Sci.* **45** 436–455.
- WILLIAMS, C. B. and GULATI, G. J. (2009). Explaining Facebook support in the 2008 congressional election cycle. OpenSIUC Working Papers, 26.
- WILLIAMS, C. B. and GULATI, G. J. (2013). Social networks in political campaigns: Facebook and the congressional elections of 2006 and 2008. *New Media Soc.* **15** 52–71.
- WITTEN, D. M. (2011). Classification and clustering of sequencing data using a Poisson model. *Ann. Appl. Stat.* **5** 2493–2518. MR2907124
- ZHANG, Y., POUX-BERTHE, M., WELLS, C., KOC-MICHALSKA, K. and ROHE, K. (2018). Supplement to “Discovering political topics in Facebook discussion threads with graph contextualization.” DOI:10.1214/18-AOAS1191SUPP.

Y. ZHANG
K. ROHE
DEPARTMENT OF STATISTICS
UNIVERSITY OF WISCONSIN MADISON
1300 UNIVERSITY AVE
MADISON, WISCONSIN 53706
USA
E-MAIL: yilin.zhang@wisc.edu
karlrohe@stat.wisc.edu

M. POUX-BERTHE
K. KOC-MICHALSKA
AUDENCIA BUSINESS SCHOOL
COMMUNICATION AND CULTURE DEPARTMENT
1 RUE MARIVAUX
44003 NANTES
FRANCE
E-MAIL: m.poux-berthe@live.fr
kkocmichalska@audencia.com

C. WELLS
DIVISION OF EMERGING MEDIA STUDIES
BOSTON UNIVERSITY
BOSTON, MASSACHUSETTS 02215
USA
E-MAIL: cfwells@bu.edu