JOURNAL OF SOCIAL COMPUTING
ISSN 2688-5255 03/06 pp157−165
Volume 2, Number 2, June 2021
DOI: 10.23919/JSC.2021.0008

Uncovering the Online Social Structure Surrounding COVID-19

Philip D. Waggoner*, Robert Y. Shapiro, Samuel Frederick, and Ming Gong

Abstract: How do people talk about COVID-19 online? To address this question, we offer an unsupervised framework that allows us to examine Twitter framings of the pandemic. Our approach employs a network-based exploration of social media data to identify, categorize, and understand communication patterns about the novel coronavirus on Twitter. The simplest structure that emerges from our analysis is the distinction between the internal/personal, external/global, and generic threat framings of the pandemic. This structure replicates in different Twitter samples and is validated using the variation of information measure, reflecting the significance and stability of our findings. Such an exploratory study is useful for understanding the contours of the natural, non-random structure in this online space. We contend that this understanding of structure is necessary to address a host of causal, supervised, and related questions downstream.

Key words: COVID-19; networks; community detection; twitter; exploratory data analysis

1 Introduction

The effects of the COVID-19 pandemic have been vast. Each wave of the virus has wrought continued and surprising havoc on nations and people around the world, devastating incomes, disrupting education of children and young adults, unsettling the provision of medical care, and exacerbating deep-seated societal and racial disparities.

A quickly growing repository of research has aimed to understand these multifaceted consequences. For example, researchers have examined the impacts of the pandemic on the stock market[1], educational outcomes[2], and mental health[3]. Other areas of research have investigated the efficacy of governmental interventions in slowing the pandemic[4], and studied ethnic, racial, and income-based disparities in COVID-19 risks and severity[5]. Further, and especially in the early stage of the pandemic as we focus on in this paper, some have found that there are differences in attitudes toward and perceptions of the pandemic[6, 7].

Adding to the multidimensionality of the COVID-19 pandemic is the complex information environment, marked by rapid communications in the news and on social media. In particular, the diverse array of actors, network connections, and information on social media offer a window into how people are responding to the pandemic. Several studies have sought to understand the networks and accounts involved in the spread of virus-related conspiracy theories and misinformation[8]. Others have leveraged social media data in the context of the pandemic to explore discrimination[9] and general sentiment[10]. Public health research has also employed social media data to predict outbreaks using symptom reports[11] and travel patterns[12].

Though the store of COVID-19 and social media-related research is vast, we suggest that in order to fully appreciate the scope and influence of the information environment during the pandemic, we must first understand the structure of communications about the virus. That is, an understanding of how people communicate about the pandemic on social media is

• Philip D. Waggoner is with the Institute for Social and Economic Research and Policy, Columbia University, New York City, NY 10027, USA. E-mail: pdw2119@columbia.edu.
• Robert Y. Shapiro and Samuel Frederick are with the Department of Political Science, Columbia University, New York City, NY 10027, USA. E-mail: rys3@columbia.edu; sdf2128@columbia.edu.
• Ming Gong is with the Department of Quantitative Methods in the Social Sciences, Columbia University, New York City, NY 10027, USA. E-mail: ming.gong@columbia.edu.
∗ To whom correspondence should be addressed.
Manuscript received: 2021-02-20; revised: 2021-06-02; accepted: 2021-06-26

essential to understand how people react and adapt to and make sense of such crises. Thus, we propose to explore the structure of online communication concerning the virus through analysis of the texts as (big) data conveyed through social media. In our initial step of contextual mining, the structures of immediate interest are the most salient and separable (sub)topics related to the virus. These structures can provide insight into how the virus is "framed" online, i.e., the perspective or standard for judgment that stands out in social media communications, in line with classical work on media and framing[13]. While our exploratory effort will not fully determine these, it is an important first step in making sense of the complexity of social media communication and in providing a deeper understanding of the pandemic's multifaceted and widespread effects.

2 Data & Method

Our data consist of thousands of tweets related to COVID-19 or the coronavirus. To get these noisy data into analyzable format, we first preprocessed, then staged the text. With a cleaned and staged corpus, we transformed the data to be passed to several network models. Using the transformed dataset, we leveraged three community detection algorithms to explore the structure (modularity) of the online Twitter space relating to COVID-19. These steps are detailed in the following subsections.

2.1 Preprocessing the text

We began with a corpus of 8.4 million tweets posted between late March and April 2020, when the virus reached its initial peak in America. Due to computational expense and to pursue greater computational efficiency, we drew and operated on four random samples of 40 000 English tweets from this initial dataset, producing a total of 160 000 tweets for analysis parsed across four corpora. In addition to increased computational efficiency, this approach allows for an informal validation check throughout where multiple samples and multiple fits of models are directly compared and contrasted. The expectation is similarity across samples given the size of the random samples of tweets. This and other validation efforts are described in depth later in the paper. Of note, these data were originally scraped using the rtweet package, and are archived at Kaggle[14, 15]. The corpus includes tweets selected based on use of at least one of the following hashtags #coronavirus #coronavirusoutbreak, #coronavirusPandemic, #covid19, #covid_19, #epitwitter, #ihavecorona, #StayHomeStaySafe (post 4/10/20), and #TestTraceIsolate (post 4/10/20).

With our raw random samples, we then preprocessed each of the four corpora of 40 000 tweets identically. First, we removed common English stopwords, punctuation, numbers, URLs, and also set text to lower case. Twitter is a notoriously noisy source of data given "linguistic noise, … message brevity, … and lack of labeled" text[16] . For example, tweets often contain special characters or terms that do not contribute substantive linguistic meaning making preprocessing a challenge requiring frequent inspection and character/term dropping (e.g., ♥, ●, ♂, ♀, "emic", and so on)[§]. We rinse and repeat this iterative cleaning process until we obtain a cleaned corpus of substantively useful content relating to the phenomenon of interest, which in our case is COVID-19/coronavirus, broadly defined. Finally, we stripped all of the white space left behind from the text cleaning in line with text mining best practices[17]. See a sample of some of the most frequently occurring words from the cleaned corpora in Table 1.

2.2 Staging and transforming the text

We staged each of the cleaned corpora in three steps. First, we built a Term Document Matrix (TDM) with terms as rows, tweets as columns, and elements as frequencies of terms in tweets. Second, we transformed the TDM into a Term-Term Matrix (TTM), with terms in rows and columns giving term combinations across all tweets. Finally, with our data in the form to allow for connections between terms used, we transformed the TTM into an adjacency matrix. An adjacency matrix, A_{ij}, is a square matrix where columns and rows act as vertices (or "nodes") of the network. In our case vertices are terms, allowing for an understanding of term usage of multiple terms across tweets, where 1 = connected, and 0 ≠ connected vertices for vertices, v_i and v_j. Edges connecting vertices capture term frequencies. Thus, we encode A_{ij},

Table 1 Frequently occurring words.

Outbreak	Help	Virus
Lockdown	Home	Trump
Deaths	Support	World

[§] Of note, given that we pre-filtered to collect only tweets using a COVID-related hashtag, we dropped related words (COVID, coronavirus, etc.) to focus on words used in association with COVID-19 and the coronavirus, instead of the terms themselves to give a clearer sense of how people talk about the virus online, given that we know they are talking about it to begin with.

$$A_{ij} = \begin{cases} 1, & \text{if } v_i \text{ and } v_j \text{ are connected;} \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

2.3 Networks and community detection

To leverage community detection to uncover structure, we need to first build a network. With A_{ij}, we built an undirected and weighted network, $G = (V, E)$, to explore the structure of the COVID-19 Twitter space, where V is defined as the full set of vertices, $v_i \in \{1, \ldots, V\}$, and E is defined as the full set of edges, $e_i \in \{1, \ldots, E\}$. The edges connect all vertices, $\{v_i \text{ and } v_j, \forall i \neq j\}$. Of note, given the size and noisiness of the data, we limited the terms in the network to those that were mentioned at least 1250 times across the 160 000 tweets in our corpora. The motivation behind this decision was to home in on terms that are not only more frequently used, but also more frequently used with each other, giving greater clarity on the structure that underlies the COVID-19 Twitter space. By winnowing the space in such a way to explore the more important and frequently co-occurring words, the resultant networks in Figs. 1 and 2 are clearer with fewer nodes and edges.

Given $G = (V, E)$, we leveraged a suite of local community detection algorithms to more explicitly explore the structure of this space, as well as the contours of this topology. Structure is defined by communities (i. e., "modules" or "clusters") of vertices in a network that are densely connected to each other, while retaining sparse connections between communities. Density in our context suggests people in one module are using similar terms to talk about COVID-19, relative to other densely connected communities in other modules which use unique terms to discuss COVID-19, and so on for all communities found in the network. By being local, then smaller groups of vertices (subgraphs) are considered on an iterative basis.

The local community detection algorithms we used are greedy optimization of modularity[18], Louvain[19], and walktrap[20]. The first algorithm uses a greedy search to look for similarties across vertices. Rather than exhaustively searching the entire space, the greedy optimization of modularity approach to detecting communities is based on a hierarchical structure of progressively similar vertices. The agglomerative (or "b

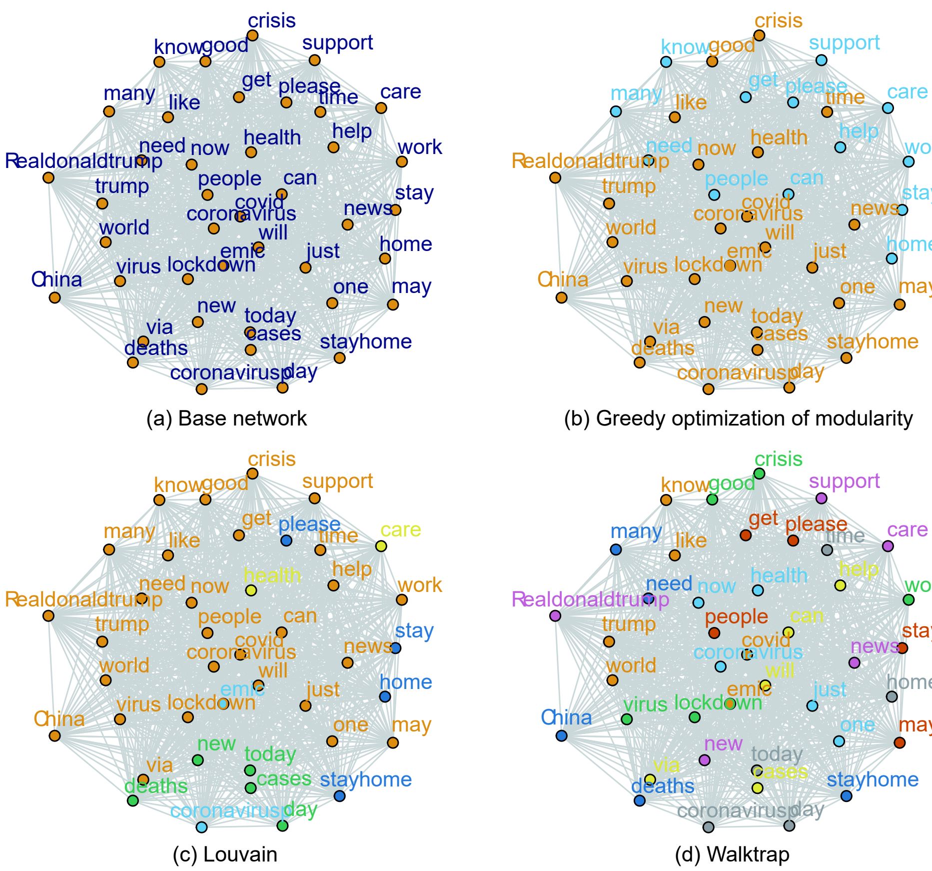

(a) Base network

(b) Greedy optimization of modularity

(c) Louvain

(d) Walktrap

Fig. 1 Community detection across three algorithms for Sample 1.

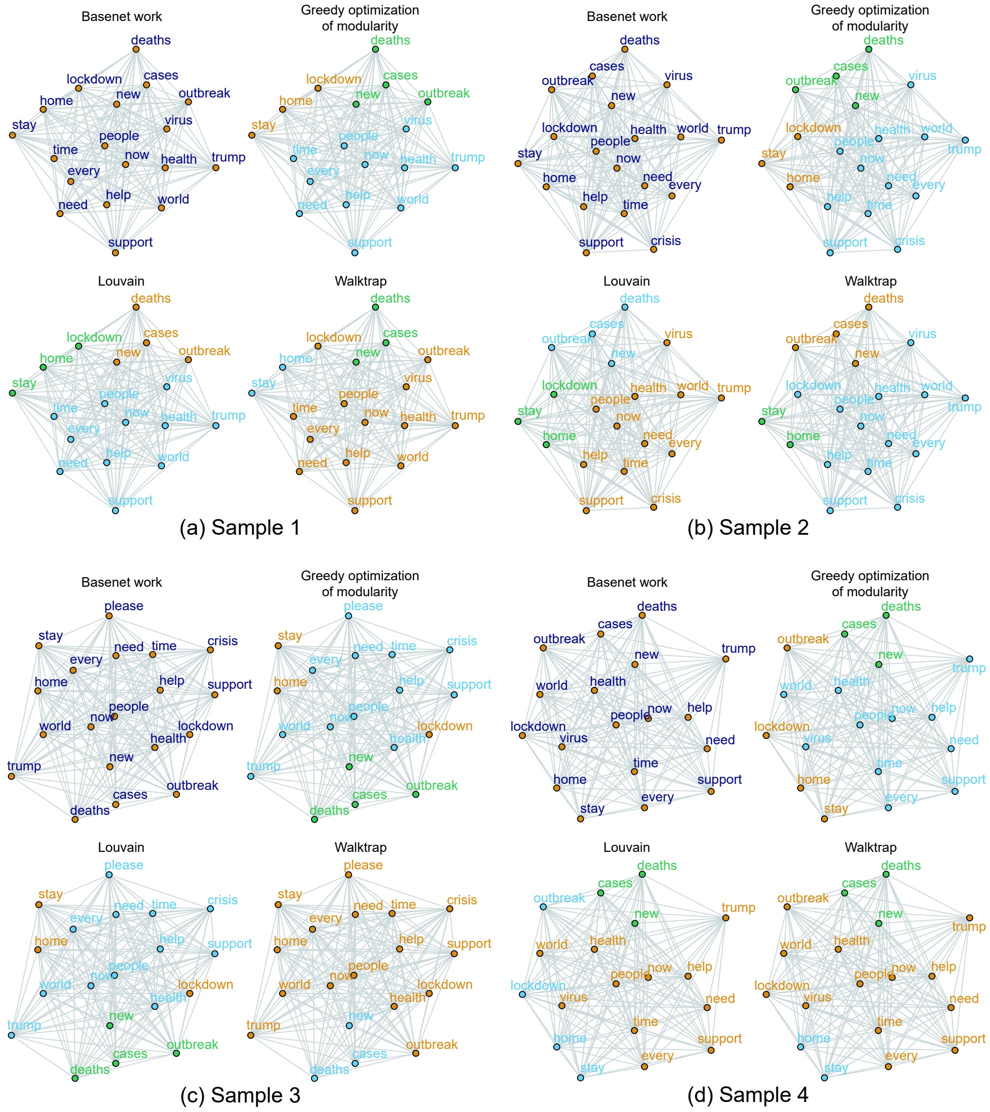

Fig. 2 Community detection across three algorithms and four samples.

ottom-up") approach allows for structure to emerge on the basis of similarity, and by implication and dissimilarity. The process is very similar to hierarchical agglomerative clustering[21]. The goal is maximal intra-module density and also inter-module sparsity. The second algorithm, Louvain, seeks to learn clustering structure locally and is also greedy. But the innovation of the Louvain algorithm is based in an additional step of building a subgraph based on the first iteration (formerly global modularity), and iteratively finding new communities to maximize modularity. Finally, the third algorithm, walktrap, seeks to learn clustering in a local way and is very similar to the prior two algorithms. The key difference is in constructing the communities, the walktrap algorithm takes a series of short random walks between vertices, v_*. Shorter walks signal a higher probability that the vertex is in a community with v_i and v_j via local modularity, and longer walks suggest it likey belongs to a different community. In sum, these local and greedy approaches to community detection offer extremely computationally efficient approaches to searching a space, and uncovering structure in a network.

Finally, after the network and community detection algorithms, we validate the community detection results in several ways in line with best practices for statistical validation of community detection algorithms[22, 23]. Specifically, Refs. [22, 23] offer several approaches to validate the robustness and significance of structure uncovered from community detection. In line with their recommendation, we validate results in two ways and present the validation results in the section following the results: comparing structure to random noise, and then comparing each algorithm to each other.

3 Finding

As our goal is to uncover the structure of a largely unknown space, the first stage offers a starting place to disentangle how people discuss COVID-19 on Twitter. We are looking for different clusters of similar terms used to discuss COVID-19. These similarities (within clusters) and differences (between clusters) found across the results from the three community detection algorithms provide progress toward our goal of searching for structure in a data space that is still new and rapidly evolving. The strikingly consistent finding across all algorithms and samples is that three communities emerge based on language used in tweets. We pull these patterns apart in the following discussion.

Community detection results are presented in Figs. 1 and 2, with communities denoted by color, Fruchterman Reingold layout is used for consistent node placement to allow for more direct comparison across all algorithms.

First, Fig. 1 includes results from the first of four samples, and allows for a zoomed in look at the results. Of note, the three communities are evident. One community contains words lockdown, home, and stay. Another community contains words deaths, new, cases, and outbreak. And the third community contains many words such as people, world, need, and help. These communities are largely consistent across the different algorithms. The greatest difference is in the walktrap algorithm (as shown in Fig. 1d). But the greedy optimization (Fig. 1b) and Louvain (Fig. 1c) algorithms show identical communities of words. See Table 2 for a complete view of the words in each community by algorithm.

Community 1 focuses on the personal domain with terms "home", "stay", and "lockdown" (for two of the three algorithms), whereas Community 2 focuses externally on the non-personal, or global domain with frequent co-occurrence of terms like "support", "world", "trump", and so on. Finally, Community 3 focuses more generically on terms associated with the threat that comes from COVID-19 with terms like "cases", "deaths", and "new". It is striking that the configuration of words in each community is highly stable, with only two words differing across communities ("lockdown" and "outbreak"). The differences are consistently with the walktrap algorithm. The greedy optimization of modularity and Louvain algorithms, though, are identical in word configurations. Taken together, the stability in communities suggests that people discuss COVID-19 online in either personal (Community 1), global (Community 2), or threat (Community 3) terms.

Zooming out, we can see this deep consistency across all three algorithms, as well as all samples of tweets in Fig. 2. The key take away is the same as that previously discussed following Table 2 , which is that three communities were found across all samples and all algorithms, with the sole exception of the walktrap algorithm in Sample 3 (Fig. 2c).

Further, the configuration of words in each community are highly consistent. In many cases, the words in each community are identical (e.g., "stay" and "home" appearing together in all iterations corresponding with the personal community/ Community 1).

In sum, the framing of COVID-19 in online/social media discussion is quite consistent, where people seem to be discussing COVID-19 in one of three ways: personal (Community 1), global (Community 2), or generic threat (Community 3). Such frames could be further interpreted as signaling devices. It is frequently noted that social media are used as an avenue to make signals of many kinds, whether mobilization efforts[24], message delivery[25], or political preferences[26]. And though it was found partisan differences in attitudes toward the pandemic in the earlier days of COVID-19[6], partisan division and affective polarization have only continued to heighten and intensify[27, 28]. Presumably these attitudes had not reached a level that could be picked up in our analysis of some of the earliest days of

Table 2 Words by community across algorithms.

Community	Greedy optimization	Louvain	Walktrap
1	home, stay, **lockdown**	home, stay, **lockdown**	home, stay
2	now, time, people, help, every, need, support, world, health, virus, trump	now, time, people, help, every, need, support, world, health, virus, trump	now, time, people, help, every, need, support, world, health, virus, trump, **outbreak**, **lockdown**
3	cases, deaths, new, **outbreak**	cases, deaths, new, **outbreak**	cases, deaths, new

Note: Words that are not identical across all three algorithms are in bold.

the pandemic. More research on the partisan aspect, as well as how and whether these frames have evolved in the later stages of the pandemic, especially as more pandemic-focused policymaking has occurred, would be a useful follow up to our work.

Regardless, we can see from the results that the framing of COVID-19 on social media is also in line with presentation of a signal. That signal could be urging people to focus internally/personally and keep others safe (Community 1/personal), or using social media as an outlet to reach and discuss COVID-19 in global terms and on a global scale (Community 2), or as a reporting mechanism, such as querying and describing the number of new cases, outbreaks, and deaths due to COVID-19 (Community 3). Such an understanding of Community 3, especially given our coverage of the earlier days of COVID-19, could be using social media to signal a generic threat from COVID-19, distinct from a personal or global casting of impact. In short, our results offer a foundational starting place to contextualize and categorize linguistic trends and patterns of discussing COVID-19 in an online environment.

4 Validation

This section focuses on validating findings to this point. Beyond the near identical community configurations of words relating to the different spheres, whether personal, global, or descriptive, we turn now to offer several formal checks of validation of these patterns across all algorithms. We proceed with two checks in this vein: first, comparing patterns across each algorithm to patterns found from random noise; and second, comparing stability of communities across each algorithm to each other.

Validation of community detection results is an important part of any community detection analysis, given the potential for different configurations of communities from different algorithms or from different samples of data. We focus our validation efforts on a recent measure, Variation of Information (VI)[29], which is defined as the information lost balanced against the information gained from two partitions of a single graph. As summarized in Ref. [30], VI balances the entropy for each cluster/module, C and C', in a common data space,

$$\mathrm{VI}(C,C') = H(C|C') + H(C'|C) \tag{2}$$

where $H(C)$ defines the entropy associated with a given cluster C,

$$H(C) = -\sum_{k=1}^{K} P(k)\log P(k) \tag{3}$$

where $k \equiv C_k$, and $P(k)$ is the probability an observation belongs to a given cluster, $k \in K$.

Building on this framework defining VI, the first approach to validation is to compare VI for the clustering/modularity found from each algorithm compared to clustering found from a random graph, but "with the same degree distribution of the original graph, but with completely random edges"[23]. Regarding interpretation, "low values represent more similar clusters and high values represent more different clusters"[23]. See the results for each algorithm compared to the random/null version across various perturbations of the graph in Fig. 3.

Perturbations, as defined in Ref. [30], are different size changes to the original network. Large percentages of perturbations correspond to larger changes to the original network. The idea is to explore whether modularity/clustering changes across different versions of the original network. This provides a useful baseline to compare to the real communities discovered in the analysis described in the previous section.

Notice in Fig. 3, across all algorithms, the curves for the real communities are higher than those for the random version at various perturbations of the original network. Taken with the comparison across each algorithm and in line with Ref. [23], this suggests that clusters found in the greedy optimization and walktrap algorithms are largely unstable and near 50%, meaning the "found community structure is a result of chance fluctuations and it is not plausible"[30]. Thus, with the significantly lower VI scores for the Louvain algorithm in Fig. 3b, we might conclude at this point that the clusters found from the Louvain algorithm are more stable and trustworthy across many versions of the original network and in comparison to the other algorithms. Yet, recall that the results from the Louvain and greedy optimization algorithms were nearly idenitcal across all four samples. Thus, this suggests that while the Louvain algorithm is more efficient and trustworthy, we can also trust the results from the greedy optimization algorithm, given the similarity with Louvain. Yet, it seems that the walktrap algorithm is the least trustworthy at this point given the differences in both word configurations from Figs. 1 and 2, compared to the more reliable results from the Louvain algorithm.

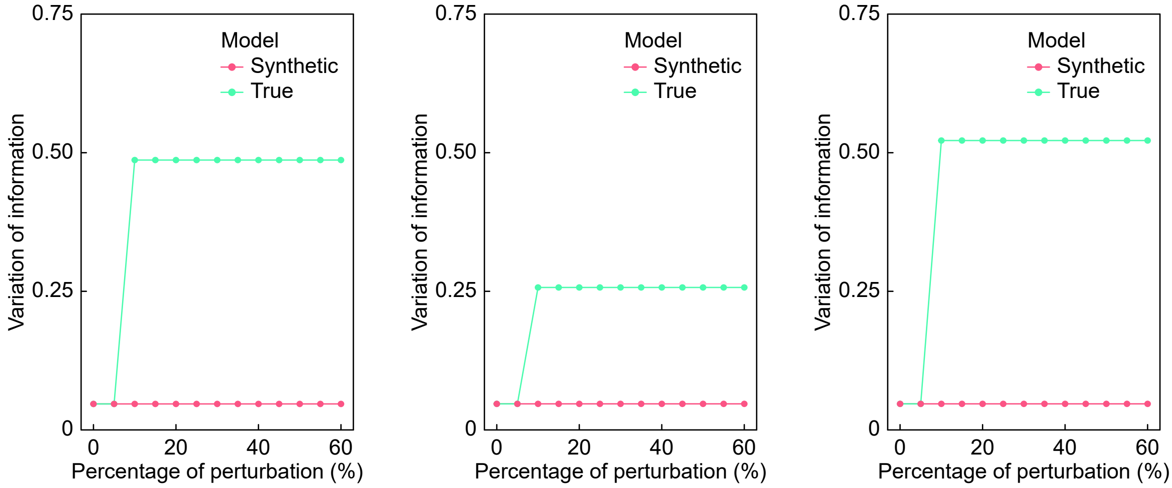

(a) Greedy optimization (b) Louvain (c) Walktrap

Fig. 3 Community detection across three algorithms for real communities and synthetic (random) version.

To deepen our validation efforts, we turn lastly to directly compare VI scores across all algorithms and at various perturbations of the original network. The results are presented in Fig. 4.

Of note, we can see more directly that the Louvain algorithm has consistently lower VI scores across multiple perturbations. This offers direct evidence of the increased efficiency and trustworthiness of the Louvain algorithm (and thus the word configuration results from the main analysis) compared to the other two algorithms. And recall, the core finding from the Louvain (and greedy optimization) algorithm(s), was the stable presence of three communities across all four samples, showing that people tend to discuss COVID-19 online under one of three frames: personal, global, or generic threat.

5 Concluding Remark

Our study at this juncture is an exploratory one that allows for an unsupervised and assumption-free look at noisy Twitter data in the context of an ongoing, rapidly-developing, and complex global pandemic. Our task is to uncover the ways in which people frame COVID-19 online, focusing on the earlier days of the pandemic.

The simplest structure that emerges is a distinction between the internal/personal, external/global, and generic threat dimensions of the pandemic. As a result, we have built a general, but consistent framework to understand and categorize differences in Twitter patterns. Such an exploratory study helps by providing

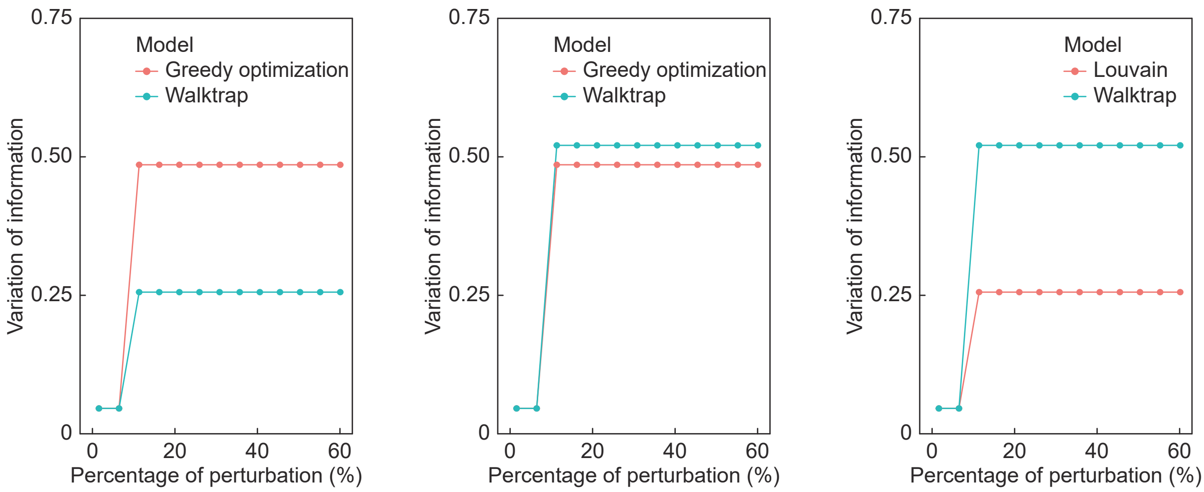

(a) Greedy optimization vs. Louvain (b) Greedy optimization vs. Walktrap (c) Louvain vs. Walktrap

Fig. 4 Community detection across three algorithms.

a step in understanding the contours of the natural, non-random structure in this online space.

We are hopeful our research will act as a launching place for other similarly situated studies to go deeper in pulling this structure apart, especially as more overt public opinion studies relating to the pandemic begin to surface, e.g., Ref. [28]. Such future work would allow for addressing related questions, such as a supervised task of measuring sentiment, explicitly probing the political aspects of discussing COVID-19 on Twitter, and also alternative approaches to uncovering communities in this space (e.g., global instead of local approaches to community detection).

Further, we encourage researchers to build on our findings by exploring nuance in these online discussion networks. Specifically, we take the tweets at face value in this research. But there is room to expect fake news, spam, and other malicious uses of Twitter to be impacting the discussion space at some level[31, 32]. Future work aimed explicitly at these and related topics would provide valuable extensions of our work. Such studies would also allow for a deeper dive into the personal, global, and descriptive structure we uncovered, but explicitly accounting for the potential bias flowing from misinformation and fake news on Twitter.

Finally, a common problem with network studies of this sort is relying on hashtags to filter the data space, as hashtags are complex[33], and can be attempts to gain popularity[34], rather than signaling genuine discussion. While we made no assumptions of motivation behind the formation of tweets on the part of the user, future work might consider parsing and exploring Twitter and relate online discussion data in a different way, such as tweets or discussions among specific communities like academia, finance, or government.

References

[1] S. R. Baker, N. Bloom, S. J. Davis, K. Kost, M. Sammon, and T. Viratyosin, The unprecedented stock market reaction to COVID-19, *The Review of Asset Pricing Studies*, vol. 10, no. 4, pp. 742–758, 2020.

[2] T. Gonzalez, M. A. de la Rubia, K. P. Hincz, M. ComasLopez, L. Subirats, S. Fort, and G. M. Sacha, Influence of COVID-19 confinement on students' performance in higher education, *PLoS One*, vol. 15, no. 10, p. e0239490, 2020.

[3] D. Y. Li, H. Chaudhary, and Z. Zhang, Modeling spatiotemporal pattern of depressive symptoms caused by COVID-19 using social media data mining, *Int. J. Environ. Res. Public Health*, vol. 17, no. 14, p. 4988, 2020.

[4] S. Flaxman, S. Mishra, A. Gandy, H. J. T. Unwin, T. A. Mellan, H. Coupland, C. Whittaker, H. Zhu, T. Berah, J. W. Eaton, et al., Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe, *Nature*, vol. 584, no. 7820, pp. 257–261, 2020.

[5] M. A. Raifman and J. R. Raifman, Disparities in the population at risk of severe illness from COVID-19 by race/ethnicity and income, *Am. J. Prev. Med.*, vol. 59, no. 1, pp. 137–139, 2020.

[6] S. K. Gadarian, S. W. Goodman, and T. B. Pepinsky, Partisanship, health behavior, and policy attitudes in the early stages of the COVID-19 pandemic, *PLoS One*, vol. 16, no. 4, p. e0249596, 2021.

[7] M. K. Gusmano, E. A. Miller, P. Nadash, and E. J. Simpson, Partisanship in initial state responses to the COVID-19 pandemic, *World Med. Health Policy*, vol. 12, no. 4, pp. 380–389, 2020.

[8] R. Kouzy, J. Abi Jaoude, A. Kraitem, M. B. El Alam, B. Karam, E. Adib, J. Zarka, C. Traboulsi, E. W. Akl, and K. Baddour, Coronavirus goes viral: Quantifying the COVID-19 misinformation epidemic on twitter, *Cureus*, vol. 12, no. 3, p. e7255, 2020.

[9] H. Budhwani and R. Y. Sun, Creating COVID-19 stigma by referencing the novel coronavirus as the "Chinese virus" on twitter: Quantitative analysis of social media data, *J. Med. Internet Res.*, vol. 22, no. 5, p. e19301, 2020.

[10] M. O. Lwin, J. H. Lu, A. Sheldenkar, P. J. Schulz, W. Shin, R. Gupta, and Y. P. Yang, Global sentiments surrounding the COVID-19 pandemic on twitter: Analysis of twitter trends, *JMIR Public Health Surveill.*, vol. 6, no. 2, p. e19447, 2020.

[11] B. A. Panuganti, A. Jafari, B. MacDonald, and A. S. DeConde, Predicting COVID-19 incidence using anosmia and other COVID-19 symptomatology: Preliminary analysis using google and twitter, *Otolaryngol. Head Neck Surg.*, vol. 163, no. 3, pp. 491–497, 2020.

[12] D. Bisanzio, M. U. G. Kraemer, T. Brewer, J. S. Brownstein, and R. Reithinger, Geolocated twitter social media data to describe the geo-graphic spread of SARS-CoV-2, *J. Travel Med.*, vol. 27, no. 5, p. taaa120, 2020.

[13] S. Iyengar, *Is Anyone Responsible? How Television Frames Political Issues*. Chicago, IL, USA: University of Chicago Press, 1994.

[14] S. Smith, Coronavirus (COVID-19) tweets – early April, https://www.kaggle.com/smid80/coronavirus-covid19-tweetsearly-april, 2020.

[15] S. Smith, Coronavirus (COVID-19) tweets – late April, https://www.kaggle.com/smid80/coronavirus-covid19-tweets-late-april, 2020.

[16] L. Derczynski, A. Ritter, S. Clark, and K. Bontcheva, Twitter part-of-speech tagging for all: Overcoming sparse and noisy data, in *Recent Advances in Natural Language Processing*, Hissar, Bulgaria, 2013, pp. 198–206.

[17] T. Kwartler, *Text Mining in Practice with R*. Chichester, UK: John Wiley & Sons, 2017.

[18] A. Clauset, M. E. J. Newman, and C. Moore, Finding community structure in very large networks, *Phys. Rev.E*, vol. 70, no. 6, p. 066111, 2004.

[19] V. D. Blondel, J. L. Guillaume, R. Lambiotte, and E. Lefebvre, Fast unfolding of communities in large networks, *J. Stat. Mech.*: *Theory Exp.*, vol. 2008, no. 10, p. P10008, 2008.
[20] P. Pons and M. Latapy, Computing communities in large networks using random walks, in *Proc. 20th Int. Symp. Computer and Information Sciences*, Istanbul, Turkey, 2005, pp. 284–293.
[21] P. D. Waggoner, *Unsupervised Machine Learning for Clustering in Political and Social Research*. Cambridge, UK: Cambridge University Press, 2020.
[22] B. Karrer, E. Levina, and M. E. J. Newman, Robustness of community structure in networks, *Phys. Rev.E* , vol. 77, no. 4, p. 046119, 2008.
[23] V. Policastro, D. Righelli, A. Carissimo, L. Cutillo, and I. De Feis, ROBustness in network (Robin): An R package for comparison and validation of communities, arXiv preprint arXiv: 2102.03106, 2021.
[24] B. Enjolras, K. Steen-Johnsen, and D. Wollebaek, Social media and mobilization to offline demonstrations: Transcending participatory divides? *New Media Soc.*, vol. 15, no. 6, pp. 890–908, 2013.
[25] S. Muralidharan, L. Rasmussen, D. Patterson, and J. H. Shin, Hope for Haiti: An analysis of facebook and twitter usage during the earthquake relief efforts, *Public Relat. Rev.*, vol. 37, no. 2, pp. 175–177, 2011.
[26] A. Ceron, L. Curini, S. M. Iacus, and G. Porro, Every tweet counts? How sentiment analysis of social media can improve our knowledge of citizens' political preferences with an application to Italy and France *New Media Soc.*, vol. 16, no. 2, pp. 340–358, 2014.
[27] J. Clinton, J. Cohen, J. Lapinski, and M. Trussler, Partisan pandemic: How partisanship and public health concerns affect individuals' social mobility during COVID-19, *Sci. Adv.*, vol. 7, no. 2, p. eabd7204, 2021.
[28] J. N. Druckman, S. Klar, Y. Krupnikov, M. Levendusky, and J. B. Ryan, Affective polarization, *local contexts and public opinion in America*, *Nat. Hum. Behav.*, vol. 5, no. 1, pp. 28–38, 2021.
[29] M. Meilă, Comparing clusterings—An information based distance, *J. Multivariate Anal.*, vol. 98, no. 5, pp. 873–895, 2007.
[30] A. Carissimo, L. Cutillo, and I. De Feis, Validation of community robustness, *Comput. Stat. Data Anal.*, vol. 120, pp. 1–24, 2018.
[31] G. K. Shahi, A. Dirkson, and T. A. Majchrzak, An exploratory study of COVID-19 misinformation on twitter, *Online Soc. Netw. Media*, vol. 22, p. 100104, 2021.
[32] K. C. Yang, C. Torres-Lugo, and F. Menczer, Prevalence of low-credibility information on twitter during the COVID-19 outbreak, arXiv preprint arXiv: 2004.14484, 2020.
[33] O. Tsur and A. Rappoport, What's in a hashtag?: Content based prediction of the spread of ideas in microblogging communities, in *Proc. 5th ACM Int. Conf. Web Search and Data Mining*, Seattle, WA, USA, 2012, pp. 643–652.
[34] J. Y. Huang, Y. Y. Tang, Y. Hu, J. J. Li, and C. J. Hu, Predicting the active period of popularity evolution: A case study on twitter hashtags, *Inf. Sci.* , vol. 512, pp. 315–326, 2020.

Philip D. Waggoner is a research scholar at the Institute for Social and Economic Research and Policy (ISERP), Columbia University, USA. His research interests are in computational social science, machine learning, and statistical computing. For more visit his website: https://pdwaggoner.github.io/.

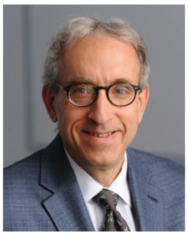

Robert Y. Shapiro is the Wallace S. Sayre Professor of Government and a professor of international and public affairs at Columbia University, USA. His research interests are in political science, public opinion, and public policy. For more visit his website: https://www.sipa.columbia.edu/faculty-research/faculty-directory/robert-shapiro.

Samuel Frederick is a PhD candidate at the Department of Political Science, Columbia University, USA. His research interests are in American politics and social media data.

Ming Gong is a master student at the Department of Quantitative Methods in the Social Sciences, Columbia University, USA. Her research interests are in computational social science and natural language processing.