Influencing the Influencers: Evaluating Person-to-Person Influence on Social Networks Using Granger Causality *

Richard Kuzma, Iain J. Cruickshank, and Kathleen M. Carley

CASOS Institute, Carnegie Mellon University,
{rkuzma2, icruicks}@andrew.cmu.edu,
kathleen.carley@cs.cmu.edu
http://www.casos.cs.cmu.edu/

Abstract. We introduce a novel method for analyzing person-to-person content influence on Twitter. Using an Ego-Alter framework and Granger Causality, we examine President Donald Trump (the Ego) and the people he retweets (Alters) as a case study. We find that each Alter has a different scope of influence across multiple topics, different magnitude of influence on a given topic, and the magnitude of a single Alter's influence can vary across topics. This work is novel in its focus on person-to-person influence and content-based influence. Its impact is two-fold: (1) identifying "canaries in the coal mine" who could be observed by misinformation researchers or platforms to identify misinformation narratives before super-influencers spread them to large audiences, and (2) enabling digital marketing targeted toward upstream Alters of super-influencers.

Keywords: Influence, Social Media, Granger Causality

1 Introduction

Influence in social networks remains one of the mainstays of social network analysis. Recently, with the advent of online social networks and the influence that online misinformation and disinformation can have, social influence has taken on added importance and dimensions. In this work, we build upon the foundations of social influence by examining social influence at the Ego network level. Using an Ego-Alter framework we identify an Ego (Donald Trump), examine which Twitter accounts he retweets most to identify key Alters (e.g. Ivanka Trump,

---

Linsdey Graham), identify tweet topics from the Ego using unsupervised machine learning, build a supervised machine learning classifier to classify tweets of Alters into those topics [2,3]. Then for a given topic, we measured the influence of each Alter on the Ego by measuring the Granger causality on the Alter's tweet time series on the Ego's tweet time series.

We devised this method in order to answer the question, "who influences the influencers?" Given the role that social media super-influencers may play in the spread of (mis)information online, identifying who influences them, what topics they are influenced on, and how much they are influenced could be key to future misinformation studies.
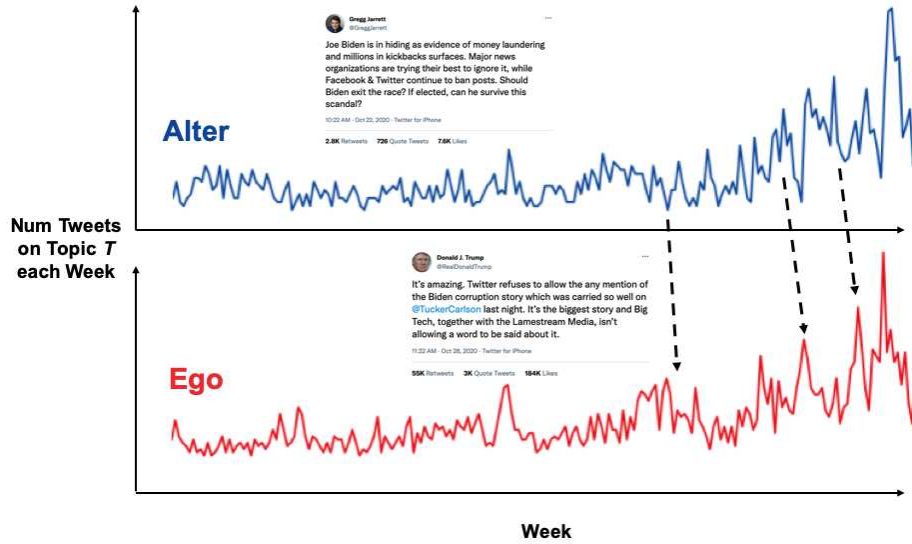

Fig. 1. Example of an Alter's Tweets affecting the Ego's Tweets for a given topic. Certain topics in the Alter's tweets are then Tweeted, with a time-delay by the Ego.

Our work is novel from other applications of Granger causality for a few reasons. Most Granger causality research on social media focuses on (1) population-level influence, (ours focuses on person-to-person influence); (2) sentiment analysis, (ours focuses on topics analysis); and (3) impact of social media behavior on real world events such as financial markets or protests (our work focuses on influence within online behavior on Twitter) [4,7,27]. Other work using Granger causality to examine person-to-person influence on social media examined sharing of explicit URLs between Twitter accounts, but not the content or topic of tweets as we do [35].

2   Background

The last few years have seen incredible advances in, and need for, social cybersecurity. Social Cybersecurity is an emerging field that looks at the intersection of human behavior in a cyber mediated information environment [10]. In particular, recent events like the COVID-19 pandemic and use of online disinformation in various national elections have highlighted the need to better understand the ways in which mis- and disinformation are spread online [8, 36]. Previous works have generally focused on the role of bots and trolls in their roles in spreading mis- and disinformation [5, 6, 10, 24, 26, 37]. More recently, coordination among cyborg or fully human accounts has been studied as means by which malicious actors can propagate mis- and disinformation [20, 30, 31]. The insight of coordinated spread by social media accounts is even being used to investigate cross-platform mis- and disinformation spread [34]. To date, much of the research in online mis- and disinformation campaigns has focused on particular types of accounts and coordination among accounts.

One area of the spread of online mis- and disinformation that is less well-studied is the role of online social media influencers. Social media influencers, or those online accounts with the ability to reach hundreds of thousands, if not millions, of other social media users with their content represent an important component of the social media landscape, by virtue of their reach. Recent work has shown that social media influencers can often have an outsized impact on the information quality of any given online discussion. For example, one study showed that the former U.S. President Donald Trump was responsible for over 30% of the COVID-19 misinformation on Twitter [15]. As another example, the 'dirty dozen', a collection of only twelve accounts, are responsible for the vast majority of anti-vaccination content online [14]. Additionally, recent disinformation campaigns, like Secondary Infektion have deliberately sought to target social media influencers to get them to spread disinformation [1, 17]. Thus it is clear that social media influencers can play a tremendous role in the spread of mis- and disinformation.

It is also important to note that the study of influence in social networks has a long history with much research into the topic. The study of influence in social networks pre-dates online social networks and has been used to understand everything from individuals' behaviors to collective action of groups [13, 18, 19]. Many of the models designed for understanding social influence work at the adjacency matrix-level of the social networks, and are fundamentally matrix based models [16, 28]. As such, these models tend to look more at population levels of influence. Less studied within the social network analysis community is the role and exercise of social influence of Ego networks.

An important consideration when studying influence is the role of causality. Granger causality is a statistical method that examines whether X has predictive power in describing Y [22]. X_t is said to Granger-cause Y_t if Y_{t+1} can be better explained by past values of Y_t and past values of X_t than by the presence of past values of Y_t alone.

2.1   Data and Definitions

We retrieved tweets from Donald Trump from the Trump Twitter Archive [9]. This repository contains 56,571 tweets and retweets from Donald Trump's twitter account. Reports indicate that not all tweets from the @RealDonaldTrump twitter account come directly from Donald Trump. Prior to the presidency, it was believed that tweets from an Android device were from Donald Trump while other platforms (iPhone, web) were not. Once Donald Trump became president, he switched to a specially-issued iPhone. At least one staffer was trusted with @RealDonaldTrump twitter privileges. Tweets came from Donald Trump, direct dictations to staffers, and presumably from the close, trusted staffers themselves attempting to channel Donald Trump's head space [12]. Given the importance of the Twitter audience to Donald Trump, we believe that all tweets from the account, even those from staffers instead of Trump himself, accurately reflect Donald Trump's wishes and are worth studying as if they came from Trump himself.

3   Methods

Fig. 2. Overview of Method. Our method consists of three, main steps: gathering the data and constructing the ego-later network, finding topics in the textual data and building a classifier of those topics, and then testing for causality in the ego-alter network on those topics.

Our method has three stages: gathering data, identifying topics of conversation, and determining granger causality.

3.1   Gathering Data

We gathered tweets from the Trump Twitter Archive, which retained copies of all of Trump's tweets and retweets from before Trump's Twitter suspension. CSV and JSON files with all captured tweets are available at the Trump Twitter Archive. We limited tweets to the time that Donald Trump was President of the United States and active on Twitter, January 20th, 2017 to January 8th, 2020. We hypothesized that the information Trump ingested affected the information he output. In other words, what he retweeted affected what he tweeted. We found

12 accounts from individuals Trump retweeted the most (the Alters). We excluded Trump's retweets of @WhiteHouse and of himself (@RealDonaldTrump). These Alters included family members, conservative media pundits, and members of Congress. We used SNScrape, an open-source tool, to get the tweet IDs of these 12 Alter accounts [25]. The original tweets of these accounts were available, but data was not available on who these Alters retweeted. We used these tweet IDs to get the tweet content and metadata through a process known as hydration, using Hydrator, another open-source tool [29].

3.2   Identifying Topics

We cleaned Trump's tweets by extracting URLs, mentions, and "RT" (used for manual retweets when Trump quoted another twitter user) from tweets. We also used the gensim Python library to preprocess tweets by removing punctuation and making all words lowercase. We then tokenized the tweets and found common bigrams and trigrams using the gensim library. Bigrams are two-word sequences like "crooked hillary" or "sleepy joe" and trigrams are three-word sequences like "make america great." Following tokenization, we embedded the tweets using the Universal Sentence Encoder (USE) from Google Research which was available at the TensorFlow Hub [11, 21]. We also featurized all tweets using the term frequency-inverse document frequency (TF-IDF) from the scikit learn library [32].

We chose the USE because it was built for short sequences ranging from a sentence to a short paragraph in length, which is similar to the length of a tweet, and performs well in semantic similarity tasks, determining if one sequence is similar to another. We performed k-medoids clustering [38] on the USE embeddings to identify topic clusters and within those topic groups examined the most important words by averaging the TF-IDF scores of all documents within a given USE cluster.

We chose to cluster only on the Ego's tweets because including Alter tweets in the topic classifier would dilute the Ego's tweet content with Alter content when forming clusters. With 12 Alters in our case study, it is possible topics would drift toward Alter content rather than Ego content, making a classifier predict whether a given Alter's tweet was like other Alter tweets rather than similar to the Ego's tweet. We settled on eight topics for Trump's tweets through a combination of silhouette score [33] from the k-medoid clustering of USE embeddings and interpretation of the TF-IDF top words for each cluster.

We categorized each Alter's tweets into one of the categories of Trump's tweets. To do this we train a Support Vector Classifier on Trump's tweets, using the USE embeddings as features and the unsupervised topic clusters as the label. Our final model had a weighted F-1 score of 0.91.

For each Alter, we clean the tweets using the same cleaning pipeline as Trump's tweets then embed the tweets using USE embeddings. We then use the model trained to predict Trump's tweet categories to predict the Alter's tweet categories. Now that the Alter's tweets are categorized, we turn them into eight time series, one for each of the topic clusters. We use weeks as time-steps

Topic	Top 50 words
Fake News	fake_news, report, media, fake_news_media, great, cnn, witch_hunt, watch, russia, story, bad, say, news, democrat, people, fbi, time, trump, know, corrupt, mueller, collusion, big, lie, today, look, fake, dems, election, history, thank, mueller_report, investigation, want, make, ukraine, fox_news, fact, poll, cover, rating, good, try, lamestream_media, country, like, use, fox, new, real
Congress	thank, democrat, vote, senate, dems, people, nancy_pelosi, great, republican, want, senator, american_people, time, impeachment, year, joe_biden, work, know, try, wow, country, say, need, congress, like, watch, house, nothing_democrat, think, biden, pelosi, let, right, happy_birthday, way, good, happen, make, big, chuck_schumer, pass, adam_schiff, remember, american, lose, thing, help, come, care, fight
MAGA	great, big, vote, country, china, economy, new, job, thank, people, today, make, year, want, tariff, good, win, come, deal, congratulation, work, farmer, trade, agree, best, time, democrat, strong, help, mexico, world, fight, need, way, number, look, history, pay, high, day, trade_deal, united_state, state, american, america, company, president, love, like, usa
Democrat Obstruction	democrat, impeachment, republican, trump, job, dems, want, vote, president, impeach, say, election, time, law_order, day, presidential_harassment, approval_rating_republican_party, thank, reason, president_trump, think, congress, house, thank, try, republican_party, today, make, year, big, win, justice, know, like, great, watch, need, attack, daca, people, way, work, political, case, antifa, police, crime, house_democrat, right, break, fast
Campaigning	great, state, thank, people, true, vote, pennsylvania, governor, fake_news, election, country, work, make, know, interview, today, say, president, like, job, joe_biden, live_press_brief_coronavirus_task_force, watch, north_carolina, want, fact, great_job, big, voter_fraud, ballot, allow, joe, fraud, mayor, let, good, stop, time, way, book, report, right, sleepy_joe, biden, use, story, voter, test, case, new_york
America First	president, great, people, country, democrat, support, military, help, big, american, thank, america, strong, want, day, need, good, united_state, vote, today, protect, bad, border, know, biden, world, safe, nation, republican, year, border_security, job, iran, state, fight, joe_biden, important, way, stand, win, time, make, come, leader, history, economy, veteran, work, china, peace
Praising Trump	thank, trump, president, president_trump, great, make_america_great, america, win, today, love, american, congratulation, country, donald_trump, big, vote, day, florida, way, united_state, job, people, white_house, biden, work, work_hard, proud, incredible, great_job, year, make, honor, tonight, great_honor, live, election, time, support, rally, million, best, come, welcome, like, help, happy, wisconsin, usa, melania, sign
Investigations	president, whistleblower, fbi, schiff, say, impeachment, adam_schiff, make, know, witch_hunt, democrat, case, people, ask, investigation, witness, want, lie, report, congress, evidence, release, hoax, comey, use, come, order, day, american, break, new, testify, start, think, obama, happen, hear, investigate, leak, like, time, fact, impeachment_hoax, ukraine, tell, bad, country, email, story, lawyer

Fig. 3. Top 50 words and phrases used by Donald Trump in each topic (ranked by TF-IDF score). Clusters determined by USE embeddings.

because we found days contained too few tweets to make an effective time series for each topic.

3.3   Granger Causality

We create a time-series with units of weeks on the x-axis and number of user's tweets on a given topic on the y-axis. The Ego, E, and Alters, A_1 to A_n, each each have m time series, where m is the number of topics. Then, for a given Alter (e.g. Ivanka Trump) and a given topic (e.g. Campaigns) we want to find if the Alter's tweets granger-cause (have predictive power to describe) the the tweets of the Ego (e.g. Donald Trump).

We first considered granger causality at the level of days, but found that most users did not have enough data when viewing topics at the day level. We think weeks are a viable time scale because intuition suggests that an Alter may tweet about a topic of importance to them repeatedly. In future work, it would be worth considering a 2-3 day rolling window

Testing for granger causality requires that both time series are stationary, that is, the mean, variance, and autocorrelation do not change with time. Sta-

tionary was checked with the Augmented Dickey Fuller test [23]. If a time series is not stationary, it may be made stationary by differencing the time series. This requires subtracting the number of tweets at time t from the number of tweets at time $t-1$. If the time series of the Alter's or Trump's tweets for a given topic are not stationary, both time series were differenced and stationarity re-examined. This way both time series always had the same amount of differencing.

If the time series could be made stationary, a granger causality test can then be done. Again, this test determines if Y_t (Trump's tweets on a topic) is better explained by Y_{t-n} (Trump's previous tweets on that topic) alone or Y_{t-n} and X_{t-n} (Alter's previous tweets on that same topic).

4 Results

Different Alters influence the Ego on different topics. Even among Alters whose tweets granger-cause Trump's tweets, some Alters have a more immediate effect (1 week) while others affect Trump's tweets with six to eight weeks of lag.

Members of Congress (Mark Meadows, Jim Jordan, GOP Chairwoman) had influence on the "Congress" topic with one week lag between their tweets and Trump's. Jim Jordan and GOP Chairwoman also influenced the "Investigations" topic with a one week lag.

Conservative pundits (Mark Levin, Gregg Jarrett, and Tom Fitton) influenced the most topics (6-7), though not all with immediate effect. Their tweets on "Fake News," "Campaigning,", and "Investigations" (along with pundit Charlie Kirk, though less so with Tom Fitton) had 1-3 week lags between their tweets and Trumps.

Donald Trump Jr had influence on "Fake News," "Congress," "Campaigning," and "America First" topics, all with 1-3 week lags. Ivanka Trump only influenced "America First" and "Praising Trump" but with 5-6 week lags. Eric Trump had influence on more topics, but a mix of immediate lag (1-3 week) and delayed lag (6-7 week) influence.

Among Alters that influenced five or more topics (Eric Trump, Mark LEvin, Gregg Jarrett, and Tom Fitton), the immediacy of influence was not consistent. Often the Alter had a few topics with immediate lag influence (1-2 week) and a few with delayed lag influence (7-8 week).

We chose to do individual hypothesis testing and felt the low p-values would not substantively change with multiple hypothesis testing, but future work may choose to multiple hypothesis testing for added rigor.

5 Discussion

In this paper we ask "who influences the influencers?" and introduce a novel method for examining person-to-person influence on Twitter using an Ego-Alter framework for finding which Alters influence the Ego's posts on social media, on which topics these Alters influence the Ego, and the time delay between Alter

Alter	Fake News	Congress	MAGA	Democrat Obstruction	Campaigning	America First	Praising Trump	Investigations	Avg Lag /Alter
Eric Trump	-	-	0.01 (1)	0.008 (7)	< 0.001 (3)	0.016 (1)	0.027 (6)	-	3.6
Ivanka Trump	-	-	-	-	-	0.021 (5)	0.026 (6)	-	5.5
Donald Trump Jr	0.05 (1)	0.07 (1)	-	-	0.028 (3)	0.001 (1)	-	-	1.5
Dan Scavino	-	-	-	-	-	-	-	-	-
Mark Meadows	-	0.004 (1)	-	-	-	-	-	-	1
Jim Jordan	-	0.001 (1)	-	-	0.003 (2)	-	-	0.004 (1)	1.3
Lindsay Graham	-	-	-	-	-	-	-	-	-
GOP Chairwoman	-	< 0.001 (1)	-	-	-	0.043 (4)	0.036 (4)	0.007 (1)	2.5
Mark Levin	0.031 (3)	0.044 (7)	0.002 (2)	0.038 (7)	< 0.001 (1)	-	-	0.031 (3)	3.83
Gregg Jarett	0.047 (1)	0.003 (2)	< 0.001 (5)	-	0.012 (3)	0.017 (7)	0.029 (2)	< 0.001 (1)	3
Charlie Kirk	-	-	-	0.039 (1)	-	-	-	< 0.001 (1)	1
Tom Fitton	0.005 (1)	0.023 (6)	0.027 (8)	-	-	0.002 (1)	0.038 (3)	0.017 (6)	4.17
Avg Lag /Topic	1.5	2.7	4.5	5	2.4	3.17	4.2	2.17	

Fig. 4. P-Values for granger causality test. Num lags in parenthesis (max 8 weeks). Alters shown in three groups: (1) family/close advisors, (2) members of Congress, and (3) political commentators.

posts and Ego posts. We show this specifically with tweets from Donald Trump and the first-order Alters he interacts with on Twitter, but believe this work could be extended to examine other influencers on Twitter and their networks, to include second and third-order Alters farther away from the Ego's immediate network, and attempted on other social media platforms (Parler, Facebook).

Our results show substantial variation in the influence of Alters on an Ego. While our work confirms the general theories of social influence in that social influence affects an individual's behavior, we also note what influences an individual and when they influence that individual are heterogeneous entities. Some Alters have a greater scope of influence. Of Donald Trump's eight tweet topics, some Alters influence only one or two topics while others influence seven of the eight topics. Some Alters have narrow, but powerful, influence. Charlie Kirk, a conservative talk show host, influences Trump on only two topics, "investigations" and "democrat obstruction." Trump's tweets lag Kirk's by only one week, showing that Kirk has high influence on these topics. Five other Alters influence the "investigations" topic and three of them also have a one-week lag on their influence, so Kirk is not special in this category. On "democrat obstruction," however, only two other Alters with influence. Kirk has a one-week lag between his "democrat obstruction" tweets and Trumps, while the other Alters have a seven-week lag, indicating less influence on this topic.

Different Alters influence different topics and should not be viewed as substitutable. Tom Fitton, a conservative activist and President of Judicial Watch, is very influential in the "fake news," "america first," and "praising trump" categories whereas Jim Jordan, a Congressman from Ohio, is highly influential on topics of "congress," "campaigning," and "investigations."

Some Alters influence different topics differently. This suggests that some Alters are more influential on certain topics than others. For instance, Gregg Jarrett, a conservative American television commentator, influences Trump's tweets on seven of eight topics. On "fake news" and "investigations" his tweets influence Trump's after a one-week lag. On "MAGA" and "America First" his tweets influence Trump's after a five-week and seven-week lag, respectively. Jarett would be a good influence vector for "fake news" and "investigations" but not as good for "MAGA" and "America First"

5.1   Future Work

Despite the useful results and demonstration of using granger-causality in understanding peer-to-peer influence on social media, there are limitations to this work. We examine a specific case of influence surrounding Donald Trump, without any comparison to how these Alters affected other Twitter users or how other Egos are influenced by their respective Alters . We focused on this base case because the Trump Twitter Archive permitted us to study the retweets of the Ego to identify the most-retweeted Alters. Generally, retweets of specific Twitter accounts are not available through the Twitter API nor are they available through Python libraries like SNScrape, making data collection difficult. If retweets of Alters were available, research could be done on whether second-order Alters (e.g. someone who Ivanka Trump retweets a lot) has influence on the Ego (Donald Trump) through the first-order Alter (Ivanka Trump).

We see two uses of this influence work. First, in the disinformation space, using Alters as "canaries in the coal mine" who could be observed by misinformation researchers or platforms to notice new misinformation narratives before super-influencers catch hold of these narratives and spread them to millions of people. Second, in the marketing space, advertisers may not have the budget to pay a super-influencer to market their product, but by paying the Alters of the Ego, the advertiser may be able to influence the super-influencer to support the product being marketed.

References

1. Secondary infektion (2021), https://secondaryinfektion.org/
2. Arnaboldi, V., Conti, M., Passarella, A., Pezzoni, F.: Analysis of ego network structure in online social networks. International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing 10, 31–40 (2012)
3. Arnaboldi, V., Conti, M., Passarella, A., Pezzoni, F.: Ego networks in twitter, an experimental analysis pp. 3459–3464 (2013)
4. Bastos, M.T., Mercea, D., Charpentier, A.: Tents, tweets, and events: The interplay between ongoing protests and social media. Journal of Communication 65(2) (2015)
5. Beskow, D.: Finding and Characterizing Information Warfare Campaigns. Ph.D. thesis (2020)
6. Bessi, A., Ferrara, E.: Social bots distort the 2016 u.s. presidential election online discussion. First Monday 21(11) (2016)

7. Bolllen, J., Mao, H., Zeng, X.: Twitter mood predicts the stock market. Journal: of Computational Science 2(1), 1–8 (2011)
8. Brangham, W.: "tech heads testify on misinformation in the aftermath of jan 6 riots. " PBS. org, https://www.pbs.org/newshour/show/tech-heads-testify-on-misinformation-in-the-aftermath-of-jan-6-riots, Mar 25 (2021)
9. Brendan: Trump twitter archive v2 (2021), https://www.thetrumparchive.com/
10. Carley, K.M.: Social cybersecurity: an emerging science (2020)
11. Cer, D, e.a.: Universal Sentence Encoder. ArXiv (2018), https://arxiv.org/abs/1803.11175
12. Chute, N.: He writes trump's tweets and has been with trump's campaign since day one (2019)
13. Coussi-Korbel, S., Fragaszy, D.M.: On the relation between social dynamics and social learning. Animal Behaviour 50(6), 1441–1453 (1995)
14. Dozen, T.D.: Center for Countering Digital Hate (2021)
15. Evanega, S, e.a.: Quantifying sources and themes in the covid-19 'infodemic' (2020)
16. Flache, A., Mäs, M., Feliciani, T., Chattoe-Brown, E., Deffuant, G., Huet, S., Lorenz: Models of social influence: Towards the next frontiers. Journal of Artificial Societies and Social Simulation 20(4) (2017), http://jasss.soc.surrey.ac.uk/20/4/2.html⟩.doi:10.18564/jasss.3521
17. Foundation, M.: Disinformation about belarus spreads in the balkans via online portals and social media (2020)
18. Friedkin, N., Johnsen, E.: Social influence network theory: A sociological examination of small group dynamics (structural analysis in the social sciences) (2011)
19. Friedkin, N.E., Proskurnikov, A.V., Tempo, R., Parsegov, S.E.: Network science on belief system dynamics under logic constraints. Science 354(6310), 321–326 (2016), https://science.sciencemag.org/content/354/6310/321
20. Giglietto, F., Righetti, N., Rossi, L., Marino, G.: It takes a village to manipulate the media: coordinated link sharing behavior during 2018 and 2019 italian elections. Information, Communication & Society 23(6), 867–891 (2020)
21. Google: Universal sentence encoder [computer software] (2018)
22. Granger, C.W.J.: Investigating causal relations by econometric models and cross-spectral methods. Econometrica 37(3), 424–438 (1969)
23. Greene, W.H.: Econometric Analysis. Macmillan Publishing Company (199)
24. Hui, P.M., Yang, K.C., Torres-Lugo, C., Monroe, Z., McCarty, M., Serrette, B.D., Pentchev, V., Menczer, F.: Botslayer: real-time detection of bot amplification on twitter. Journal of Open Source Software 4(42), 1706 (2019), https://doi.org/10.21105/joss.01706
25. JustAnotherArchivist: snscrape [computer software] (2021), https://github.com/JustAnotherArchivist/snscrape
26. M, H.W., S, G., A, D., M, R., L, U., HA, S., DH, E., L, L., Curtis: Bots and misinformation spread on social media: Implications for covid-19. Journal of Medical Internet Research 23(5) (2021)
27. Mittal, A., Goel, A.: Stock prediction using twitter sentiment analysis (2011), http://cs229.stanford.edu/proj2011/GoelMittal-StockMarketPredictionUsingTwitterSentimentAnalysis.pdf
28. Nedić, A., Olshevsky, A., Uribe, C.A.: Graph-theoretic analysis of belief system dynamics under logic constraints. Scientific Reports 9 (2019)
29. the Now, D.: Hydrator [computer software], https://github.com/docnow/hydrator
30. Pacheco, D., Hui, P., Torres-Lugo, C., Truong, T., B, F., A, Menczer, F.: Uncovering coordinated networks on social media: Methods and case studies (2020)

31. Pacheco, D., Flammini, A., Menczer, F.: Unveiling coordinated groups behind white helmets disinformation. Companion Proceedings of the Web Conference 2020 (WWW '20), Association for Computing Machinery pp. 611–616 (2020)
32. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in python. Journal of Machine Learning Research 12, 2825–2830 (2011)
33. Shahapure, K.R., Nicholas, C.: Cluster quality analysis using silhouette score. In: 2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA). pp. 747–748 (2020)
34. Starbird, K., Wilson, T.: Cross-platform disinformation campaigns: Lessons learned and next steps (2020)
35. Steeg, G., Galstyan, A.: Information transfer in social media. Proceedings of the 21st international conference on World Wide Web p. 509–518 (2012)
36. T, N., N, K., C, C., A, H., A., H.: The danger of misinformation in the COVID-19 crisis. Missouri Medicine 117(6) (2020)
37. Uyheng, J., Carley, K.M.: Characterizing network dynamics of online hate communities around the covid-19 pandemic. Applied Network Science 6(20) (2021)
38. Velmurugan, T., et al.: Efficiency of k-means and k-medoids algorithms for clustering arbitrary data points. Int. J. Computer Technology & Applications 3(5), 1758–1764 (2012)