

# Influence Dynamics Among Narratives

## A Case Study of the Venezuelan Presidential Crisis

Akshay Aravamudan<sup>1</sup>, Xi Zhang<sup>1</sup>, Jihye Song<sup>2</sup>, Stephen M. Fiore<sup>2</sup>, and Georgios C. Anagnostopoulos<sup>1</sup>

<sup>1</sup> Department of Computer Engineering & Sciences,  
Florida Institute of Technology, Melbourne, Florida, USA  
{aaravamudan2014, zhang2012}@my.fit.edu, georgio@fit.edu  
<sup>2</sup> University of Central Florida, Orlando, Florida, USA  
chsong@knights.ucf.edu, sfiore@ist.ucf.edu

**Abstract.** It is widely understood that diffusion of and simultaneous interactions between narratives — defined here as persistent point-of-view messaging — significantly contributes to the shaping of political discourse and public opinion. In this work, we propose a methodology based on Multi-Variate Hawkes Processes and our newly-introduced Process Influence Measures for quantifying and assessing how such narratives influence (Granger-cause) each other. Such an approach may aid social scientists enhance their understanding of socio-geopolitical phenomena as they manifest themselves and evolve in the realm of social media. In order to show its merits, we apply our methodology on Twitter narratives during the 2019 Venezuelan presidential crisis. Our analysis indicates a nuanced, evolving influence structure between 8 distinct narratives, part of which could be explained by landmark historical events.

**Keywords:** Influence Dynamics · Social Media Narratives · Multi-Variate Hawkes Processes · Granger Causality · Venezuelan Presidential Crisis

## 1 Introduction

Discourse on social media platforms has become increasingly relevant in political contexts throughout the world [11, 19]. There have been quite a few studies that have illustrated the power that social media yields in shaping public opinions and influencing political dialogues [10, 13]. While it has been argued that the Internet has not revolutionized how politics is conducted [9], it does provide an efficient medium for rhetoric and discussion born out of existing cultural, economic, and political situations [4]. As such, understanding the nature of how ideas, opinions,

---

Akshay Aravamudan, Xi Zhang, Jihye Song, Stephen M. Fiore, and Georgios C. Anagnostopoulos. Influence dynamics among narratives. In Robert Thomson, Muhammad Nihal Hussain, Christopher Dancy, and Aryn Pyke, editors, *Social, Cultural, and Behavioral Modeling*, 204–213. Cham, 2021. Springer International Publishing. doi: 10.1007/978-3-030-80387-2\_20.

and calls to action diffuse over social media provides insights into how real-world events may unfold, especially in the context of areas experiencing political and economic turmoil. Such insights may help interested parties in developing a more comprehensive portrait of developing socio-geopolitical contentions.

The Venezuelan presidential crisis is an indicative example of significant political events. It began on January 10<sup>th</sup>, 2019, when Nicolás Maduro was sworn in for a second presidential term following a disputed election in May of 2018. This development prompted international reactions with nations siding with either Nicolás Maduro or Juan Guaidó, then president of Venezuela’s National Assembly. The situation escalated after Guaidó declared himself as the interim president on January 23<sup>rd</sup>, 2019. All the while, there was significant participation on social media, which drove conversations globally. During this period, Twitter was heavily employed, resulting in competing opinions and, thus, rendering it as an effective tool for political communication [16].

We study this in the context of narratives associated with this event. We develop an interdisciplinary approach that unites social science theorizing on narratives with computational approaches to information evolution. Our narrative approach combines a parsimonious definition of online narratives (see [2]) as “... *recurring statements that express a point of view.*” We use this in conjunction with stances taken on a subset of narratives (*i.e.*, support or disagree). In this way, we are able to model the mutual influence of complementary and competing narratives evolving over time. As such, we put forward an approach for quantifying and assessing the co-evolving influences among narratives under the prism of Granger causality as it is defined in the context of Temporal Point Processes (TPPs). This allows one to discern concrete influence motifs between online narratives. Such motifs serve as an additional facet of characterizing online discussions, which may prove useful to social scientists studying online discourse. We next describe work related to our approach. Then in Sec. 3, we provide an overview of our dataset. Sec. 4 describes our methodology in some detail. As proof of concept, we then apply our methodology to study the influence dynamics of 8 concurrent narratives connected to the Venezuelan presidential crisis using Twitter data as outlined in Sec. 3. Finally, we provide and comment on our findings in Sec. 5. In particular, we are able to identify influence patterns that can be correlated to milestones of the crisis.

## 2 Related Works

Our work on narrative modelling falls under a broader umbrella of dynamic topic modelling for temporally sequenced documents, where topics correspond to narratives and documents to tweets. A first example of such an approach is temporal LDA [17], which extends traditional Latent Dirichlet Allocation (LDA) and allows for learning topic transitions over time. A couple of other examples that we mention here are based on TPPs, which are used to model topic dynamics. Lai et al. [12] propose using a marked Hawkes process to uncover inter-topic relationships. Similarly, Mohler et al. [14] propose a Hawkes-Binomial topic

model to detect mutual influence between online Twitter activity and real-world events. While our work also falls under this general strand, it focuses solely on learning topic dynamics (see Sec. 4), since the topics themselves are ultimately provided by Subject Matter Experts (SMEs) as described in Sec. 3.

### 3 Data

The Twitter dataset used for this study is a subset of a larger social media dataset curated by a data provider as part of a larger grant program and is described in [2] and [3]. In accordance with the requirements of the organization funding the research, the data were anonymised prior to sharing it with researchers to protect user privacy. Overall, it encompasses over 7 million tweets, overwhelmingly in Spanish, from December 25<sup>th</sup>, 2018, to February 1<sup>st</sup>, 2019, a period which coincides with the commencement of the presidential crisis.

Each tweet has been annotated with potentially multiple narrative labels. The narratives were derived by applying topic modelling on the tweets’ textual content via Non-negative Matrix Factorization combined with tf-idf statistics [18]. These topics were subsequently refined and formalized by SMEs. A subset of the tweets was then manually annotated by the same SMEs and were then used to train a BERT-based multilingual cased multi-label classification model [15] that was then used to annotate the remaining tweets. A similar procedure was followed to label tweets as pro- or anti-Maduro.<sup>1</sup>

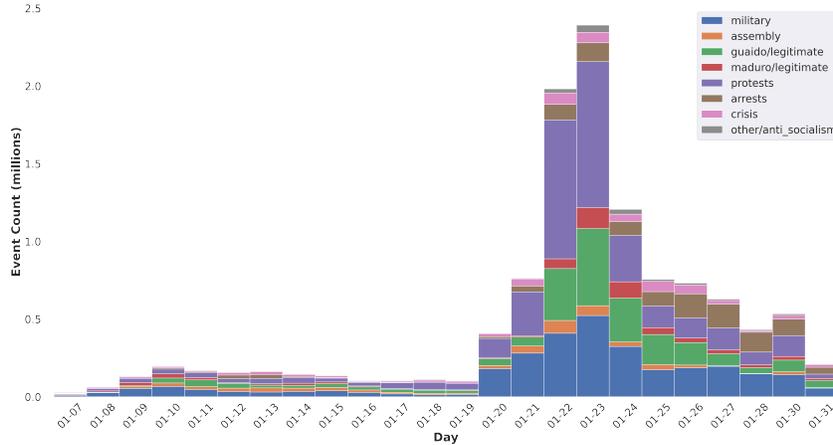
**Table 1.** Stance distribution per narrative of the Venezuela Twitter data.

Narrative	Total Tweets	% anti-Maduro	% pro-Maduro
<i>military</i>	1,534,242	67.38	21.87
<i>assembly</i>	252,448	95.79	1.82
<i>guaido/legitimate</i>	1,014,726	95.43	3.02
<i>maduro/legitimate</i>	304,127	2.92	96.72
<i>protests</i>	1,746,615	85.87	2.42
<i>arrests</i>	570,574	97.73	0.74
<i>crisis</i>	305,291	73.25	3.58
<i>anti-socialism</i>	101,716	78.04	14.77

<sup>1</sup> Narrative and stance labelling was carried out by the data provider and was provided to us as is with very limited description of the process followed.

<sup>2</sup> <https://www.theguardian.com/world/2019/jan/21/venezuela-claims-foiled-attempted-military-uprising>

<sup>3</sup> <https://www.theguardian.com/world/2019/jan/23/venezuela-protests-thousands-march-against-maduro-as-opposition-sees-chance-for-change>



**Fig. 1.** Histogram of Twitter event (tweet) counts per narrative in 2019. One observes a burst of activity between January 19<sup>th</sup> and 21<sup>st</sup>, during which there was a small-scale coup initiated by 27 soldiers.<sup>2</sup> The peak activity, occurring on January 23<sup>rd</sup>, appears strongly related to massive protests, which demanded Maduro to step down.<sup>3</sup> During the same day, we also witness a significant increase in anti-Maduro tweets.

We only considered narratives that were present in at least 100,000 tweets. As a result, we analyzed a total of 8 narratives. The daily event counts of these narratives in the time period of interest is shown in Fig. 1, while Tab. 1 shows the distribution of stances per narrative. The *military* narrative includes discussions about the Venezuelan army, security services, or other organizations that reported to Maduro’s government. *Assembly* includes any mentions of Venezuela’s National Assembly. *Guaido/legitimate* and *maduro/legitimate* consist of tweets that expressly support the legitimacy of Guaidó and Maduro, respectively. *Protests* includes tweets that mention anti-Maduro demonstrations, public gatherings, or rallies. *Arrests* includes tweets that refer to people who had been imprisoned at the time. Moreover, the *crisis* narrative label refers to the Venezuelan humanitarian crisis<sup>4</sup> and finally, *anti-socialism* includes tweets that mention socialism, communism, or leftism as the primary cause of the humanitarian crisis.

## 4 Methodology

### 4.1 Multi-Variate Hawkes Processes

In this work, we employed Multi-Variate Hawkes Process (MVHPs) [8] — a particular system of TPPs — to characterize the temporal dynamics of tweets

<sup>4</sup> <https://www.reuters.com/article/us-venezuela-politics-un/venezuelans-facing-unprecedented-challenges-many-need-aid-internal-u-n-report-idUSKCN1R92AG>

comprising the narratives of interest. MVHPs have been widely used for modelling social media events (*e.g.*, [7, 20]) as they are capable of describing self- and mutually-exciting modes of event generation. TPPs are completely characterized by their conditional intensity  $\lambda(t | \mathcal{H}_{t-})$  at time  $t$  given their past observations  $\mathcal{H}_{t-}$ , which is referred to as the history of the process. The quantity  $\lambda(t | \mathcal{H}_{t-})dt$  yields the expected number of events such a process will generate in the interval  $(t, t + dt)$ . Assume an MVHP comprising  $P$  processes. Its  $i^{\text{th}}$  process features a conditional intensity of the form

$$\lambda_i(t | \mathcal{H}_{t-}) = b_i(t) + a_{i,i} \sum_{t_k^i \in \mathcal{H}_{t-}^i} \phi_{i,i}(t - t_k^i) + \sum_{\substack{j \in \mathcal{P} \\ j \neq i}} \alpha_{i,j} \sum_{t_k^j \in \mathcal{H}_{t-}^j} \phi_{i,j}(t - t_k^j) \quad (1)$$

where  $i \in \mathcal{P} \triangleq \{1, 2, \dots, P\}$ ,  $\mathcal{H}_{t-}^i$  is the history of the  $i^{\text{th}}$  process and  $\mathcal{H}_{t-} \triangleq \cup_{i \in \mathcal{P}} \mathcal{H}_{t-}^i$  is the history of the multi-variate process. In Eq. 1,  $b_i(\cdot) \geq 0$  is a history-independent intensity component, which we will refer to as base or background intensity. The quantity  $\phi_{i,j}(t - t_k) \geq 0$  reflects the (typically, momentary and, subsequently, diminishing) increase in intensity that the  $i^{\text{th}}$  process will experience at time  $t$  due to an event of the  $j^{\text{th}}$  process, which occurred at time  $t_k$ . The functions  $\phi_{i,j}(\cdot)$  are referred to as memory kernels. Also, for fixed base intensity and memory kernels, the non-negative  $a_{i,j}$ 's constitute the MVHP's model parameters. Finally, apart from the background intensity, one can distinguish two additional terms in Eq. 1: **(i)** a second, self-exciting term, through which past events of the process may generate further future events and **(ii)** a third, cross-exciting term, through which events of the other processes may do the same. An MVHP's base intensity and memory kernel functions are typically inferred non-parametrically, while the  $a_{i,j}$  parameters are learned via (sometimes, penalized) maximum likelihood estimation.

For the purposes of studying influences among narratives, we used MVHPs that featured one process per narrative. Each process utilised a constant background intensity  $b_i$  and process-independent memory kernels  $\phi_{i,j}(t) = e^{-t}$ , which is a typical choice for Hawkes processes. Moreover, we opted to fit such MVHPs using overlapping time frames within the period of interest for two reasons: **(i)** initial experimentation gave strong indications—via probability-probability plots—that fitting a single MVHP to the entire period would result in a bad-fitting model. It appears that the tweet dynamics under consideration are not well approximated by MVHPs at long time ranges, but only at shorter ones. And, **(ii)** modeling tweet dynamics on overlapping (instead of disjoint) time frames has a beneficial smoothing effect on the estimated model parameters, when comparing time-adjacent models.

## 4.2 Quantifying & Comparing Influence Effects

Based on the notion of causality for continuous stochastic processes, Eichler et al. [6] extend the definition of Granger (predictive) causality to MVHPs and show that, if  $a_{i,j} > 0$ , then the  $j^{\text{th}}$  process Granger-causes the  $i^{\text{th}}$  process. This allows

one to discern influences within and between processes by mere inspection or, more formally, by testing against  $a_{i,j} = 0$ . Moreover, the base intensity of each process can be naturally interpreted as a nonspecific cause, which is external to the system of processes. One can reject the absence of such unaccounted-for cause by testing against  $b_i = 0$ . Of course, the entirety of this causal reasoning is predicated on the absence of other confounding processes (or factors, in general), which is an assumption that we will also adopt for our studies. Thence, in our exposition we will make use of the notion of apparent influences among processes as a stand-in for Granger-causal effects sans confounding.

While testing for  $a_{i,j} = 0$  determines a potential influence via a binary decision, quantifying the magnitude of the influence effect based on the  $a_{i,j}$  parameters in order to draw comparisons is a much more nuanced issue. Some prior work, such as [1] and [12], for example, attempt to directly interpret the parameters' magnitudes as corresponding magnitudes of influence for the purpose of comparisons. This approach, though, is fraught with problems: **(i)** the first one is semantic in nature: the physical meaning of the  $a_{i,j}$ 's is intrinsically linked to the specific form of the memory kernels they multiply and, hence, is often far from straightforward to describe. And, **(ii)** solely relying on the  $a_{i,j}$ 's values to quantify the influence effect completely ignores memory kernels, which are equal contributors in shaping event dynamics; based on this approach, this renders comparing influence effects problematic.

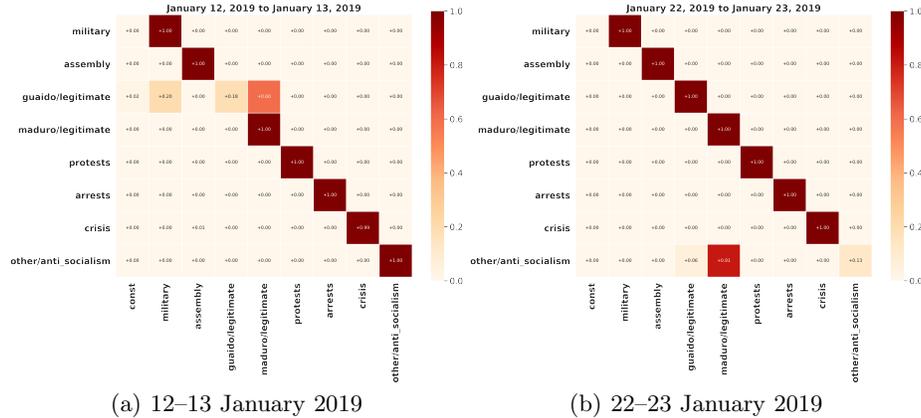
In order to circumvent these shortcomings, we introduce and advocate the use of Process Influence Measures (PIMs) for quantifying the magnitude of process-to-process influences. In particular, we define the  $j$ -to- $i$  PIM  $\pi_{i,j} \in [0, 1]$  as the probability that an event of the  $j^{\text{th}}$  process is the most likely cause for an event of the  $i^{\text{th}}$  process. These probabilities can be easily estimated from the results provided by a fitted model. This is due to the following fact that readily stems from the theory of TPPs, whose conditional intensity is represented additively: if  $E_k^i$  is the  $k^{\text{th}}$  event of process  $i$ , which occurred at time  $t_k^i$ , and  $\mathcal{E}^j$  denotes the set of all events of process  $j$ , then

$$\mathbb{P} \{ E_k^i \text{ was caused by any earlier event in } \mathcal{E}^j \} = \frac{a_{i,j} \sum_{t_\ell^j \in \mathcal{H}_{t_k^i}^j} \phi_{i,j}(t_k^i - t_\ell^j)}{\lambda_i(t_k^i | \mathcal{H}_{t_k^i}^i)} \quad (2)$$

$$\mathbb{P} \{ E_k^i \text{ was caused by } b_i(t) \} = \frac{b_i(t_k^i)}{\lambda_i(t_k^i | \mathcal{H}_{t_k^i}^i)} \quad (3)$$

for  $i, j \in \mathcal{P}$ . It is important to note that these probabilities jointly depend on model parameters, memory kernels as well on relative event timings and, thus, capture multiple facets of process dynamics. One can iterate over all events of process  $i$  and, each time, record the process  $j$ , which exhibits the largest probability given by Eq. 2. Then, the PIM  $\pi_{i,j}$  is estimated via the frequency with which an event of process  $j$  was the likeliest cause for an event of process  $i$ .

## 5 Results & Discussion

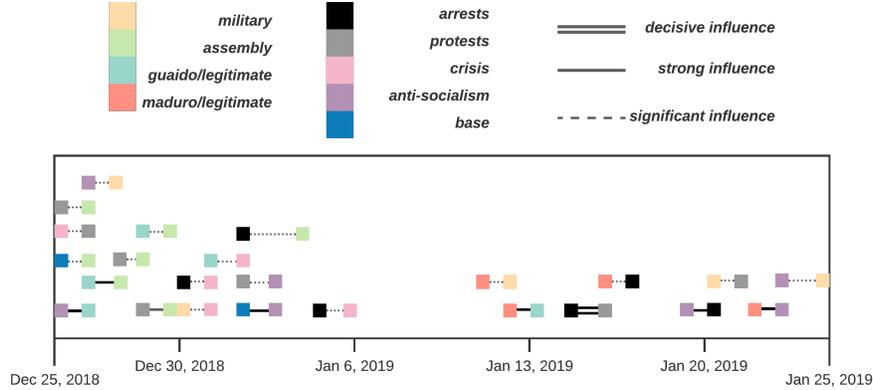


**Fig. 2.** PIM heat maps for two time frames. Columns show influence sources, while rows depict the narratives we studied. These maps illustrate self-driving narratives (prominent diagonal entries), as well as inter-narrative influences (sizeable off-diagonal entries).

We used the setup described in Sec. 4 to infer intra/inter-narrative influences from the available data. We split the observed time period into 35 two-day time frames that were overlapping by one day.<sup>5</sup> For each such time frame, we trained a MVHP and ensured that it was exhibiting (at least) a reasonably good fit as judged by probability-probability plots. For investigating influences among narratives, each MVHP encompassed a process per narrative. Furthermore, we computed and tabulated PIMs as described in Sec. 4.2. Fig. 2 provides PIMs as heat maps for two indicative time frames: one before Guaidó’s self-proclamation as legitimate president and one after. These particular examples illustrate self-reinforcing narratives ( $\pi_{i,i} \approx 1$  diagonal values), as well as narratives being influenced by one or more of the remaining narratives (appreciable  $\pi_{i,j}$  off-diagonal values). We characterize influences in terms of PIMs as follows: “significant” influences, when  $\pi_{i,j} \in (0.2, 0.6]$ , “strong” influences, when  $\pi_{i,j} \in (0.6, 0.99]$  and “decisive,” when  $\pi_{i,j} > 0.99$ . Fig. 3 highlights such influences over the time frame that our study considered.

Judging from the results in Fig. 3, we notice that *anti-socialism* strongly influences *guaido/legitimate* between December 25<sup>th</sup> and 26<sup>th</sup>, which may reflect Guaidó’s support from Western governments and the growing dissatisfaction of

<sup>5</sup> In particular, we chose two-day time frames to reduce the computational burden of training. Also, we used an hourly timescale to represent event time stamps to maintain the numerical stability of our training algorithm.



**Fig. 3.** Timeline indicating noteworthy influences based on our estimated PIMs. Weaker cross-narrative influences with PIM values  $\pi_{i,j} \in [0.0, 0.2]$  and self-influences of any kind have been omitted for clarity. Note that, during the period of January 6<sup>th</sup> through the 11<sup>th</sup>, narratives only influence themselves ( $\pi_{i,i} \approx 1$ ).

the Venezuelan populace with existing conditions under Maduro’s government, which were prevalent even before the crisis. Moreover, after Maduro’s inauguration on January 10<sup>th</sup>, we notice that *maduro/legitimate* significantly influences *military*. The tweets during these days were posted after the first open cabildo — a political action convention — held by then president of the National Assembly, Guaidó, where he was voted in as acting president. This could have motivated Maduro supporters to favor a strengthening of his claim to the presidency and, by extension, favor the military’s response to the unrest.<sup>6</sup> Between January 12<sup>th</sup> and 13<sup>th</sup>, *maduro/legitimate* significantly affects *guaido/legitimate* showing a reactionary response to the mobilisation of Maduro’s military in response to rising support for Guaidó. This also coincides with Guaidó’s brief arrest<sup>7</sup> by the Bolivarian Intelligence Service on January 13<sup>th</sup>. Finally, we notice that between January 20<sup>th</sup> and 21<sup>st</sup>, *military* significantly influences *protests*, which overlaps with when some National Guardsmen rose against Maduro.<sup>8</sup> This was followed by a widespread and vociferous protests against Maduro. Between January 24<sup>th</sup> and 25<sup>th</sup>, we observe a significant influence from *anti-socialism* to *military*, which may be attributed to the wake of anti-Maduro protests<sup>9</sup> on January 23<sup>rd</sup> that called for the military to relinquish their allegiance to Maduro.

<sup>6</sup> <https://www.bbc.com/news/world-latin-america-47036129>

<sup>7</sup> <https://www.nytimes.com/2019/01/13/world/americas/venezuela-juan-guaido-arrest.html>

<sup>8</sup> <https://www.nytimes.com/2019/01/21/world/americas/venezuela-maduro-national-guard.html>

<sup>9</sup> <https://www.theguardian.com/world/2019/jan/23/venezuela-protests-thousands-march-against-maduro-as-opposition-sees-chance-for-change>

Finally, a potentially interesting extension of our work could specifically address strategic platform manipulation that may drive narratives during political crises. For instance, an operation involving Facebook and Instagram accounts attributed to a U.S. communications firm was flagged and removed for coordinated inauthentic behavior targeting Venezuela<sup>10</sup> and was found to have primarily posted anti-Maduro content [5]. Additionally, Twitter announced the removal of a large Venezuelan state-backed operation from its platform.<sup>11</sup> Given the role of the Venezuelan political crisis in international geopolitics, future work could examine how domestic and foreign actors are leveraging social media to manipulate narratives for political objectives. To achieve this, we could use a more elaborate MVHP to model such actors as separate processes.

In summary, we have empirically demonstrated via our Venezuelan case study that PIMs, as introduced in Sec. 4.2, furnish an unambiguous and interpretable dynamic characterisation of intra-/inter-narrative influences. Furthermore, we have showcased how such dynamics may be explained by landmark events — exogenous to social media — within broader historical contexts. In this capacity, our work provides an additional, important lens for studying influence between narratives, which, we hope, may prove useful to social scientists.

## Acknowledgments

This work was supported by the U.S. Defense Advanced Research Projects Agency (DARPA) Grant No. FA8650-18-C-7823 under the *Computational Simulation of Online Social Behavior (SocialSim)* program of DARPA’s Information Innovation Office. Any opinions, findings, conclusions, or recommendations contained herein are those of the authors and do not necessarily represent the official policies or endorsements, either expressed or implied, of DARPA, or the U.S. Government. Finally, the authors would like to thank the manuscript’s anonymous reviewers for their helpful comments and suggestions.

## References

1. Alvari, H., Shakarian, P.: Hawkes process for understanding the influence of pathogenic social media accounts. In: 2019 2nd International Conference on Data Intelligence and Security (ICDIS). pp. 36–42 (2019), <https://doi.org/10.1109/ICDIS.2019.00013>
2. Blackburn, M., Yu, N., Berrie, J., Gordon, B., Longfellow, D., Tirrell, W., Williams, M.: Corpus development for studying online disinformation campaign: A narrative + stance approach. In: STOC@LREC (2020)
3. Blackburn, M., Yu, N., Memory, A., Mueller, W.G.: Detecting and Annotating Narratives in Social Media: A Vision Paper. ICWSM, US (Jun 2020), <https://doi.org/10.36190/2020.23>

<sup>10</sup> <https://about.fb.com/news/2020/09/august-2020-cib-report/>

<sup>11</sup> [https://blog.twitter.com/en\\_us/topics/company/2019/further\\_research\\_information\\_operations.html](https://blog.twitter.com/en_us/topics/company/2019/further_research_information_operations.html)

4. Calderaro, A.: Social media and politics. In: Outhwaite, W., Turner, S. (eds.) *The SAGE Handbook of Political Sociology*, 2v, chap. 44, pp. 781–796. Sage Publications Ltd., first edn. (Dec 2017)
5. Cryst, E., Ponce de León, E., Suárez Pérez, D., Perkins, S.: Bolivarian Factions: Facebook Takes Down Inauthentic Assets. Tech. rep., Stanford’s Freeman Spogli Institute for International Studies (Sep 2020)
6. Eichler, M., Dahlhaus, R., Dueck, J.: Graphical modeling for multivariate Hawkes processes with nonparametric link functions. *Journal of Time Series Analysis* **38**(2), 225–242 (2017), <https://doi.org/10.1111/jtsa.12213>
7. Farajtabar, M., Gomez-Rodriguez, M., Wang, Y., Li, S., Zha, H., Song, L.: Co-evolve: A joint point process model for information diffusion and network co-evolution. In: *Companion Proceedings of the The Web Conference 2018*. p. 473–477. WWW ’18, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE (2018), <https://doi.org/10.1145/3184558.3186236>
8. Hawkes, A.G.: Spectra of some self-exciting and mutually exciting point processes. *Biometrika* **58**(1), 83–90 (1971), <https://doi.org/10.2307/2334319>
9. Hindman, M.: *The Myth of Digital Democracy*. Princeton University Press (2009)
10. Jost, J.T., Barberá, P., Bonneau, R., Langer, M., Metzger, M., Nagler, J., Sterling, J., Tucker, J.A.: How social media facilitates political protest: Information, motivation, and social networks. *Political Psychology* **39**(S1), 85–118 (2018), <https://doi.org/10.1111/pops.12478>
11. Jungherr, A.: Twitter use in election campaigns: A systematic literature review. *Journal of Information Technology & Politics* **13**(1), 72–91 (2016), <https://doi.org/10.1080/19331681.2015.1132401>
12. Lai, E., Moyer, D., Yuan, B., Fox, E., Hunter, B., Bertozzi, A., Brantingham, P.: Topic time series analysis of microblogs. *IMA Journal of Applied Mathematics* **81**, hwx025 (07 2016), <https://doi.org/10.1093/imamat/hxw025>
13. Margetts, H., John, P., Hale, S., Yasseri, T.: *Political turbulence: How social media shape collective action*. Princeton University Press (2015)
14. Mohler, G., McGrath, E., Buntain, C., LaFree, G.: Hawkes binomial topic model with applications to coupled conflict-Twitter data. *The Annals of Applied Statistics* **14**(4), 1984 – 2002 (2020), <https://doi.org/10.1214/20-AOAS1352>
15. Pires, T., Schlinger, E., Garrette, D.: How multilingual is multilingual BERT? In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. pp. 4996–5001. Association for Computational Linguistics, Florence, Italy (Jul 2019), <https://doi.org/10.18653/v1/P19-1493>
16. Sytnik, A.: Digitalization of diplomacy in global politics on the example of 2019 Venezuelan presidential crisis. In: Alexandrov, D.A., Boukhanovsky, A.V., Chugunov, A.V., Kabanov, Y., Koltsova, O., Musabirov, I. (eds.) *Digital Transformation and Global Society*. pp. 187–196. Springer International Publishing, Cham (2019), [https://doi.org/10.1007/978-3-030-37858-5\\_15](https://doi.org/10.1007/978-3-030-37858-5_15)
17. Wang, Y., Agichtein, E., Benzi, M.: Tm-lda: Efficient online modeling of latent topic transitions in social media. In: *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. p. 123–131. KDD ’12, Association for Computing Machinery, New York, NY, USA (2012), <https://doi.org/10.1145/2339530.2339552>
18. Xu, W., Liu, X., Gong, Y.: Document clustering based on non-negative matrix factorization. In: *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. p. 267–273. SI-

## Influence Dynamics Among Narratives

- GIR '03, Association for Computing Machinery, New York, NY, USA (2003), <https://doi.org/10.1145/860435.860485>
19. Zhang, W., Johnson, T.J., Seltzer, T., Bichard, S.L.: The revolution will be networked: The influence of social networking sites on political attitudes and behavior. *Social Science Computer Review* **28**(1), 75–92 (2010), <https://doi.org/10.1177/0894439309335162>
  20. Zhou, K., Zha, H., Song, L.: Learning social infectivity in sparse low-rank networks using multi-dimensional hawkes processes. In: Carvalho, C.M., Ravikumar, P. (eds.) *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research*, vol. 31, pp. 641–649. PMLR, Scottsdale, Arizona, USA (29 Apr–01 May 2013), <http://proceedings.mlr.press/v31/zhou13a.html>