5^{th} International Conference on Computational Social Science IC^2S^2
July 17-20, 2019, University of Amsterdam, The Netherlands

Measuring the Influence of Online Misinformation: A Hierarchy of Social Media Data

Dimitra (Mimie) Liotsiou, Philip N. Howard

University of Oxford, Oxford, UK
{mimie.liotsiou, philip.howard}@oii.ox.ac.uk

Keywords: misinformation, social media, social influence, causality, Big Data

Extended abstract

In recent years, the proliferation of misinformation and propaganda campaigns on social media around elections and key political events has been recognized as an important threat to democracy internationally. Researchers have attempted to measure the reach of and engagement with such campaigns on social media [3, 7]. However, measuring the influence of online misinformation on actual election outcomes remains a challenge, as it requires more data than social media companies typically make externally available [2, 5].

To help address this problem, we propose here a hierarchical typology of social media data, encoding the usefulness of each type of data for measuring the influence of online misinformation on actual election behavior (Figure 1). For each type of data, we also present how accessible social media platforms make it to external researchers.

The typology orders data in Levels from 0 to 4, going from less to more useful for measuring the influence of misinformation. Level 0 corresponds to data on misinformation content on social media, e.g. posts or ads. This may indicate an attempt to influence elections, but on its own does not indicate whether this attempt succeeded. For that, one needs to know that social media users were exposed to this misinformation (Level 1), and/or that out of those, some engaged with it (Level 2). But engagement with misinformation does not guarantee it will change one's election behavior – for that to be investigated, one needs data on the actual election behavior (Level 3), and, to isolate the influence (causal effect) of this misinformation, one generally also needs to account for other relevant causes of the election behavior (Level 4) [5, 6]. Social media platforms have internal data on all levels, particularly for Levels 0-2 (3 stars in Figure 1), but often also for Levels 3 and 4 to some extent (2 stars) [1]. They make externally available some data from Level 0 and 2 (1-2 stars), but generally not from Level 1 (exposure data, e.g. who or how many people viewed a post), and not from Levels 3 and 4.

We illustrate this typology further by drawing upon our findings from an analysis of the Russian Internet Research Agency's (IRA) social media campaign targeting US voters around the 2016 US presidential elections [4]. For this we used public and non-public data ("Special Data" column in Figure 1) from social media platforms (Facebook, Instagram, Twitter, YouTube), containing posts and ads, with aggregate engagement numbers. This analysis contributes novel insights, including that the IRA's campaign was very extensive (millions of tweets, tens of thousands of Facebook and Instagram posts), and had very wide reach (e.g. more than 189 million Instagram likes and comments in 2013-2018). It was strategically focused on key divisive issues (e.g. immigration, race), and often targeted users based on their interests in these, as well as their demographics and location. We also found large spikes in IRA activity around key political events like candidate debates, and that activities continued and often rose even after the election, in 2017. This special data spanned many years (2012-2018), included some

exposure data (Level 1) for Facebook ads only, and covered all public IRA accounts detected by the platforms – data that would not have been accessible through the platforms' public APIs.

More useful

Level	Types of Data	Public API Data	Special Data	Internal Data
4	Election behavior + other causes			★★
3	Election behavior			★★
2	Engagement	★	★★	★★★
1	Exposure		★	★★★
0	Misinformation campaign posts, accounts, ads	★	★★	★★★

Figure 1: Hierarchical typology of social media data, in terms of their usefulness for measuring the influence of online misinformation on election behavior (0: low, 4: high), along with their accessibility (public, special, internal data). The number of stars is proportional to the quality and quantity of the data.

Going forward, given that the platforms' provision of special data enabled crucial conclusions to be reached on the IRA's activities and reach across platforms over time, continuing and expanding external researchers' access to such previously internal data would have immense benefits for deepening our understanding of online misinformation campaigns. One avenue worth pursuing is offering further data on user exposure to misinformation (Level 1), to get a fuller picture of how many people were exposed to misinformation, even if they did not actively engage with it. Further, data on the demographics (e.g. location, age, race) of those who were exposed to and/or who engaged with IRA content would help determine whether the IRA succeeded in actually reaching target demographics, particularly in key swing states.

And in order to get closer to measuring the influence of misinformation campaigns on actual election behavior and outcomes, researches would need platforms to also share data from Levels 3 and 4. This would help answer questions such as e.g. what was the effect of IRA social media content on whether one voted (or not) in the 2016 US presidential elections? For this, one does not only need voting records (Level 3), but also data on other causes of voting behavior (Level 4) such as race, location, and interests (attributes that the IRA often used to target social media users, as we found), so as to isolate the influence of IRA content.

Acknowledgements

This work is funded by the European Research Council through the grant "Computational Propaganda: Investigating the Impact of Algorithms and Bots on Political Discourse in Europe," Proposal 648311, 2015-2020, Philip N. Howard, Principal Investigator, and the grant "Restoring Trust in Social Media Civic Engagement," Proposal Number: 767454, 2017-2018, Philip N. Howard, Principal Investigator. We are grateful for additional support from the Open Society Foundation and Ford Foundation. Project activities were approved by the University of Oxford's Research Ethics Committee, CUREC OII C1A 15 044, C1A 17 054. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the researchers and do not necessarily reflect the views of the funders, or the University of Oxford.

References

[1] Robert M Bond, Christopher J Fariss, Jason J Jones, Adam DI Kramer, Cameron Marlow, Jaime E Settle, and James H Fowler. A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415):295–298, 2012.

[2] Christian Caryl. There's a way to know if Russia threw the election to Trump. *The Washington Post*, 2019.

[3] Ritam Dutt, Ashok Deb, and Emilio Ferrara. "Senator, we sell ads": Analysis of the 2016 Russian Facebook ads campaign. In *International Conference on Intelligent Information Technologies*, pages 151–168. Springer, 2018.

[4] Philip N Howard, Bharath Ganesh, Dimitra Liotsiou, John Kelly, and Camille François. The IRA, social media and political polarization in the United States, 2012-2018. *Working Paper 2018.2. Oxford, UK: Project on Computational Propaganda. comprop.oii.ox.ac.uk*, 2018.

[5] Dimitra Liotsiou, Luc Moreau, and Susan Halford. Social influence: From contagion to a richer causal understanding. In *International Conference on Social Informatics*, volume 10047, pages 116–132. Springer, 2016.

[6] Cosma Rohilla Shalizi and Andrew C Thomas. Homophily and contagion are generically confounded in observational social network studies. *Sociological methods & research*, 40(2):211–239, 2011.

[7] Samuel C Woolley and Philip N Howard. *Computational Propaganda: Political Parties, Politicians, and Political Manipulation on Social Media*. Oxford University Press, 2018.