10-2012

Social Media and the Elections

Takis Metaxas
pmetaxas@wellesley.edu

Eni Mustafaraj
emustafa@wellesley.edu

OVERLINE

Social Media and the Elections

Panagiotis T. Metaxas and Eni Mustafaraj

Manipulation of social media affects perceptions of candidates and compromises decision-making.

In the United States, social media sites, such as Facebook, Twitter, and YouTube, are currently being used by two out of three people (*1*), and search engines are used daily (*2*). Monitoring what users share or search for in social media and the Web has led to greater insights into what people care about or pay attention to at any moment in time. Furthermore, it is also helping segments of the world population to be informed, to organize, and to react rapidly. However, social media and search results can be readily manipulated, which is something that has been underappreciated by the press and the general public.

In times of political elections, the stakes are high, and advocates may try to support their cause by active manipulation of social media. For example, altering the number of followers can affect a viewer's conclusion about candidate popularity. Recently, it was noted that the number of followers for a presidential candidate in the United States surged by over 110 thousand within one single day, and analysis showed that most of these followers are unlikely to be real people (*3*).

We can model propaganda efforts in graph-theoretic terms, as attempts to alter our "trust network": Each of us keeps a mental trust network that helps us decide what and what not to believe (*4*). The nodes in this weighted network are entities that we are already familiar with (people, institutions, and ideas), and the arcs are our perceived connections between these entities. The weights on the nodes are values of trust and distrust that we implicitly assign to every entity we know. A propagandist is trying to make us alter connections and values in our trust network, i.e., trying to influence our perception about the candidates for the coming elections, and thus "help us" decide on candidates of their choice.

The Web, as seen by search engines (*5*), is similarly a weighted network that is used to rank search results. The hyperlinks are considered "votes of support", and the weights are a computed measurement of importance assigned to Web pages (the nodes in the graph). It is also the target of propaganda attacks, known as "Web spam" (*6*). A Web spammer is trying to alter the weighted Web network by adding connections and values that support his or her cause, aimed at affecting the search engine's ranking decisions and thus the number of viewers who see the page and consider it important (*4*).

"Google bomb" is a type of Web spam that is widely known and applicable to all major search engines today. Exploiting the descriptive power of anchor text (the phrase directly associated with a hyperlink), Web spammers create associations between anchor words or phrases and linked Web pages. These associations force a search engine to give high relevancy to results that would otherwise be unrelated, sending them to the "top ten" search results. A well-known Google bomb was the association of the phrase "miserable failure" with the Web page of President G. W. Bush initially and later with those of Michael Moore, Hillary Clinton, and Jimmy Carter (*7*). Another Google bomb associated candidate John Kerry with the word "waffles" in 2004. A cluster of Google bombs was used in an effort to influence the 2006 congressional elections. Google has adjusted its ranking algorithm to defuse Google bombs on congressional candidates by restricting the selection of the top search results when querying their names (*8*). During the 2008 and 2010 elections, it proved impossible to launch any successful Google bombs on politicians, and it is hoped that the trend will continue.

During the 2010 Massachusetts Special Election (MASEN) to fill the seat vacated by the death of Senator Ted Kennedy, we saw attempts to influence voters just before the elections, launched by out-of-state political groups (*9*). Propagandists exploited a loophole introduced by the feature of including real-time information from social networks in the "top ten" results of Web searches. They ensured that their message was often visible by repeatedly posting the same tweet. A third of all election-related tweets sent during the week before the 2010 MASEN were tweet repeats (*9*). All search engines have since reacted by moving real-time results out of the organic results (results selected purely by information retrieval algorithms) and into a separate search category.

"Twitter bombs," however, are likely to be launched within days of the elections. A Twitter bomb is the act of sending unsolicited replies to specific users via Twitter in order to get them to pay attention to one's cause. Typically, it is done effectively by means of "bots," short-lived programs that can send a large quantity of tweets automatically. Twitter is good at shutting most of them down because of their activity patterns and/or users' complaints. However, bombers have used fake "replies" to spam real users who are not aware of their existence. For example, in the 2010 MASEN, political spammers created nine fake accounts that were used to send about 1000 tweets before being blocked by Twitter for spamming (*9*). Their messages were carefully focused, however, targeting users who in the previous hours were discussing the elections. With the retweeting help of similarly minded users, >60,000 Twitter accounts were reached within a day at essentially no cost. Twitter bombs, unfortunately, have become common practice.

A more sophisticated effort to create a fake grassroots movement [often referred to as "astroturf" (*10*)] was the creation of a "Prefab tweet factory" (*11*). Designed to evade Twitter's spam detection, a spammer created daily sets of tweets targeting journalists and urging other similarly minded users to tweet. The effect of this spam was to give the impression to the targeted journalists that their reporting was monitored and was not appreciated by "the public," and thus applied pressure to the reporters to modulate their views (*11*). We do expect to see such low-budget prefab tweets in the next elections and whenever opportunity for putting pressure on journalists arises (*12*).

One of the effective (but expensive) ways to spam is to buy online search ads (appearing at the top of the search results as "sponsored" search in search

Department of Computer Science, Wellesley College, 106 Central Street, Wellesley, MA 02481 USA. Email: pmetaxas@wellesley.edu; emustafa@wellesley.edu

engines and as "promoted" tweets in Twitter) that appear in queries including names and characteristics of the political opponent. When a citizen searches for a candidate's name or other related terms, the prominently placed and aptly worded ad will encourage them to click on it, thus transporting them to a page designed and maintained by the opponent. An example in a limited form occurred in the 2008 elections, when a site unfavorable to a candidate (TheRealBobRoggio.com), appeared as an advertisement when searching for the name of the candidate (13). The contents of these ads can be adjusted rapidly, allowing experimentation with titles and contents that will draw maximal attention. The selection of the best ad can be further refined to match the profile of the specific user with the use of data collected and mined by a process often described as "microtargetting" (14). A newer advertising tool will be the use of "promoted trends" in Twitter (15) to attract the attention of a wider, yet focused, audience. These techniques may be effective and legal but they are expensive, compared with the spamming techniques we mentioned above.

Yet more ways to spread spam may be through the use of photographs and videos that ridicule the opponent. Search engines usually allocate a prominent place in their organic results for images and videos of well-known people, including political candidates. Their selection in the search results depends on the keywords associated with them (not with their visual contents) and with the popularity in clicks they achieve. Insulting-while-funny pictures typically attract the curiosity of the users and can go viral, allowing propagandists to pass their message, while avoiding any automatic filtering by the search engines (16). Although this was observed during the 2010 elections (16), there is some evidence that search engines are working to clean their organic results, by asking users to report images they find offensive.

Owing to their popularity and ease of access, social media data have been used to attempt to predict future events, such as movie box-office revenues (17, 18), product sales (19), stock market fluctuations (20), and even electoral results. Predicting movie box-office revenues using Twitter (4) and Yahoo search (5) data can be extremely accurate if the predictions are based on unambiguous parameters and a careful consideration of potential confounders. Predicting election results via Twitter data (which are readily accessible) has been applied to reality TV competitions (21) and a few political elections (22–24). And it is easy to find Twitter polls promoted by newspapers for the current U.S. election (25). However, using social media for predicting political elections is highly controversial. There is no agreement among researchers yet on the measures responsible for any successful prediction (e.g., tweet volume or tweet content). The time period of data collection has also been variable, ranging from weeks to months before the elections and ending days to weeks before the elections. In most cases, researchers have filtered their data on the basis of decisions clearly made after the elections were over and the results were known (including which parties' tweets were included) (23). This has led to an inability to replicate reported success rates (23, 24). Representativeness is currently the most important problem (21). Just having a large number of tweets does not mean that there has been representative sampling of the voting population [e.g., in political conversations, 1% of the Twitter accounts are often responsible for 30% of the tweet volume (11)].

Even more than in previous elections, we should expect that all candidates and political parties will use social media sites to create enthusiasm in their troops, raise funds, and influence our perception of candidates (or our perception of their popularity). We should be aware of how that works and be prepared to search for the truth behind the messages.

References and Notes

1. Pew Foundation, 65% of online adults use social networking sites (2011); http://bit.ly/OWHYwA
2. Pew Foundation, Search engine use 2012 (2012); http://bit.ly/T9f814
3. M. C. Petermann, The Twitter underground economy: A blooming business (2012); http://bit.ly/W84BhK
4. P. T. Metaxas, in Lecture Notes in Business Information Processing (Springer-Verlag, New York, 2010), vol. 45, pp. 170–182; http://bit.ly/ffYsuC
5. A. Broder et al., Comput. Netw. 33, 309 (2000).
6. C. Castillo, B. D. Davison, Found. Trends Inform. Retriev. 4, 377 (2010).
7. Google bomb, Wikipedia (2012); http://bit.ly/TjOeKa
8. P. T. Metaxas, paper presented at Workshop on Information and Decision in Social Nets, MIT Media Lab, Cambridge, MA, 31 May to 1 June 1, 2011; http://bit.ly/NRmNox
9. P. T. Metaxas, E. Mustafaraj, in Proceedings of the Web-Sci10: Extending the Frontiers of Society On-Line, Raleigh, NC, 26 to 27 April 2010; http://bit.ly/qlnTmW
10. The "Truthy" project, http://truthy.indiana.edu/.
11. E. Mustafaraj et al., in Proceedings of IEEE PASSAT/Conf. on Social Computing, Minneapolis, MN, 9 to 11 October 2011 (IEEE, 2011), pp. 103–110; http://bit.ly/oZFqas
12. Journalists and media outlets, who typically have large number of followers, should be careful when retweeting without comment. Even though some journalists declare in their Twitter profile that their "Retweet does not imply agreement," this distinction is missed by the vast majority of the user population who view retweet as endorsement.
13. E. Mustafaraj, P. T. Metaxas, in Proceedings of ACM Information Retrieval and Advertising Workshop, Boston, MA, 23 July 2009 (ACM, NY, 2009); http://bit.ly/RjcJYn
14. T. Edsal, Let the nanotargeting begin, New York Times Online (2012); http://nyti.ms/QfX792
15. See video clip of Z. Moffatt on Wall Street Journal, WorldStream Video; http://on.wsj.com/RjaxQH
16. E. Mustafaraj, P. T. Metaxas, C. Grevet, Proc. IEEE Comput. Sci. Eng.4, 320 (2009); http://bit.ly/RjAkLF
17. S. Asur, B. A. Huberman, in Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (IEEE Computer Society, Los Alamitos, CA, 2010), pp. 492–499.
18. S. Goel, J. M. Hofman, S. Lahaie, D. M. Pennock, D. J. Watts, Proc. Natl. Acad. Sci. U.S.A. 107, 17486 (2010).
19. H. Choi, H. Varian, Predicting the present with Google trends; http://bit.ly/RdWiMB
20. J. Bollen, H. Mao, X. J. Zeng, J. Comput. Science 2, 1 (2011).
21. F. Ciulla et al., EPJ Data Science 1, 8 (2012).
22. D. Gayo-Avello, A meta-analysis of state-of-the-art electoral prediction from Twitter data, http://bit.ly/Pqwkd7 (2012).
23. A. Jungherr, P. Jurgens, H. Schoen, Soc. Sci. Comput. Rev. 30, 229 (2012); http://bit.ly/OtbOqO
24. P. T. Metaxas, E. Mustafaraj, D. Gayo-Avello, in Proceedings of IEEE PASSAT/Conference on Social Computing (IEEE, 2011), pp. 165–171; http://bit.ly/qUHJyv
25. M. T. Moore, Twitter index tracks sentiment on Obama, Romney, USA Today, 1 August 2012; http://usat.ly/Rjb1Gq