

Social Bots in Japan's 2014 General Election: Preliminary Results from a Corpus-Linguistic and Qualitative Study of Computational Propaganda on Twitter

Stefan Evert, Philipp Heinrich, Fabian Schäfer
Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany
stefan.evert@fau.de, philipp.heinrich@fau.de, fabian.schaefer@fau.de

Keywords: Social Bots, Japan, 2014 General Election, Twitter, Right-wing Extremism

Summary: We present results on identification and behavioral analysis of social bots in a sample of more than 500,000 tweets. Typical forms of bot activity include massive retweeting and repeated posting of (nearly) the same tweet, sometimes used in combination. We focus on the second method and present (i) a case study on two patterns of bot activity; (ii) methodological considerations on the automatic identification of such patterns and the prerequisite near-duplicate detection; and (iii) qualitative insights into the purposes behind the usage of political bots in Japan's general election of 2014.

Introduction: Japanese is the second most-used language on twitter.¹ Nevertheless, despite the size of the Japanese twittersphere, social media did not play an essential role in Japan's electoral campaigns until 2014. It was only after the 2009 general election that "politicians first jumped on the Twitter bandwagon", inspired by "the success that Obama was having with the platform."² This is due to plain judicial reasons, since online campaigning during the official election campaign period (usually 10 days before the election) had been legally banned in Japan until the revision of the Public Offices Election Law in 2013. Our paper presents preliminary results from a corpus-linguistic study of twitter data collected in the weeks before and after the day of the most recent general election in Japan (December 14th, 2014), a snap election announced by PM Abe Shinzō only one month before actually taking place. Although various researchers have already looked into the use of social media by Japanese politicians and the various political parties for campaigning, the role of social bots in Japanese politics has not been thoroughly researched yet. We understand social bots as "computer-generated programs that post, tweet, or message of their own accord"; if the main objective of their activity is a political one, one speaks of "political bots", being used to "manipulate public opinion, choke off debate, and muddy political issues".³

Sampling and Method: In the period from December 8th–30th, 2014, we collected 542,584 tweets via the public Twitter API, based on a set of keywords relating to the election, which were then tokenized and annotated with part-of-speech tags using MeCab. A preliminary inspection of sorted concordances with the CQPweb concordancer revealed a large number of near-duplicate

¹ <https://www.statista.com/statistics/267129/most-used-languages-on-twitter/>

² Williams, J.A., and D. Miller. "Netizens Decide 2014? A Look at Party Campaigning Online." In: *Japan Decides 2014: The Japanese General Election*, edited by R. Pekkanen, S.R. Reed, and E. Scheiner. (Palgrave, 2015).

³ Howard, P. N., and B. Kollanyi (2016). "Bots, #Strongerin, and #Brexit: Computational Propaganda during the UK-EU Referendum." *Working Paper 2016.1*. Oxford: Project on Computational Propaganda.

tweets, often pushing a certain political agenda. This observation suggested targeted campaigns by networks of social bots, attempting to inflate the frequency and importance of certain topics by posting many copies of the same tweet (with small modifications, e.g. different Web links) or by massive retweeting. We therefore implemented a simple and efficient algorithm for automatic identification of duplicates: tweets were normalized by deleting whitespace, punctuation (incl. @ and #), URLs and RT markers; a hash structure was then used to recognize exact duplicates of the normalized tweets. The algorithm identified 431,050 tweets (79.4%) as duplicates. Even if retweets—which account for 56.7% of the data and are duplicates by definition—are excluded, there are still 143,869 (61.2%) duplicates among the remaining 234,914 tweets.

Our further exploration of social bots was guided by two serendipitous findings: (A) a collocation analysis of #解散 (i.e. “dissolution of parliament”) showed that it appeared surprisingly often in the same sequence with #自民党 (Liberal Democratic Party; LDP), which makes sense, since the LDP was in power before the election, and #セクハラ (sexual harassment); (B) among the most frequent nouns in the corpus, we noticed Minami-Chōsen (南朝鮮), a somewhat disrespectful term for South Korea also used by right-wing extremists. In our analysis of these tweets, we provide evidence for the activity of bots during the election campaign. Although it is not possible to link them to a certain party or provide proof for an actual effect on real users (most of the fake accounts had only very few followers or none at all), we argue that bots could be either used to (i) “push” certain topics, claims, or slogans so that they become trending topics or hashtags and thus have a better chance of being picked up by the mass media, or to (ii) bandwagon on existing hashtags to multiply extremist positions. In both cases, social bots may be considered successful even if they do not reach many real users, because the main targets are quantitative analytics and the mass media.

Campaign A: The sequence #自民党 #セクハラ #解散 (see above) occurs in 9,065 tweets in our data set, all posted by the same user account @Stupid00002, which is a (now suspended) fake account of the female politician Ayaka Shiomura, who was subject to sexist taunts by a male colleague of the LDP in a Lower House committee meeting in April 2014. Our duplicate analysis showed that 99.4% of these tweets are multiple, slightly modified copies of only 57 unique tweets. The 4 most frequent ones were sent more than 460 times each. The fact that we found hardly any retweets shows (i) that the campaign had little effect on genuine users and (ii) that it did not involve a retweeting botnet. In order to test the accuracy of our automatic duplicate detection, we manually grouped near-duplicates among the 57 unique tweets and compared these groups to a hierarchical clustering based on the Levenshtein (minimum edit) distance between the normalized tweets, revealing that only 7 cases of near-duplicates were missed by our simple algorithm and all but two of the near-duplicates could have been identified automatically when applying a cutoff distance of 5, which confirms the usefulness of our approach. Further analysis showed that the same 57 tweets were sent by a second account (@excreta_ZAiFX) in slightly modified form; for each unique tweet exactly the same number of copies was sent from both accounts.

Campaign B: The term Minami-Chōsen (南朝鮮) occurs in 12,389 tweets in our data set; again, only a negligible proportion of 0.3% (36 tweets) are retweets. These tweets also show an extremely high duplication ratio of 97.0%, but the copies are spread across a much broader range of 271 different user accounts. In order to determine whether this campaign involved multiple

distinct networks of social bots, we analyzed the repetition patterns of all 359 unique tweets by graphing the number of user accounts posting each tweet against the total number of copies sent (blue points in Fig. 1). There are two highly conspicuous patterns, which suggest two botnets: (i) a set of tweets with close to 13.5 copies per user account on average, but different overall frequency (indicated by the red regression line); and (ii) a tight cluster of 10 tweets that is clearly separated from all other tweets (highlighted by a green circle). For comparison, tweets that do not contain the term 南朝鮮 are shown as grey points (retweets are excluded from these counts). While pattern (i) can only be clearly identified once the topical tweets (blue points) have been recognized, the cluster of pattern (ii) is so distinct from the repetition patterns of all other tweets that it could easily have been identified by a statistical analysis of the complete set of tweets. Several additional tweets (grey points) clearly belong to these patterns as well, even though they were not matched by our keyword search. Almost all the tweets in cluster (ii) are near-duplicates of each other, most likely sent from bots and are targeted against TPP, the planned reorganization of prefectures in Japan (*dōshū-sei*), and the reception of refugees to Japan.

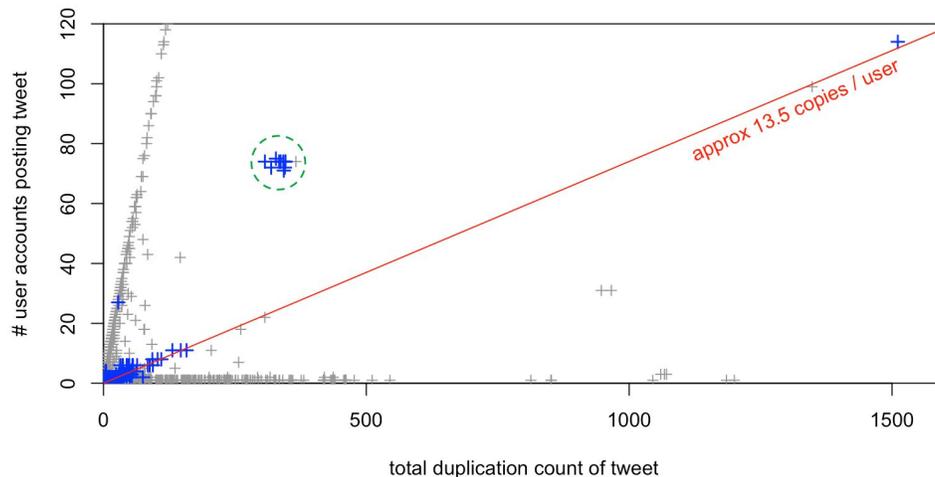


Figure 1: A correlation plot of duplication count (without retweets) against number of distinct user accounts for each unique tweet (tweets mentioning South Korea are shown in blue, others in grey) reveals four distinctive patterns of bot activity: (i) tweets with around 13.5 copies per user account (red regression line); (ii) a tight cluster of 11 tweets with remarkably similar statistics (green circle); (iii) frequent repetitions of the same tweet from a single account (forming a horizontal line at the bottom of the plot); (iv) repetition across a large number of different accounts with only a single copy from each account (steep line at the left-hand side of the plot).

Conclusion: Based on our observations, we conjecture that the dominant patterns of political bot activity in our data set could be identified automatically by combining a statistical analysis of tweet duplication patterns (which easily identifies patterns (ii)–(iv) in Fig. 1) with a thematic analysis of the tweet content (pattern (i) in Fig. 1 is clearly identified by the average number of copies per user account in combination with the keyword Minami-Chōsen). Telling from the names of fake user accounts or the tweeted content, we argue that it was most likely right-wing Internet extremists (jap. *Netto uyoku*) behind these botnets and not politicians or parties themselves, trying to propagate and multiply their ultranationalist positions by bandwagoning on existing hashtags and keywords. We intend to apply this approach to monitor political bot activity on Twitter in the run-up to the German parliamentary election of 2017.