Political Hashtag Hijacking in the U.S.

Asmelash Teka Hadgu
L3S Research Center
Hannover, Germany
teka@l3s.de

Kiran Garimella
Qatar Comput. Res. Inst.
Doha, Qatar
vgarimella@qf.org.qa

Ingmar Weber
Qatar Comput. Res. Inst.
Doha, Qatar
ingmarweber@acm.org

ABSTRACT

We study the change in polarization of hashtags on Twitter over time and show that certain jumps in polarity are caused by "hijackers" engaged in a particular type of hashtag war.

Categories and Subject Descriptors

H.3.5 [**Information Storage and Retrieval**]: Online Information Services; J.4 [**Social and Behavioral Sciences**]: Sociology

Keywords

political trends; twitter; political leaning classification; partisanship

1. INTRODUCTION

On Twitter, *hashtags* are used to label tweets as being related to a particular topic. Through them, users join virtual debates and they are used to "frame" issues. Users from opposing political camps engage in political "hashtag wars"[1] to obtain control over the terms being used. E.g., the political right established "obamacare" as the standard expression for the Affordable Healthcare Act. On Twitter, the left fought back with hashtags such as *#obamacares* or *#iloveobamacare*. Given their importance, the use of hashtags related to politics has been studied before [1, 3]. One important aspect which has not been studied, however, is the change of political polarization of hashtags over time. This helps campaign organizers to know when they are "under attack" and it helps citizens to know when a debate is dominated by political activists. We use retweets of labeled seed users, e.g., `@BarackObama`, to obtain Twitter users with an inferred political orientation. By analyzing their hashtag usage we assign a leaning to hashtags and monitor this leaning over time. "Change points" with a sudden jump in leaning are identified and we show that they correspond to the activity of "hashtag hijackers", whom we characterize in detail. The methodology in this work is generalizable to a multi-party system, e.g., U.K.

2. METHODOLOGY

The methodology to assign a leaning to hashtags is identical to [5] and similar to [1]. It is summarized as follows. We start with a set of 14/19 seed users for the left/right respectively. We then get all their public tweets and look at all people retweeting them. Retweeting users are then filtered for U.S. locations in their profiles using Yahoo! Placemaker and assigned a (fractional) leaning according to which side they retweet more. We validated this leaning against `wefollow.com`, `twellow.com` and `persecuting.us`. A user contributes fractionally to each leaning. The (mis-)classification accuracy is then weighted by this fraction. When comparing against Persecuting, both our labels and the ground truth are weighted and cases "close to the middle" contribute less. The accuracies are 98.6%, 93% and 90.4% respectively. For the labeled users their tweets are obtained and scanned for hashtags. Apolitical hashtags are removed by looking at co-occurrence with a set of seed political hashtags such as *#obama* or *#tcot*. A leaning with respect to a party p is then assigned to hasthag h in week w according to $\text{Lean}(h,w,p) = (\frac{v_p}{V_p} + \frac{2}{\sum_P V_p})/(\sum_P \frac{v_p}{V_p} + \frac{2*|P|}{\sum_P V_p})$. Here v_p denotes the aggregated user volume of a fixed (h, w) pair for a party p, V_p denotes the total user volume of all hashtags in w for p, and $|P|$ is the number of parties, two in our setting. The definition of $\text{Lean}(h,w,p)$ is a volume-based voting approach where (i) within a given week each party is given the same weight, and (ii) a regularization term reduces extreme leaning values for low volumes. User volumes, rather than tweet volumes, are used as they are more robust against a small number of outlier users.

[1] `politi.co/MILaI5`

WWW 2013 Companion, May 13–17, 2013, Rio de Janeiro, Brazil.
ACM 978-1-4503-2038-2/13/05.

3. DETECTION OF CHANGE POINTS

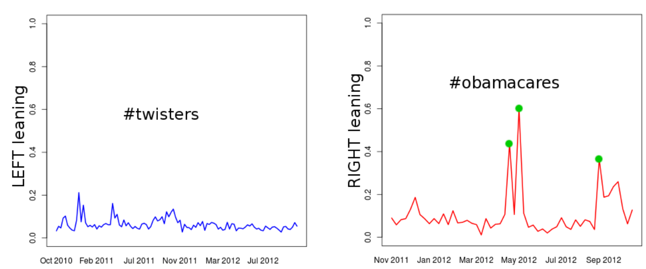

Figure 1: An example of a consistently right-leaning hashtag on the left, a left-jumping-towards-right hashtag on the right. Identified change points are highlighted in green.

Figure 1 shows an example of a hashtag with sudden changes in leaning. We will refer to such outliers as *change points*. We restrict our focus to change points corresponding both to (i) *upwards* jumps and (ii) cases where the party is "usually inactive", meaning an average leaning across all weeks of $< 1/2$. These interesting cases are directly caused by an unusually high level of hashtag usage by a given leaning, rather than by the absence thereof. Note that the set of change points depends on the leaning under consideration. To detect change points, we tried different algorithms [2] against a rule-based heuristics. In the end, we used the following rule-based approach as it gave the most consistent results.

1. Total_number_of_weeks >= 4
2. Change_from_previous_week > std
3. Change_from_previous_week > 0.20
4. Current_value − Average_value > std

These rules are only applied to (hashtag, leaning) pairs where the hashtag, averaged across all weeks, has an average leaning of $< 1/2$. If the leaning value of a hashtag at a given week meets all the criteria it is marked as a "change point". The set of change points (= (h, w) pairs) identified is denoted by CP where this set depends on the leaning under consideration.

4. ANALYSIS OF CHANGE POINTS

To quantify how change points differ, we computed a number of features. First, a "normalized volume" capturing the focus all parties assign to the hashtag. It is defined as NV $= \sum_{p \in P} \frac{f(h,w)^p}{f(*,w)^p}$. Second, the normalized volume for only the leaning at hand, i.e., NVL $= \frac{f(h,w)^p}{f(*,w)^p}$. Third, a volume-based trending score TS as defined in [4]. Fourth, the ratio of the week indices, starting from 1 and only covering weeks where the hashtag was present. Fifth and sixth again the combined normalized volume and the with-leaning normalized volume. The first three features are computed in a micro-average manner, where hashtags with more (non-)change points contribute more. The last three were computed as macro-averages for the sets of 2447/2551 hashtags with at least one change point for the left/right respectively, where each hashtag with a change point contributes the same. For the first three, we performed Student t-tests to check if the feature values differ for (h, w) pairs that are change points (*_cp) from those that are not (*_ncp). We first used a Fisher's F-Test to test for equality of variance and then used the appropriate unpaired t-test. In all the cases, the tests show statistically significant results with $p < 0.01$. The results tell the following narrative. First, change points happen for low volume, un-trending weeks. This makes sense as high volume hashtags such as #tcot would be hard to hijack. When, however, only hashtags with change points are considered the volume is slightly higher than for the non-change points and, in particular, the hijacking side puts in about 5-8 times its usual (normalized) volume. Finally, change points do not occur sooner or later than non-change points as the index ratio is close to 1.0.

Lean.	NV x1000 cp / ncp	NVL x1000 cp / ncp	TS cp / ncp	index (cp/ncp)	NV (cp/ncp)	NVL (cp/ncp)
Left	.92 / 1.55	.48 / .77	3.55 / 4.52	1.01	1.52	4.59
Right	.71 / 1.56	.41 / .78	3.53 / 4.52	1.05	1.08	7.71

Table 1: Hypothesis test for features computed on hashtags from each party separately. There were 2,827 change points jumping "left-wards" and 3,094 "right-wards" out of a total of 133,907 (h, w) pairs. The last three columns are macro-averages for the 2447/2551 hashtags h with at least on change point for the left/right.

5. ANALYSIS OF HIJACKERS

In this section, we dig deeper and (i) identify "hashtag hijackers" active in change points and (ii) give basic characteristics about their Twitter usage. We first identified (h, w) pairs as being change points as previously described. Next, for such (h, w) pairs, we looked at users from the "other leaning" (compared to h's overall normal leaning) using h in week w and looked at their usage frequencies of h during that week. A user was awarded a "hijacker point" if the user used h during a change point. Then, we ranked users according to the number of hijacker points they collected for all identified change points and considered the top 1,000 users for each leaning. To give a qualitative impression of how these hijackers differ from normal users, Figure 2 shows term clouds for the four sets of user profiles (left vs. right and normal vs. hijacker). Hijacking users can be seen to use party-related terms more often.

To find general differences between hashtag hijackers and "normal" users, we looked at their respective (i) number of tweets, the fraction of (ii) retweets in general and (iii) of seed users in particular, and the number of (iv) their followers and (v) users they are following. All of those statistics (Table 2) indicate that hashtag hijackers are more active than other politically interested users.

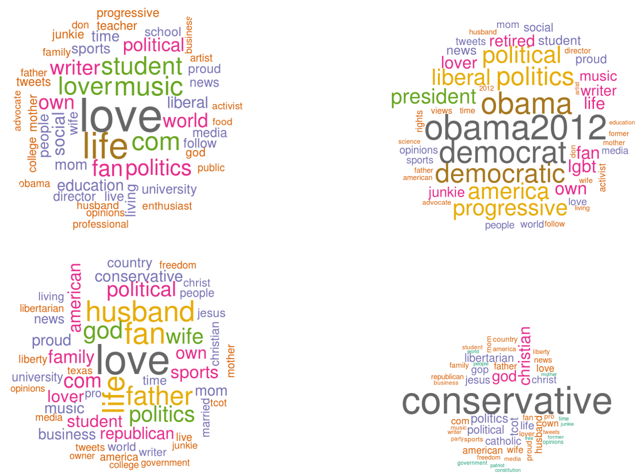

Figure 2: Clouds of terms used in profiles. Left users on top, right users on the bottom; normal users on the left, hijackers on the right.

leaning	set	tweets	retw.	seed retw.	fol-wers	fol-wing	focus	frac.*	avg. lean
L	n	4544	12.8%	0.6%	1554	573	11.3%	9.6%	0.81
	h	13620	7.6%	0.2%	2176	1268	27.1%	22.4%	0.97
R	n	3928	13.7%	1.3%	1291	684	10.8%	22.0%	0.74
	h	17306	11.6%	0.3%	5930	1969	21.4%	32.2%	0.95

Table 2: Comparison of basic statistics for the top hashtag hijackers per leaning against all users from the corresponding leaning. The 'n' indicates normal users and 'h' indicates hijackers. * – fraction of tweets containing political hashtags.

However, this could be explained by the fact that we rank users by their number of hijacking points, thereby creating a bias towards more active users. To rule out this explanation, we computed the (i) focus of a user's activity during change points as the fraction of a user's hashtag volume that occurs during change points, (ii) fraction of tweets by these users containing political hashtags, (iii) average leaning of these users. All these three features are consistently higher for hijackers, giving evidence for a non-random, conscious decision.

6. CONCLUSIONS AND FUTURE WORK

We presented a study of temporal hashtag polarization for the U.S. on Twitter. We focused on "change points" where sudden jumps in political leaning happen and showed that such jumps correspond to activity by "hashtag hijackers" – highly active and politicized users. This hijacking, a special form of hashtag wars, tends to happens for low volume hashtags which are more vulnerable.

7. REFERENCES

[1] M. Conover, J. Ratkiewicz, M. Francisco, B. Goncalves, A. Flammini, and F. Menczer. Political polarization on twitter. In *ICWSM*, 2011.

[2] V. Hodge and J. Austin. A survey of outlier detection methodologies. *Artificial Intelligence Review*, 22(2), 2004.

[3] D. M. Romero, B. Meeder, and J. Kleinberg. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *WWW*, 2011.

[4] I. Subasic and C. Castillo. The effects of query bursts on web search. In *WI*, 2010.

[5] I. Weber, V. R. K. Garimella, and A. Teka. Political hashtag trends. In *ECIR Demo*, 2013. politicalhashtagtrends.sandbox.yahoo.com.