Secular vs. Islamist Polarization in Egypt on Twitter

Ingmar Weber
Qatar Computing Research Institute
P.O. Box 5825, Doha, Qatar
Email: ingmarweber@acm.org

Venkata R. Kiran Garimella
Qatar Computing Research Institute
P.O. Box 5825, Doha, Qatar
Email: vgarimella@qf.org.qa

Alaa Batayneh
Al Jazeera English
P.O. Box 23127, Doha, Qatar
Email: BataynehA@aljazeera.net

***Abstract*—We use public data from Twitter, both in English and Arabic, to study the phenomenon of secular vs. Islamist polarization in Twitter. Starting with a set of prominent seed Twitter users from both camps, we follow retweeting edges to obtain an extended network of users with inferred political orientation. We present an in-depth description of the members of the two camps, both in terms of behavior on Twitter and in terms of offline characteristics such as gender. Through the identification of partisan users, we compute a valence on the secular vs. Islamist axis for hashtags and use this information both to analyze topical interests and to quantify how polarized society as a whole is at a given point in time. For the last 12 months, large values on this "polarization barometer" coincided with periods of violence. Tweets are furthermore annotated using hand-crafted dictionaries to quantify the usage of (i) religious terms, (ii) derogatory terms referring to other religions, and (ii) references to charitable acts. The combination of all the information allows us to test and quantify a number of stereo-typical hypotheses such as (i) that religiosity and political Islamism are correlated, (ii) that political Islamism and negative views on other religions are linked, (iii) that religiosity goes hand in hand with charitable giving, and (iv) that the followers of the Egyptian Muslim Brotherhood are more tightly connected and expressing themselves "in unison" than the secular opposition. Whereas a lot of existing literature on the Arab Spring and the Egyptian Revolution is largely of qualitative and descriptive nature, our contribution lies in providing a quantitative and data-driven analysis of online communication in this dynamic and politically charged part of the world.**

I. Introduction

Since Mohamed Bouazizi set himself on fire in Tunisia on December 17, 2010 many Arab countries have undergone turmoils and revolutions. Governments in Tunisia, Libya, Egypt and Yemen were toppled, unrest continues in Bahrain and open civil war in Syria. Though it is still under debate how much of a causal effect social media played during these unrests, Twitter and Facebook were widely utilized to organize protests with governments trying to cut off such communication links.

With the change in government, previously outlawed organizations such as the Muslim Brotherhood in Egypt have seized the opportunity to widen their online presence. In this work, we look at the case of Egypt as a young democracy to study how tensions between opposing political camps materialize in social media. Concretely, we focus on polarization between Islamist and secular forces on Twitter.

We build on observations made for US politics [1] to use retweeting of "seed" users' tweets to obtain a set of politicized tweeps.[1] The seed users consist of prominent tweeps, representing both ideologies. By analyzing the content of their engaged audience we can test a number of hypotheses related to the communication patterns and other behavior of the two camps. Our key findings are as follows.

Retweeting signifies endorsement. Using simple retweet information we could label users as either Islamist or secular with accuracy similar to inter-judge agreement.

Similar user sets. Overall, both sides attract tweeps with similar characteristics, both in terms of activity level, gender distribution and with respect to demographics (where students with a technical background are important in both sets).

Polarized hashtags can be identified. Assigning a polarity score to hashtags leads to the identification of politically charged topics.

Tension over time. Monitoring how polarized and "far from the center" the set of all hashtags is over time, i.e., how unique their usage is to a single political camp, captures a general notion of "political tension, with period of high tension correlating with periods of violence. Simple measures of volume or per-user polarity trends do not reveal such patterns.

Vocabulary congruence. The distribution of user-user hashtag similarities is more heavy tailed for Islamists, hinting at a small user group within this camp with a tendency to be more in unison. The similarity with between Islamist seed users and non-seed users follows a bimodal distribution, suggesting two sub groups.

Media preferences and polarity. Applying valence measures to tweeted URLs we can identify sites attracting viewers from a particular political camp. Apart from identifying partisan sites we also find evidence for the hypothesis that the Arabic version of Al Jazeera is somewhat closer to Islamists than the English version, though both are relatively well "centered".

Community structure in retweeter graph. Network-wise we observe and quantify a tendency for like-minded people to connect to each other. This tendency is, however, less pronounced than in the US political setting.

More devoutness, more donations, but fewer insults. Last not least we show that closeness to Islamists comes with (i) a higher propensity to use religious terms, (ii) a higher propensity to use charity related terms and (iii) a *lower* propensity to use derogatory expressions referring to other religions. All three trends are robust with respect to the inclusion of a range of other variables.

[1] We use the term "tweep" as a short hand for any Twitter user, regardless of experience or activity level.

II. Related Work

Related work can be loosely classified into two topics. First, there are studies on political polarization of politicians, party manifestos, search terms and many other objects. Generally, this line of work focuses on left-vs.-right polarization in US politics. Second, there is a large body of work on the Egyptian Revolution and on the Arab Spring in general. Though this line of work typically includes a discussion on the role of social media during the revolution (and during ongoing protests) it does not use data from Twitter to test hypotheses related to polarization. Each of the two topics is discussed in more detail in the following.

Political Polarization. One of the best known studies on political polarization looked at US political blogs [2]. They observed a clear split of the blogosphere into two communities, corresponding to Democrats and Republicans. The retweet graph constructed from our data set shows a somewhat similar characteristics. A related study was presented in [1]. Though set in a very different context, namely, US politics with a left-vs.-right polarization as opposed to Egyptian politics with a secular-vs.-Islamist polarization, some of their methods are similar to ours. In particular, we modify their formula to compute the valence of hashtags and we also investigate the congruence of the retweet and mention networks. The rest of our analysis, e.g., looking at the polarity of URLs or studying the interdependence between polarity and usage of certain classes of terms, not only differs in the overall setting but also in the questions asked. The problem of classifying tweeps at scale is studied in [3]. The authors use a machine learning approach to determine user categories such as political affiliation, ethnicity identification and affinity for a particular business. Such an approach could be used to further expand our current set of identified users. We believe, however, that the observed trends would not change if more data of similarly high quality is added. Though we are focusing on polarity among tweeps, one of the earliest related works quantified the political leaning of *politicians*. The NOMINATE score [4] uses voting records in combination with dimensionality reduction. Such techniques can be used to quantify the statement that US American politics follows a 1-dimensional left-to-right schema. The leaning of web search queries has been investigated in [5]. Though the methodology used to identify the leaning relies on blogs with a political leaning, the formula for their "leaning" is similar to ours. A similar approach has also been applied to data from Twitter to study left-vs.-right polarization in the US. Though we are using a very similar setup to [6] to assign a "leaning" to hashtags, the research questions addressed are completely different. To estimate party positions of unknown texts, word frequencies have been extracted from labeled sets (e.g., party programs) [7], [8]. Such purely content-based approaches could be applied as an alternative to our link-based approach.

Egypt and Egyptian Revolution. The work conceptually closest to ours is [9] where Mostak uses geo-referenced data from Twitter in combination with fine-grained geographic census data to test the hypotheses that Islamism is more widespread in low income areas. To quantify a tweep's level of affinity for political Islamism, three measures are proposed. The first measures the vocabulary difference between the tweep's tweets and content mined from an Arabic-language forum associated with the Muslim Brotherhood. The second looks at the degree to which the user stops tweeting at the five daily prayer times mandated by Islam. And the last looks at the number of leading Muslim Brotherhood politicians that the user follows. This last approach is similar to our use of retweet relationships. However, the general questions asked and the methodology applied differ completely.

Several studies have looked at the role that social media played during the Egyptian revolution [10], [11], [12], [13], [14], [15], often from a qualitative point of view or relying on surveys. None of them has, however, tried to quantify the polarization online in the post-revolution setting. The work in [16] differs as it puts more emphasis on computer science methodology and algorithms for sentiment and response analysis. Using their methodology they observe that "Twitter discussion on the Egyptian revolution was much different from other Twitter discussions in terms of who was tweeting and what was said". The blogosphere in Egypt and the Arab world has also been the focus of academic research. An application of using the Habermasian public sphere as a theoretical framework is presented in [17], a study of two political blogs and how they cover Muslim-Christian tensions. Etling et al. [18] present a descriptive study of approximately 3,000 hand-coded blogs in a large network of 35,000 Arabic-language blogs. They observe that Bloggers are focused mainly on domestic political issues with the Israeli-Palestine conflict being of universal interest. Outgoing links are largely to Web 2.0 sites such as YouTube and Wikipedia, followed by pan-Arab mainstream media sources, such as Al Jazeera. Alanie et al. [19] present a qualitative analysis of the role blogs played within the context of the Egyptian revolution using blog data. They find that blogs provided a "counter-narrative" to the government government-supplied version of events during the 18-day uprising. In [20] the point is raised that the role of Al Jazeera should also not be ignored and that, in fact, Al Jazeera is tightly linked with peer-produced material, as exemplified by their Creative Commons[2] and "The Stream"[3] initiatives. Two survey-based studies on political engagement and (i) Internet usage in Egypt and the (ii) how people imagine networks across dimensions of class, religion and other factors are presented in [21] and [22] respectively.

III. Data Set

Our data acquisition starts with the creation of a list of "seed users" which are manually labeled as secularist or Islamist. Part of the list was taken from [9] with additional entries added by an Egyptian expert. The final list of seed users is shown in Table I. For each of the seed users we obtained (up to) their most recent 3,200 tweets. For each of these tweets we then obtained (up to) 200 retweeting users. Both limits are imposed by the Twitter API. For each of the identified users we obtained their Twitter "bio" and, in particular, the "location" field. Where this field was non-empty we used Yahoo! Placemaker[4] to detect geographic place references. Yahoo! Placemaker works for input strings in several languages, including English and Arabic. The 7,088 users with a place reference from Egypt were kept for further

[2]http://cc.aljazeera.net/

[3]http://stream.aljazeera.com/

[4]http://developer.yahoo.com/boss/geo/docs/free_YQL.html

analysis. All their public tweets (up to 3,200 per user) were then downloaded once around January 2013 and again around March 2013, resulting in a total of 16,889,153 tweets.

Secularists		Islamists	
Twitter Name	Screen_Name	Twitter Name	Screen_Name
Mohamed ElBaradei	@ElBaradei	Muhammad Morsi	@MuhammadMorsi
Alaa Al-Aswany	@alaaaswany	Fadel Soliman	@FadelSoliman
Ayman Nour	@AymanNour	Essam Al Erian	@EssamAlErian
Wael Abbas	@waelabbas	Almogheer	@almogheer
Belal Fadl	@belalfadl	Hazem Salah	@HazemSalahTW
Dr. Hazem Abdelazim	@Hazem_Azim	Khaleed Abdallah	@KhaleedAbdallah
MohamedAbuHamed	@MohamedAbuHamed	Melhamy	@melhamy
HamzawyAmr	@HamzawyAmr	Dr Mohamed Aly	@dr_mohamed_aly
E3adet_Nazar	@E3adet_Nazar	Mustafa Hosny	@MustafaHosny
GameelaIsmail	@GameelaIsmail	El_Awa	@El_Awa
shabab6april	@shabab6april		
waelabbas	@waelabbas		

TABLE I. SEED TWITTER ACCOUNTS FOR EGYPTIAN POLITICS.

Users were fractionally assigned to one of the two sides (Islamist or secularist) according to which seed users they retweeted. A user who retweeted X distinct secular and Y distinct Islamist seed users would have a secular score of $X/(X+Y)$ and an Islamist score of $Y/(X+Y)$. Depending on which is larger we label users as "secularist" or "Islamist" with 154 tied cases labeled as "center". When a single value is needed, e.g., for correlation and dependence analysis, we used the Islamist score. To detect retweets we used the public tweets of each user and searched for "^RT @seed_user:". This way the recall of retweets was improved as we could obtain more than 200 retweet events per source tweet. In a further attempt to link the analysis to offline variables, we compiled a list of gender-specific first names both in Arabic and English. There were 939 male 387 female Arabic names. This dictionary was then compared against the so-called "real name" field in Twitter profiles. In 4,077 out of 7,088 cases a match was found and the inferred gender was noted. We used -1, 0, +1 to encode male, unknown and female respectively. Tweets were classified as English, Arabic or other/unknown using a language detection tool [5]. The per-user macro-averaged language statistics as well as other statistics about the 6,934 non-center users are in Table II.

	Secularists	Islamists
# tweeps	5,215	1,719
Av. followers	1,468	884
Av. friends	457	501
Av. tweets	7,587	7,830
# men	2,382	786
# women	556	201
% English	16.9	16.2
% Arabic	71.3	72.7

TABLE II. STATISTICS ABOUT TWEEPS RETWEETING SEED USERS. THE NUMBER OF TWEETS IS PER THE PROFILE PAGE.

To see if the political label inferred through the retweeting information is accurate, we asked two judges, both native Arabic speakers fluent in English and well aware of the political situation in Egypt, to label a set of 99 users, 49-50 from either camp, into "Islamist - supports the Muslim Brotherhood or their ideology", "Secularist - supports the secular opposition or their ideology" or "Unknown - apolitical or cannot tell from profile page".[6] Judges were provided with hyperlinks to the tweeps' Twitter profiles.

The two judges labeled 52—29, 15—26 and 32—44 tweeps as Islamist, secularist and unknown respectively. The first judge disagreed with the ground truth for 17/52 of his Islamists and 2/15 of his secularists, ignoring the "unknown" label. For the second judge the disagreement was 4/29 and 5/26. Combined, 77% of non-unknown labels agreed with the inferred label. For the cases where both judges had provided non-unknown label their mutual agreement was 36/45=80% - indicating that our simple method performs similar to human experts. We believe the agreement could be improved by providing judges with "evidence" for the label, e.g., retweets of the seed users, rather than relying on them to filter through hundreds of tweets for relevant clues.

Note that we deliberately chose to *exclude* the seed users from our analysis as otherwise signals related to polarity might be artificially amplified.

IV. WHO ARE THEY?

Table II shows that with respect to activity, gender and even language distribution the two camps are largely similar. Surprisingly, this even holds for most of the terms used in their profiles, as shown in Figure 1. Despite the similarities, certain stereotypical differences related to "liberal" or "muslim" can be observed.

Fig. 1. Word clouds of terms used in profiles by Islamists (left) and secularists (right).

V. WHAT DO THEY TWEET ABOUT?

To get a handle on what the two different camps were tweeting about we monitored the hashtags they use and assigned a "valence" to them. This valence quantifies the polarity of the hashtags used and is very similar to methods used in existing work [1][6]. Let v_I denote the aggregated (fractional) user volume of a hashtag h for the Islamists. Here, a user with an Islamist score of 0.8 would contribute 0.8 to this aggregate if he ever used h. Let V_I denote the total Islamist user volume of all hashtags. We define v_S and V_S similarly. Then, we compute the polarity of h as

$$\mathrm{Pol}(h) = \frac{\frac{v_I}{V_I} + \frac{2}{V_I+V_S}}{\frac{v_I}{V_I} + \frac{v_S}{V_S} + \frac{4}{V_I+V_S}}, \tag{1}$$

where a polarity of 1.0 is fully Islamist and 0.0 is fully secularist. We also use suffixes i and s as in 0.7s to denote the absolute value polarity that is predominantly (i)slamist or (s)ecular. Table III shows the top 5 hashtags in terms of polarity. Only hashtags with a minimum of 50 distinct users were included in the computation. Interestingly, not all hashtags are polarized and popular apolitical examples include #ff and #iphone5, with nearly perfectly balanced polarity of .49 and .48 respectively. Co-occurrence based techniques [6] could be used to filter out hashtags without a political context, but we deliberately chose to keep them for our analysis. In a sense, they are the glue that keeps society together or, more

[5] http://code.google.com/p/language-detection/

[6] The initial set contained 100 users but one user's profile was removed at the time of labeling.

formally, they bring down overall polarity (see the following section).

Hashtag (Islamist)	Leaning	Hashtag (Secular)	Leaning
مؤُتمر_نصرة_الأَحوَاز (Al-Ahwaz_Support_Conference)	0.948i	مرسي_كذَاب_قوي (Morsy_is_a_big_liar)	0.982s
جبهة_الخرَاب (destruction_front)	0.947i	feb10	0.982s
جرَائم_المعَارضة_السلمية (crimes_of_the_peaceful-_opposition)	0.943i	عَايزين_البوب_رئِيس_وزرَاء (we_want_the_Bob_(El-Baradei)_as_prime_minister)	0.981s
بَاسم_عودة (Basem_Ouda)	0.941i	حسن_البَاندَا_في_التَاريخ (Hassan_Al-Panda_in_history)	0.981s
يحيَى_عيَاش (Yahya_Ayyash)	0.940i	علي_قنديل (Ali_Qandeel)	0.979s

TABLE III. TOP 5 HASHTAGS, IN TERMS OF LEANING.

VI. IS SOCIETY DRIFTING APART?

Can Twitter provide signals for growing political tensions? We decided to test this hypothesis by looking at hashtag usage in a global manner. The basic idea is that if both camps live "in their own bubble" with their own language then this indicates tension. Concretely, we looked at how polarized hashtags over all are. A hashtag with a full 1.0 (= 1.0i) Islamist or a full 0.0 (= 1.0s) secular leaning is fully polarized at a value of 1.0. Similarly, a hashtag used by both sides in the corresponding proportions is not be polarized at all at a value of 0.0. Tracking this value for all hashtags at a given point in time quantifies the overlap between the two speech bubbles. We chose a weekly granularity for our analysis, going back "only" to March 2012. The reason for this choice is that Twitter introduced an Arabic language interface in March 2012[7] and that our data showed different adoption rates for Arabic hashtags before and after this date with a sudden jump happening in a single week.

Figure 2 shows the overall hashtag polarity over time, as well as the number of distinct hashtags in use in a given week. Only hashtags used by at least three distinct users in a given week are considered. The peak in polarity at the end of 2012 seems to coincide with the political struggle over the constitution and a planned referendum on the topic. Outbreaks of violence are marked and refer to the following events. **a** - Assailants with rocks and fireboms gather outside Ministry of Defence to call for an end to military rule. **b** - Demonstrations break out after President Morsi grants himself increased power to protect the nation. Clashes take place between protestors and Muslim Brotherhood supporters. **c,d** - Continuing protests after the November 22^{nd} declaration. **e** - Demonstrations in Tahrir square, Port Said and all accross the country. **f,g** - Demonstrations at Tahrir square. [8]

Figure 2 illustrates that hashtag usage can serve as a "barometer" for tension in society and that the mere number of distinct hashtags being used does not. Quite strikingly, all outbreaks of violence happened during periods where the hashtag polarity was comparatively high. To see if a similar correlation could be obtained by a user-based rather than hashtag-based polarity we also assigned a polarity score to each user that was *active* in a given week. An active user had to use a hashtag that was used by at least three users in that week. Averaging across all these hashtags the user used, a within-week leaning was then computed for each user, regardless of their retweet-derived leaning. This leaning was then mapped to a polarity value in [0,1] in the same way as for hashtags. Figure 3 shows the corresponding user polarity over time, as well as the number of distinct, active users in a given week. Neither this user polarity nor the number of active users reveal a relationship to the violent events. Interestingly, only the hashtag polarity derived from collective language usage (Figure 2) seems to correlate strongly.

VII. TWEETING IN UNISON WITHIN EACH CAMP?

To quantify how much "in unison" the different groups were we looked at the cosine similarity for the sets of hashtags used between user pairs. A large within-group similarity would indicate a stronger vocabulary coherence and a large group-with-seed-users similarity would indicate a stronger influence of "party lines".

For the Islamist users we observed a pairwise average similarity of .13, with a median of .00 - indicating a small but fairly coherent group of users pulling up the average. For secularists the roles were reversed with an average similarity of .16, smaller than the median of .33 - hinting at a group of "completely out of sync" users bringing down the average. Surprisingly for us, the hashtag agreement with the seed users was larger for the secular users, with average and median values of .14 and .14 respectively, as compared to .11 and .03 for Islamists.

At least part of these observations could be explained by different levels of hashtag adoption overall. For example, a user who *never* uses hashtags will have a similarity of 0.0. To remove such effects we redid the analysis for users with at least 31 (= median) distinct hashtags. The first set of trends persisted unchanged. Namely, two non-seed Islamist users had an average (median) cosine similarity of .16 (.11) respectively - indicating a heavy tail in the distribution. For non-seed secularists the average (median) was .16 (.21) - again hinting at out-of-sync users. The similarity of these active hashtag users with respect to the seed users changed and was now more similar for both sets, with averages (medians) of .11 (.12) and .16 (.09) for secularists and islamists respectively.

Figure 4 shows the distribution of similarities for users actively using hashtags to their respective seed users. The left plot suggests a bimodal distribution with one mode corresponding to tweeps less similar to the seed users and a second mode with users more closely "following the party narrative".

VIII. WHAT DO THEY READ ONLINE?

Other researchers have used data from Twitter to study the bias of popular media outlets [23], [24]. We use a similar approach, though we use tweeted URLs rather than follower information to gauge audience bias. Concretely, we apply the same polarity formula from hashtags to URL (sub-)domains.

[7] http://blog.twitter.com/2012/03/twitter-now-available-in-arabic-farsi.html

[8] See http://en.wikipedia.org/wiki/Timeline_of_the_2011_Egyptian_revolution_under_Mohamed_Morsi_\%28July\%E2\%80\%93October_2012\%29#October and http://en.wikipedia.org/wiki/Timeline_of_the_2011_Egyptian_revolution_under_Mohamed_Morsi_\%28from_November_2012\%29#November for a detailed breakdown. Major events in the period before our study are also described at http://www.aljazeera.com/indepth/spotlight/egypt-one-year-on/2012/01/2012124164311269954.html.

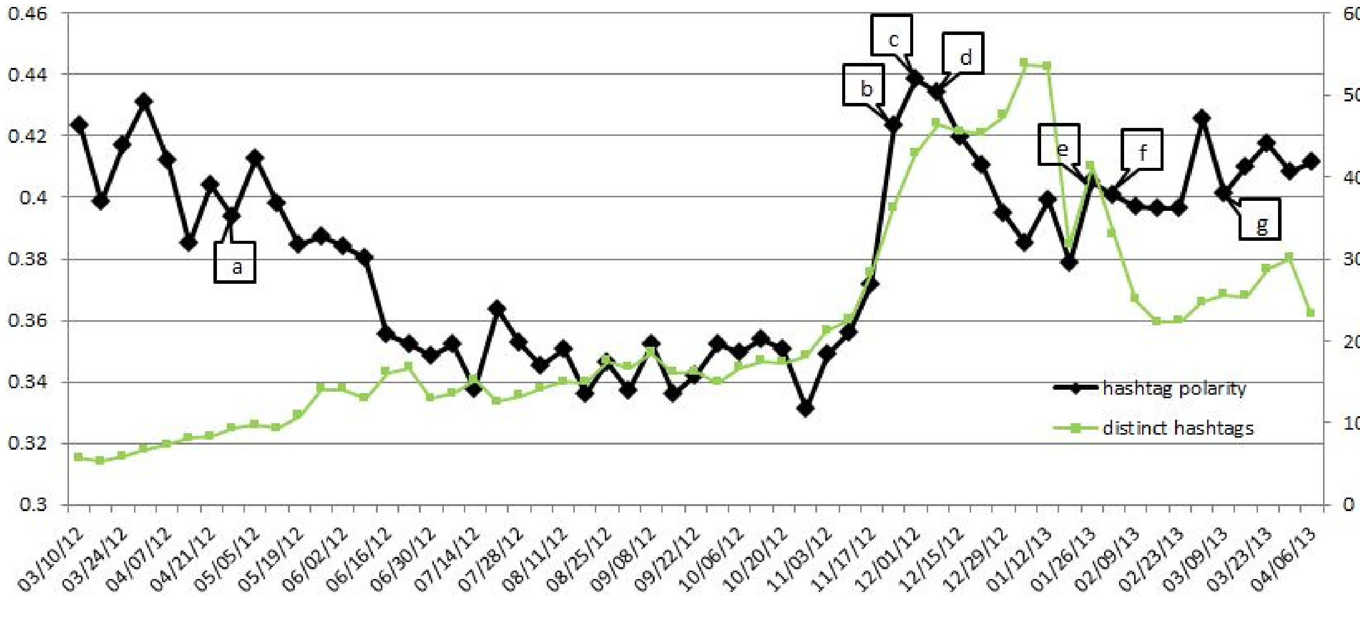

Fig. 2. Overall hashtag polarity over time.

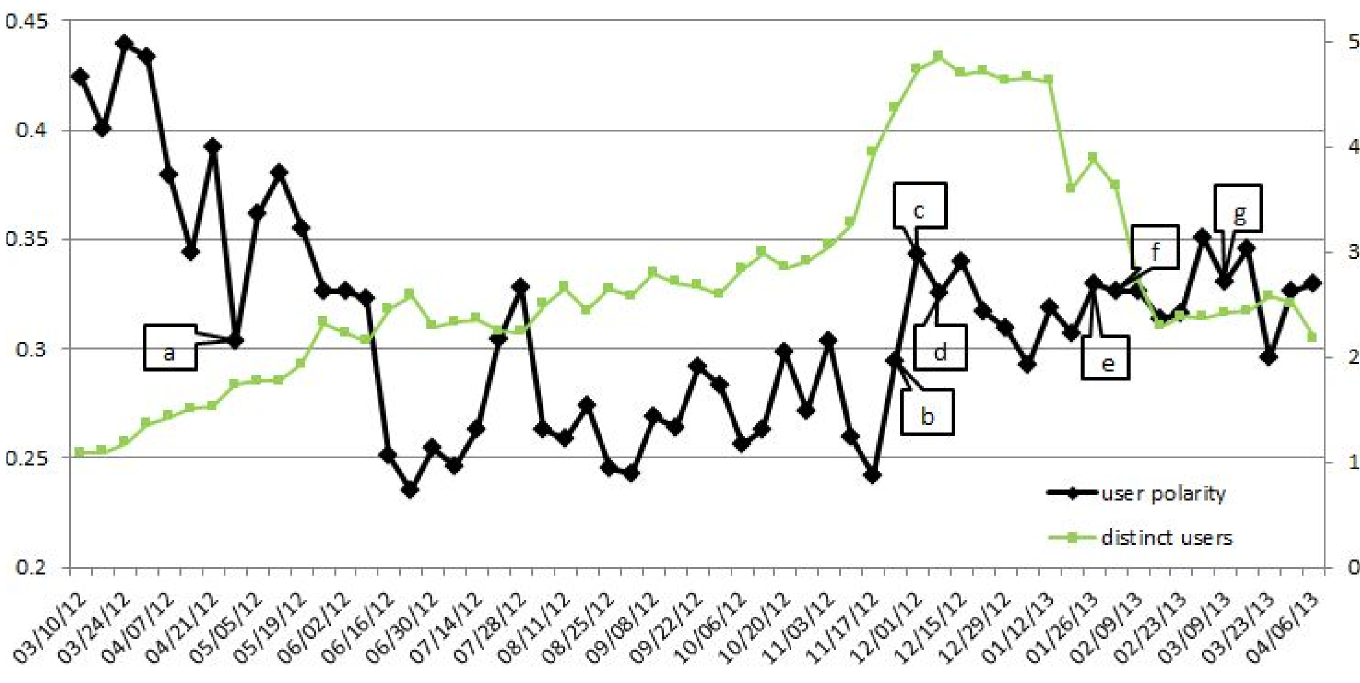

Fig. 3. Overall user polarity over time.

Inspecting the most polarized domains the reader will find stereotypical differences. Interestingly, a large fraction of strongly polarized domains consist of blogs, in line with other work that shows the importance of this media form for political discourse in Egypt.

Though neither of the two Al Jazeera domains is highly polarized, we looked at them in detail in an attempt to quantify existing observations about the differences between the Arabic and English versions [25]. We found that the Arabic version aljazeera.net has a very slight Islamist leaning (0.61i, 261i vs. 516s tweeters) whereas the English one aljazeera.com has a secularist leaning (0.63s, 21i vs. 114s tweeters).

As another case study, URLs from youtube.com can be found in the majority of tweeps' streams and, correspondingly, the leaning for the domain is very close to center (0.52s, 1,296i vs. 4,208s), though highly polarized individual videos exist.

IX. How Do They Connect?

For the case of US politics the strong community structure of blog inter-linkage [2] and Twitter retweet networks [1] has been observed before. Here, we wanted to see if something similar holds for retweet and mention networks in Egyptian politics.

Islamists			Secularists		
misr25.tv	0.957i	18/1	mella5er.blogspot.co.uk	0.977s	1/194
forum.islamstory.com	0.948i	33/4	saveegypt.net	0.975s	0/58
lojainiat.com	0.946i	15/1	sandmonkey.org	0.972s	0/51
islamstory.com	0.942i	118/21	weekite.blogspot.com	0.970s	0/48
mustafahosny.com	0.933i	14/1	arabawy.org	0.970s	0/44
melhamy.blogspot.com	0.927i	27/5	amrabdelazez.blogspot.com	0.968s	0/44
albayan.co.uk	0.926i	14/2	alalamalislami.com	0.958s	0/33
hussein-hamed.com	0.925i	18/3	akhbarbaladna.net	0.958s	0/33
egy-nahda1.blogspot.com	0.925i	18/3	mcndirect.com	0.957s	0/32
dostourmasr2012.com	0.923i	21/4	insaneyat.wordpress.com	0.951s	0/28

TABLE IV. TOP 10 POLARIZED DOMAINS FOR BOTH CAMPS. THE COUNTS ARE FOR DISTINCT ISLAMIST/SECULAR TWEEPS RESPECTIVELY.

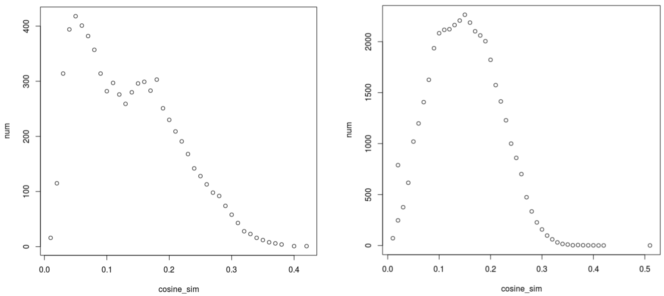

Fig. 4. User-user hashtag usage cosine similarities for tweeps with at least 31 (the median) distinct hashtags. Left is for Islamist tweep-seed similarity, right for secularist tweep-seed similarity.

To this end we constructed graphs where nodes were tweeps and (binary) directed edges were present if a user (i) retweeted or (ii) mentioned another user. Table V shows the label congruence probability for retweet edges. Results are for per-edge micro-averages but per-user macro-averages differ by less than 2% in the direction of less "mixing". Results for mention rather than retweet edges are also virtually identical, again with about 2% differences towards more cross-ideology edges.[9] As the class distribution of users is biased (with 5,215/1,719/154 users labeled as Secularist/Islamist/Center respectively) there is a higher chance to retweet a secularist than an Islamist. To better understand the strength of the preference for intra-group retweets we normalized the link distribution by the number of retweets each group attracts. The corresponding results are shown in parentheses in Table V.

	Retweeted User		
Retweeter	Center	Islamist	Secularist
Center	3.6 (44.5)	22.5 (31.6)	73.8 (23.9)
Islamist	3.9 (36.8)	46.3 (50.6)	49.9 (12.6)
Secularist	1.7 (31.6)	12.1 (25.9)	86.2 (42.5)

TABLE V. EDGE STATISTICS ABOUT TWEEPS RETWEETING OTHER NON-SEED USERS. NUMBERS IN PARENTHESES ARE FOR THE CASE WHEN THE GROUP SIZES OF THE RETWEETED USERS ARE NORMALIZED FOR THE OVERALL NUMBER OF RETWEETS THEY ATTRACT.

Figure 5 shows the corresponding retweet network. Seed users are not included in the set. Red nodes indicate Islamists, blue nodes secularists with intra-ideology edges colored correspondingly. Yellow nodes indicate split center nodes and yellow edges either center or cross-ideology edges. Though not perfectly bipolar, there is a clearly visible grouping of users of matching inferred ideology. Still, the polarization in terms of retweet network structure appears weaker than in the U.S. [1].

[9]We included retweets in the mentions as they also include @username.

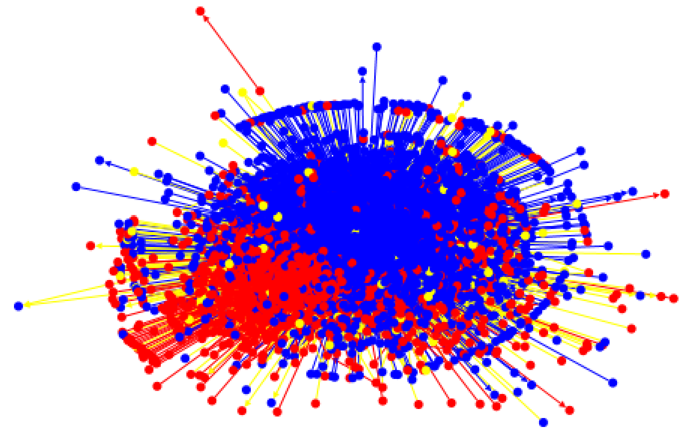

Fig. 5. Visualization of the retweet network obtained with NodeXL [26] and the Fruchterman-Reingold [27] layout algorithm. Four nodes not connected to the giant component were removed. All 6,579 remaining connected users users are included. Edges pointing from retweeting user to source user.

X. POLITICAL ISLAMISM VS. RELIGIOSITY AND OTHER CONNECTIONS

So far our analysis has looked at things such as user characteristics, polarization over time or network structure. Here, we focus on what other forms of behavior or, more concretely language use, correlate with identification with a political camp.

Our starting point here is the tendency in the West to equate political Islamism and devoutness to Islam or, at least, to see them as going hand in hand. We wanted to see if we can find statistical evidence for or against the hypothesis that these two are indeed linked. As a proxy for religiousness we chose a basic dictionary-based approach where we compiled a list of 609 terms (381 Arabic and 228 English) related to Islam. The first column in Table VI shows examples. Each tweet was matched against this dictionary and matching tweets marked as "religious". We chose a similar dictionary-based approach to look for mentions of derogatory terms referring to other religions, in particular Judaism, and to donations and charitable acts[10]. Before matching against the dictionaries, terms were normalized for variants of dialects and other Arabic-language specifics [28].

To quantify correlations between variables we first map values to their percentile ranks, mapping the median value to 0.5 and so on. This is done to cope with monotone but

[10]Almsgiving or "zakt" is one of the five pillars of Islam with devout Muslims donating significant fractions of their income. See http://en.wikipedia.org/wiki/Zak\%C4\%81t or http://islam.about.com/od/zakat/p/zakat.htm.

Religious	Derogatory	Charitable
اللّه (Allah) (1482195)	كَافر (unbeliever) (31668)	صدقه (charity) (16534)
رب (god) (914783)	تجنب (avoid non muslims) (5256)	grant (8129)
كن (be) (681039)	سَاخر (joking about something) (3989)	زكَاة (charity) (3055)
اَله (god) (349418)	الملحد (the atheist) (3840)	donat (2658)
أَحد (the only one) (344802)	الكفَار (the non muslims) (3685)	خيرية (charity) (2378)
نص (words of koran) (320775)	التكفير (the non muslim) (3236)	الزكَاة (the charity) (1830)
جن (genie) (301073)	pig (1810)	الصدقة (the charity) (1498)
ايه (verse) (290847)	متشَائِم (pessimist) (1698)	charit (1436)
سلَام (peace) (268904)	خنزير (pig) (1499)	الخيرية (the charity) (983)
عرف (know) (261256)	المنشق (the deserter) (1078)	صدقَات (charities) (859)

TABLE VI. TOP 10 FROM EACH OF THE LISTS IN TERMS OF NUMBER OF MENTIONS. THE NUMBERS IN THE PARENTHESES INDICATE THE NUMBER OF TWEETS CONTAINING THESE WORDS.

non-linear dependencies. For example, a random variable Y that satisfies *exactly* $Y = X^2$ for a variable X uniformly distributed in $[0, 1]$ only reveals a linear Pearson product-moment correlation coefficient (also known as "r") of 0.87 between X and Y, despite the perfect, deterministic dependence. After mapping values to their percentile ranks the correlation coefficient becomes 1.0. We applied this transformation to all variables, except gender which is discrete by nature. Table VII shows the correlation coefficients obtained this way.

Given the strong link between Islam and the Arabic language[11] results should not naively be compared across the two languages. Hence, we extended our analysis to see if, once language differences are accounted for, trends still persist. Similarly, certain effects could be explained by hidden dependencies on gender or overall activity. As our focus is polarization, we look in detail at the impact of the variable indicating a user's polarization on (i) the usage of religious terms, (ii) the usage of charity-related terms, and (iii) the usage of derogatory terms for other religions. As a model we choose a standard linear regression of the following form to which we applied least-squares model fitting. $y_{\text{targ}} = a_{\text{pol}} \cdot x_{\text{pol}} + a_0 + a_{1a} \cdot x_{1a} + a_{1b} \cdot x_{1b} + a_2 \cdot x_2 + a_3 \cdot x_3$

The variable of interest here is x_{pol} – the percentile of the polarization score, ranging from 0.0 (= only secular seed users retweeted) to 1.0 (= only Islamist seed users retweeted). The other variables are added one at a time: a fixed-term constant; x_{1a} and x_{1b} – percentiles of fractions of English and Arabic language tweets (added at the same time); x_2 – percentile of the total number of tweets for the user; x_3 – gender (-1 = male, 0 = unknown, +1 = female). After adding a new variable, the new coefficients a_i are re-learned and the value of a_{pol} is recorded in Table VIII.

	Feature set			
	pol.+const.	...+lang.	...+tweets	...+gender
religiosity	.23***	.11***	.08***	.07***
charity	.06***	.08***	.04**	.05*
derogatory	-.05***	-.06***	-.11***	-.07***

TABLE VIII. VALUES OF THE POLARIZATION COEFFICIENT a_{pol} IN A LINEAR REGRESSION MODEL TO PREDICT THE PERCENTILE OF THE TARGET VARIABLE. *, ** AND *** INDICATE P-SIGNIFICANCE LEVELS OF $< .05$, $< .01$ AND $< .001$ RESPECTIVELY.

For all three target variables the sign of a_{pol} is consistent across all four different models, indicating a clear direction of the effect, despite possible correlation with other factors. Concretely we observe the following, with or without the inclusion of additional factors.

(i) The relative usage of religious terms increases as a user is closer to the Islamist end of the polarity.

(ii) The relative usage of charity-related terms increases in the same direction.

(iii) The relative usage of derogatory terms referring to other religions *decreases* in the same direction.

Though we were expecting observations (i) and (ii) to hold, observation (iii) goes against the common (Western) wisdom that followers of the Muslim brotherhood are more likely to use religious hate speech. This is at least partly surprising as Mohamed Morsi famously made derogatory references to Jews in public speeches before becoming president.[12]

XI. CONCLUSIONS

We presented a quantitative study on Islamist vs. secular polarization in Egypt on Twitter. Starting from a set of annotated seed users we followed retweet edges to obtain a set of 7k users with inferred political orientation. This set then allowed us to study polarization for hashtags or domains, as well as analyzing their communication and networking patterns.

We found strong indications that a measure of global hashtag polarization, related to the overlap between hashtags used by the two political sides, works as a "barometer for tension" with high values coinciding with periods of violent outbreaks. Given a small but steady buildup of polarity before the unexpected outbreak of violence in late November 2012, there *might* be forecast potential and we plan to explore this further in the future. A similar user-based rather than hashtag-based measure did not reveal the same pattern.

Using hand-crafted dictionaries we could show consistent trends linking proximity to political Islamism to increases in word usage related to religion or charities, but a *decrease* in derogatory terms for other religions. We have been careful not to draw conclusions about *causal* connections and have restricted ourselves to observing correlations. However, time information and the fact that the cause precedes the

[11] See http://islam.about.com/od/arabiclanguage/a/arabic.htm or http://en.wikipedia.org/wiki/Arabic_language#Arabic_and_Islam for an introduction.

[12] http://www.nytimes.com/2013/01/15/world/middleeast/egypts-leader-morsi-made-anti-jewish-slurs.html

	arabic	charity	derogatory	english	gender	religious	polarity	tweets
Arabic	1.0^{***}	0.108^{***}	0.239^{***}	-0.499^{***}	-0.026^{***}	0.485^{***}	0.031^{**}	-0.061^{***}
charity	0.108^{***}	1.0^{***}	0.192^{***}	$--$	0.015^{***}	0.115^{***}	0.129^{***}	0.176^{***}
derogatory	0.239^{***}	0.192^{***}	1.0^{***}	-0.148^{***}	-0.031^{**}	0.418^{***}	-0.004^{*}	0.113^{***}
English	-0.499^{***}	$--$	-0.148^{***}	1.0^{***}	0.063^{***}	-0.432^{***}	-0.061^{***}	$--$
gender	-0.026^{***}	0.015^{***}	-0.031^{**}	0.063^{***}	1.0^{***}	$--$	$--$	-0.107^{***}
religious	0.485^{***}	0.115^{***}	0.418^{***}	-0.432^{***}	$--$	1.0^{***}	0.086^{***}	-0.046^{***}
polarity	0.031^{**}	0.129^{***}	-0.004^{*}	-0.061^{***}	$--$	0.086^{***}	1.0^{***}	0.089^{***}
tweets	-0.061^{***}	0.176^{***}	0.113^{***}	$--$	-0.107^{***}	-0.046^{***}	0.089^{***}	1.0^{***}

TABLE VII. CROSS-TABULATION OF LINEAR CORRELATIONS BETWEEN PERCENTILE-TRANSFORMED VARIABLES (EXCEPT GENDER). A $--$ INDICATES NO SIGNIFICANCE. *, ** AND *** CORRESPOND TO 5%, 1% AND 0.1% RESPECTIVELY.

consequence could lead to stronger claims with the help of Granger Causality [29] and related tools. Concerning the set of variables studied, additional ones could be included. Following the goal of [9] estimators of a user's income or education level could be constructed, possibly through the monitoring of keywords related to employment. Or the variables already present could be improved by, e.g., using tweeting behavior around prayer times to improve the measurement of a tweep's degree of devoutness. Additionally, crowd-sourcing could be used to overcome shortcomings of static dictionaries, without suffering undue noise from hard machine learning problems. In the setting of Egypt, views related to the Copts are arguably of a higher domestic relevance than views related to Jews. For both groups, however, name dictionaries could go a long way to obtain user sets with a likely affinity to either group.

In this work, we used a notion of polarity that relates to a 1-dimensional political spectrum similar to the US left-vs.-right polarization. Another commonly used notion of the term refers to *sentiment analysis*. In future work, we plan to combine the two both for Arabic [30] and English [31]. For example, when the sentiments of the two sides on a given issue drift apart then this might indicate potential for conflict. Similar, we deem it interesting to study the sentiments attached to, say, the US or Israel and their distribution between the two sides.

REFERENCES

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

[2] L. A. Adamic and N. Glance, "The political blogosphere and the 2004 u.s. election: divided they blog," in *LinkKDD@KDD*, 2005, pp. 36–43.

[3] M. Pennacchiotti and A. Popescu, "Democrats, republicans and starbucks afficionados: user classification in twitter," in *KDD*, 2011, pp. 430–438.

[4] K. T. Poole and H. Rosenthal, "A spatial model for legislative roll call analysis," *AJPS*, pp. 357–384, May 1985.

[5] I. Weber, V. R. K. Garimella, and E. Borra, "Mining web query logs to analyze political issues," in *WebSci*, 2012, pp. 330–334.

[6] I. Weber, V. R. K. Garimella, and A. Teka, "Political hashtag trends," in *ECIR*, 2013, pp. 857–860.

[7] M. Laver, K. Benoit, and J. Garry, "Extracting policy positions from political texts using words as data," *APSR*, vol. 97, no. 02, pp. 311–331, 2003.

[8] J. Slapin and S. Proksch, "A scaling model for estimating time-series party positions from texts," *AJPS*, vol. 52, no. 3, pp. 705–722, 2008.

[9] T. Mostak, "Social media as passive polling: Using twitter and online forums to map islamism in egypt," Master's thesis, Center for Middle Eastern Studies Harvard University, 2012.

[10] M. Lim, "Clicks, cabs, and coffee houses: Social media and oppositional movements in egypt, 2004–2011," *JC*, vol. 62, pp. 231–248, 2012.

[11] K. Starbird and L. Palen, "(how) will the revolution be retweeted?: information diffusion and the 2011 egyptian uprising," in *CSCW*, 2012, pp. 7–16.

[12] A. M. Attia, N. Aziz, B. Friedman, and M. F. Elhusseiny, "Commentary: The impact of social networking tools on political change in egypts "revolution 2.0"," *ECRA*, vol. 10, pp. 369–374, 2011.

[13] O. Oh, C. Eom, and H. Rao, "Collective sense-making through the twitter service during the 2011 egypt revolution," in *ICIS*, 2012.

[14] S. Khamis, "The transformative egyptian media landscape: Changes, challenges and comparative perspectives," *IJOC*, vol. 5, pp. 1159–1177, 2011.

[15] N. A. Azab, "The role of the internet in shaping the political process in egypt," *IJEP*, vol. 3, pp. 31–51, 2012.

[16] A. Choudhary, W. Hendrix, K. Lee, D. Palsetia, and W.-K. Liao, "Social media evolution of the egyptian revolution," *CACM*, vol. 55, no. 5, pp. 74–80, 2012.

[17] M. el Nawawy and S. Khamis, "Political blogging and (re) envisioning the virtual public sphere: Muslim christian discourses in two egyptian blogs," *IJPP*, vol. 16, no. 2, pp. 234–253, 2011.

[18] B. Etling, J. Kelly, R. Faris, and J. Palfrey, "Mapping the arabic blogosphere: politics and dissent online," *NMS*, vol. 12, no. 8, pp. 1225–1243, 2010.

[19] B. Al-Ani, G. Mark, J. Chung, and J. Jones, "The egyptian blogosphere: a counter-narrative of the revolution," in *CSCW*, 2012, pp. 17–26.

[20] P. N. Howard and M. M. Hussain, *Democracy's Fourth Wave? Digital Media and the Arab Spring*. Oxford University Press, 2013.

[21] A. Al-Kandari and M. Hasanen, "The impact of the internet on political attitudes in kuwait and egypt," *Telematics and Informatics*, vol. 29, pp. 245–253, 2012.

[22] R. Srinivasan, "Bridges between cultural and digital worlds in revolutionary egypt," *TIS*, vol. 29, no. 1, pp. 49–60, 2013.

[23] J. Golbeck and D. Hansen, "Computing political preference among twitter followers," in *CHI*, 2011, pp. 1105–1108.

[24] J. An, M. Cha, K. P. Gummadi, J. Crowcroft, and D. Quercia, "Visualizing media bias through twitter," in *SocMedNews*, 2012.

[25] M. M. Kraidy, "Al jazeera and al jazeera english: A comparative institutional analysis," University of Pennsylvania, Tech. Rep., 2008, http://repository.upenn.edu/ascpapers/274.

[26] M. A. Smith, B. Shneiderman, N. Milic-Frayling, E. Mendes Rodrigues, V. Barash, C. Dunne, T. Capone, A. Perer, and E. Gleave, "Analyzing (social media) networks with nodexl," in *C&T*, 2009, pp. 255–264.

[27] T. M. J. Fruchterman and E. M. Reingold, "Graph drawing by force-directed placement," *Software: Practice and Experience*, vol. 21, no. 11, pp. 1129–1164, 1991.

[28] K. Darwish, W. Magdy, and A. Mourad, "Language processing for arabic microblog retrieval," in *CIKM*, 2012, pp. 2427–2430.

[29] C. W. J. Granger, "Investigating causal relations by econometric models and cross-spectral methods," *Econometrica*, vol. 37, no. 3, pp. 424–438, 1969.

[30] M. Korayem, D. J. Crandall, and M. Abdul-Mageed, "Subjectivity and sentiment analysis of arabic: A survey," in *AMLTA*, 2012, pp. 128–139.

[31] B. Pang and L. Lee, "Opinion mining and sentiment analysis," *FTIR*, vol. 2, no. 1-2, pp. 1–135, 2007.