Journalists and Twitter:
A Multidimensional Quantitative Description of Usage Patterns

Mossaab Bagdouri

Department of Computer Science
University of Maryland
College Park, MD, USA
mossaab@umd.edu

Abstract

We conduct a large scale quantitative comparison of the usage pattern of a microblogging service by journalists, news organizations, and news consumers. Through two statistical tests of eighteen numerical features over 5,000 news producers and 1 million news consumers, we find that Arab journalists and English news organizations tend to broadcast their tweets to a large audience; that English journalists adopt a strategy of targeted and engaging communication; that journalists are more distinguishable in the Arab world than in the European English speaking countries; that print and radio journalists have a very dissimilar behavior while the television ones share some characteristics with each of them; and that British and Irish journalists are similar to a large extent. This paper is the first to provide a multidimensional bird's-eye view on the usage pattern of journalists over Twitter.

Introduction

Twitter is becoming a primary platform for breaking news (Phuvipadawat and Murata 2010; Hu et al. 2012). This makes it useful to both the traditional news producers (i.e., journalists and news outlets) and their consumers (e.g., readers). But Twitter is also dramatically changing the traditional roles of news producers and consumers. Journalists, for instance, interact more with the "crowds" to gather more information from ordinary people (Zak 2012), some of which have started to play the role of "citizen journalists" (Bowman and Willis 2003).

Understanding how journalists use Twitter will help us infer how their function is changing in our society. Several studies have been conducted to analyze how journalists interact with different stakeholders in the process of news production within the context of social media. They focus on particular regions, specialties or events, and vary between qualitative observations on a small set of journalists (typically, less than a dozen), and small scale quantitative summaries of journalists' tweets (often few thousands). The findings of these studies become, thus, difficult to generalize to the wide scale of the overall population of journalists.

Our approach in this work is to collect a large dataset of Twitter accounts related to the production and consumption of news. These accounts span over two regions (Arab world, and European English speaking countries), three user categories (journalists, news organizations, and consumers), and three media types (print, radio, and television). Through a series of statistical comparisons, we aim to answer these questions: Do journalists engage personally with their audience compared to news organizations? Do observations about English journalists—who are typically studied in previous work—apply to journalists from different regional, cultural, and lingual backgrounds (e.g., Arab journalists)? Do journalists use Twitter in a manner dissimilar from news consumers, and do these (dis)similarities hold across different regions? Are journalists a homogeneous group, or do they differ as a function of the type of the news outlet they work for? To which extent do journalists who speak the same language, but belong to different countries share similar characteristics?

Our work provides a multidimensional bird's-eye view of the journalists' use of Twitter that complements prior studies. Our findings can be used to inform the design of more customized tools for this group of professionals.

The remainder of this paper is organized as follows. We start by surveying some of the prior works that study the interaction of journalists with social media, or perform large scale analyses of news related tweets. Next, we describe the data we collected and the methodology used for the analysis of its features. Then, we perform a series of comparisons between several groups of the Twitter accounts of news producers and consumers. We finally summarize our findings in a concluding section.

Related Work

The last decade has seen a large number of studies conducted about the use of social media by professional and "citizen" journalists. Lasorsa, Lewis, and Holton (2012) analyzed 22,000 tweets taken over a period of two weeks from 500 journalists who have a high number of followers. They focused their analysis on the aspects of impartiality, gatekeeping, accountability, and transparency. To conduct the analysis, they split their sample into a set of *elite* journalists (i.e., those who work for a major news outlet), and *non-elite* journalists. They found that the former group is more resisting to

changes that accompany the emergence of microblogs, and are less motivated to share their opinions and engage their readers. As Dailey and Starbird later illustrated (2014), journalists who are less prominent are more encouraged to adopt the norms and practices of social media. Dailey and Starbird explored the role that journalists play in the time of a natural disaster. They conducted eleven interviews to debrief journalists, volunteers, and public officers about the use of a liveblog to gather and report news before, during and after the Hurricane Irene hit the rural area of Catskills, NY in August 2011. They showed that journalists rapidly adapt to the imposed constraints by coordinating the efforts of the flow of information from and to different individuals and entities. Parmelee (2013) took a grounded theory approach to code eleven interviews with political journalists working for American newspapers. He found these journalists to be more conservative about publicly sharing their opinions than what Lasorsa, Lewis, and Holton had reported. Many other papers studied the interaction between journalists and Twitter for specific aspects such as usefulness (Ahmad 2010; Vis 2013), branding (Molyneux 2014), and humor (Holton and Lewis 2011), in certain regions (Bruns 2012; Verweij and van Noort 2014; Hermida, Lewis, and Zamith 2014), for some specific events (Hermida, Lewis, and Zamith 2014) or topics (Sheffer and Schultz 2010; Parmelee 2013). To the best of our knowledge, they are all limited by the scope and the size of the pool of journalists (handful to few hundreds) and their tweets (hundreds to thousands). De Choudhury, Diakopoulos, and Naaman (2012) trained a classifier that categorizes Twitter accounts into organizations, journalists/bloggers, ordinary individuals, and other. They applied it to eight random samples of 5,000 users, each related to some event that took place in 2011. They looked at some characteristics of these different user categories, such as the number of URLs contained in the tweets, the interactivity of the accounts with replies and retweets, and the number of questions they ask. Our work can be seen as an extension to that of De Choudhury, Diakopoulos, and Naaman; as we study additional features for a larger number of news producers (5,000 journalists and news organizations, with more than 13 million tweets), from four media types, across two languages and regions; and compare them against two billion tweets from more than a million of news consumers.

Several large scale experiments have been conducted to analyze news on Twitter. Olteanu et al. (2015) mined 30 million news articles and 2 billion tweets to compare the discussion of the topic of climate change between social media and news outlets. They took an iterative approach to select keywords relevant to this topic and extract the corresponding tweets and news articles. In their quantitative analysis they found that news outlets address this topic during disasters and governmental interventions, while Twitter popularizes individual actions. Castillo et al. (2014) studied the lifecycle of online news articles in social media. They mined more than 3 million visits and 200,000 social media reactions to model the lifetime of 600 articles. They showed that 20 minutes are sufficient to predict the overall traffic of a news article. The traffic that news articles receive was also a subject of the work of Diakopoulos and Zubiaga (2014), but with a focus on the topical content of the news. They found statistical evidence that socially deviant events have high chances of being (re)shared on Twitter. While these large scale works are related to the core interest of journalism (i.e., news), they do not focus on the journalists themselves.

Methods

We collect 13 million tweets of 5,358 Twitter accounts of journalists and news organizations, in addition to two billion posts from over one million of their connections. We extract eighteen features from those accounts and their tweets, and analyze their distributions over various groups using Welch and Kolmogorov-Smirnov statistical tests.

Data

We collect a list of news related contacts from the directory of the website http://media.info, maintaining only profiles with at least one Twitter account who have posted at least one public tweet. These profiles are distributed across three dimensions that consist of three countries (United Kingdom, Ireland, and Gibraltar); four media formats (radio stations, TV channels, newspapers, and magazines); and two profile categories (journalists and organizations). The categories within the same dimension are not mutually exclusive (e.g., a journalist might work for both a magazine and a radio station). For this reason, when we study a particular dimension, we omit the accounts that have more than one affiliation.

To compare journalists from different cultures and languages, we use a set of the Twitter accounts of 1,230 Arab journalists released by Bagdouri and Oard (2015). As there is no indication of any particular distribution of this set into some subgroups, we treat the whole set as if it was a single homogeneous group.

We assume the sets of Arab and English journalists and organizations to be representative of their respective populations, acknowledging that a better data, if available, would be one in which the Twitter accounts are randomly sampled from these populations. The combined set we collected contains 5,358 Twitter accounts of journalists and news organizations. We crawl all of their 13,140,449 publicly available tweets using Twitter API.[1] Table 1 summarizes the distribution of these accounts across the different dimensions we study in this paper.

We wish to extend our analysis to potential news consumers. Hence, we crawl a large set of the accounts of the audience that has some interest in news. An obvious place to find these people is the combination of the sets of friends and followers of the news related accounts, in addition to the users they mention in their tweets. Because this combination is extremely large, making it practically impossible to crawl, we limit ourselves to the accounts that have a bidirectional follower / friend relationship with any of our journalists, in addition to those who have been mentioned in any tweet of the journalists we have downloaded.

---

[1]The API returns up to 3,200 of the most recent tweets. http://dev.twitter.com/rest/reference/get/statuses/user_timeline

							EN	AR	EN	AR
English Journalists							**EN** **Org.**	**AR** **Journ.**	**EN** **Cons.**	**AR** **Cons.**
Country			Media Format							
UK	IE	GI	Newspaper	Magazine	Radio	Television				
2,285	343	16	402	73	1,690	403				
Total			2,648				1,480	1,230	1,100,510	310,974

Table 1: Distribution of Twitter accounts in the different groups of this study.

For each language (i.e., Arabic and English), we keep only users whose dominant language is that one. The final audience set contains 1,411,484 accounts and has 2,257,844,613 tweets. We recognize that some of these can also be of journalists. However, their prevalence would be extremely low given the size of the collection, and would not have a significant impact on the aggregate statistics. Hence, we ignore their potential existence, and name this group "news consumers." In addition, with the limit of 5,000 enforced by Twitter on the number of accounts that a single user can follow,[2] journalists might prefer to follow celebrities over ordinary people. Similarly, but for reasons other than Twitter's limits, journalists might be biased towards mentioning accounts that are highly active on Twitter. Our process for collecting the Twitter accounts of news consumers ignores these potential biases. This is a limitation of this work that can be addressed in the future by randomly sampling the news consumers from the set of the followers of journalists and news organizations.

Analysis

We compare our groups of Twitter accounts using two statistical tests based on eighteen features.

Features We select eighteen numerical features of the Twitter accounts to compare the groups we study. Some of these are directly extracted from the user profile. Others are observed and aggregated over the tweets. We cluster these features into six aspects.

- Audience perception: Three features indicate how the audience might perceive or judge an account: the number of followers (⧻*Followers*), the number of lists that contain that account (⧻*Lists*), and the presence of Twitter's blue verification badge (*%Verified*). All of these features are extracted directly from the user profiles, independently from the tweets.

- Audience reaction: Two features signal how users react to the tweets of a given account: retweets (⧻*In-Retweets*) and favorites (⧻*In-Favorites*). We denote by original tweets those that are not retweets. For each of these two features, we divide the sum of the reactions to all original tweets of an account by the number of its original tweets. The resulting value corresponds to the average reactions to the tweets of that account.

- Broadcast communication: Three features indicate how an account tends to *broadcast* their tweets. ⧻*Daily Tweets* is the average number of tweets published per day by a given account. We compute this value by dividing the number of tweets we crawled over the span of their publication time. This should give a more current view of the verbosity of the account compared to simply dividing the number of tweets indicated in the profile by its lifetime. ⧻*Hashtags* is the average number of hashtags per original tweet for a given account. Hashtags are useful to reach out to people interested in the corresponding topics. Some journalists and news organizations share links to their products, such as articles and TV shows. We capture this act of *branding* through ⧻*URLs*, which is calculated in a manner similar to the hashtags.

- Targeted communication: Some accounts opt for a strategy of targeted communication. They mention some users, ask questions, reply to others' tweets, and favorite or retweet them. We capture this behavior using five features. ⧻*Mentions* indicates the average number of mentions of other users per original tweet for a given account. *%Replies* of an account is the ratio of original tweets that are replies to other people. *%Questions* indicates the ratio of original tweets that include a question mark in its Arabic or English form. We recognize that some tweets with question marks are rhetorical, and that some questions do not have a question mark. Machine learning techniques could be used to estimate whether a tweet is an answer-seeking question (Zhao and Mei 2013; Hasanain, Elsayed, and Magdy 2014), but we choose here the simpler approach of relying solely on the existence of question marks. The fourth feature in this group is the ratio of retweets emitted by an account (*%Out-Retweets*). The fifth is the number of favorited tweets, which we extract directly from the user profile (⧻*Out-Favorites*).

- Personal communication: We hypothesize that the communication of journalists is more personal, whereas that of organizations is more official. We test this hypothesis by calculating the ratio of original tweets that contain any of the first-person singular pronouns *I*, *am*, *my*, and *mine*; and the first-person plural pronouns *we*, *our* and *ours*. We denote these two features by *%I* and *%We*, respectively.

- Publication medium: Twitter users can publish their tweets through various technologies, such as mobile devices, desktop computers, and third-party applications. We catch the medium of publication through the *source* field of the tweet, which is an arbitrary string of charac-
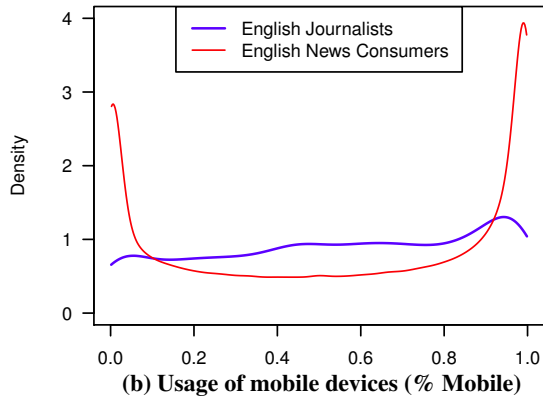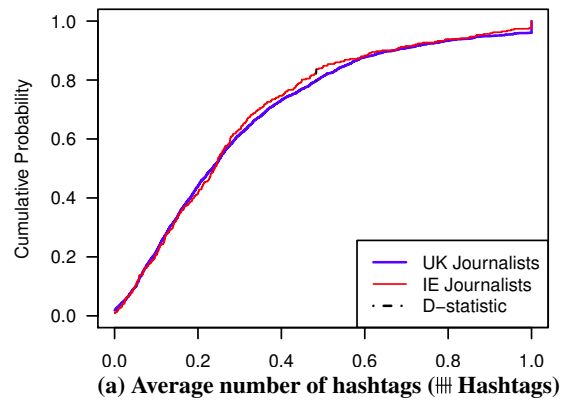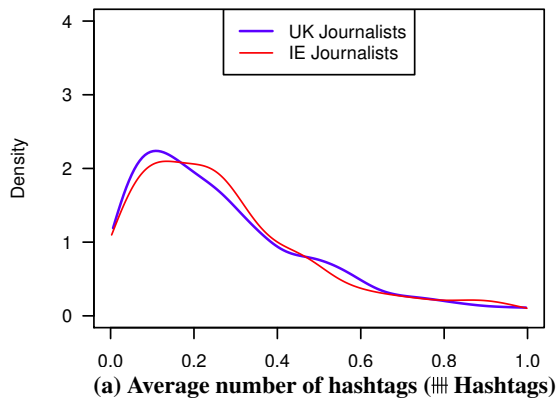
Figure 1: The density curves of two pairs of groups for two features.

Figure 2: Estimation of the D-statistic from two ECDFs.

ters. In the dataset we collected, we found 1,640 distinct sources. We sort them by decreasing order, and consider the most frequent sources that cover, together, at least 90% of all the tweets of the journalists and news organizations. We assign these 13 sources to three categories (%Mobile, %Desktop and %App), in addition to *Other*. For each of these categories, we compute, for a given account, the corresponding feature as a ratio out of the original tweets.

We focus our analysis only on these numerical features. Additional insights can be drawn from other features, such as the content of the tweets. We leave these to future work.

Statistical Tests This study is designed as a series of eight comparisons, each between two groups, using most or all of the eighteen features. Each group (e.g., English radio journalists) has its own distribution over a particular feature. The most straightforward statistical test we can run is, perhaps, Welch's t-test (Welch 1947). Welch extends Student's t-test—that is widely used to test the hypothesis of the equality of the means of two populations—by relaxing the assumption of the equality of the variances. We run Welch's t-test in each comparison of the means, and indicate the significance level with "***" for $p < 0.001$, "**" for $p < 0.01$, "*" for $p < 0.05$, and "-" where the test is not significant.

Unfortunately, Welch's t-test has two pitfalls. It assumes the normality of the distributions, and is limited to the comparison of the means. Figure 1 shows two comparisons for which the t-test finds no significant difference. In (a), the two distributions do indeed appear to be similar. But in (b), one group seems to have a uniform distribution, while the other is bimodal. To mitigate these limitations, we use the two-sample Kolmogorov-Smirnov (KS) test (Chakravarty, Roy, and Laha 1967). KS tests whether two samples are drawn from the same distribution. It does so by calculating a *D-statistic* that quantifies the differences between the two samples, and a p-value to indicate significance. Intuitively, D is the largest distance between the two empirical cumulative distribution functions (ECDF, which is a cumulative function estimated from the data) of the two samples. It takes a value of 0 when the distributions are identical, and 1 when they are completely different. Figure 2 depicts the ECDF of the same pairs of Figure 1. The *D-statistic* is significantly high in (b), and low in (a), which agrees with our observation that the distributions in (a) appear similar, while in (b) they do not. For the remainder of this paper, we report the value of the *D-statistic* in **underlined bold** when $p < 0.001$, **bold** when $p < 0.01$, *italic* when $p < 0.05$, and with no emphasis where no significance is observed.

	Feature	Journalists		Organizations		Jour vs. Org	
		Mean	SD	Mean	SD	KS	t-test
#	Followers	45K	410K	63K	329K	**0.27**	-
#	Lists	255	1,556	596	3,351	**0.27**	***
%	Verified	14.84	35.56	19.59	39.71	*0.05*	***
#	Daily Tweets	5.30	8.78	16.26	92.63	**0.27**	***
#	URLs	0.16	0.18	0.55	0.32	**0.58**	***
#	Hashtags	0.31	0.30	0.42	0.43	**0.15**	***
#	Mentions	0.90	0.39	0.50	0.43	**0.48**	***
%	Replies	44.21	19.96	11.70	14.12	**0.69**	***
%	Questions	13.50	7.12	12.22	9.25	**0.21**	***
%	Out-Retweets	22.29	15.61	19.82	17.42	**0.14**	***
#	Out-Favorites	942	3,462	753	2,364	**0.07**	*
%	I	4.01	3.93	0.54	1.00	**0.76**	***
%	We	0.36	1.69	1.31	2.99	**0.46**	***
%	Mobile	54.96	31.50	11.07	16.96	**0.65**	***
%	Desktop	29.05	26.79	30.46	29.97	*0.05*	-
%	App	6.18	16.75	28.35	32.24	**0.47**	***
#	In-Retweets	2.92	20.70	4.60	46.11	**0.23**	-
#	In-Favorites	3.80	30.14	3.43	24.49	*0.05*	-

Table 2: Journalists vs. organizations.

Results

We compare two groups of news producers (i.e., organizations and journalists), two groups of journalists from different cultures and regions (i.e., Arabic and English speaking countries), journalists with news consumers, three groups of journalists using different media (i.e., television, radio, and print), and two groups of journalists sharing the same language, but from different countries (i.e., Ireland and the United Kingdom).

Organizations Broadcast, Journalists Target

Both of the journalists and the news organizations are expected to disseminate news in Twitter. Do they do so in a similar manner? How do they interact with their audience? Is Twitter also a source of information for them? We answer these questions by comparing the distributions of the features of 2,648 English journalists against those of 1,480 English news organizations.

We start our analysis with the expectations of the audiences from these two groups of users. We first observe, in Table 2, that organizations have, on average, 39% more followers than journalists, as the number of followers are 62,680 and 45,037 respectively. But this difference is not statistically significant, as $p > 0.05$. The difference in the distributions is high ($D = 0.27$) and significant at $p < 0.001$. The same significant difference in the distributions is observed for the number of lists containing these accounts. But the mean for the organizations is much larger than that of the journalists (255 and 596 respectively), with a high significance ($p < 0.001$). These two features perhaps

suggest that people who want to get the news from Twitter expect to find them in the timelines of the organizations more than from the journalists. This preference is enhanced by Twitter's endorsement to the accounts of the organizations with the *verified* sign. In fact, this sign is present in 19.59% of the accounts of organizations, and only in 14.84% of those of the journalists. The difference in the mean is significant at $p < 0.001$. The distributions are, however, similar given the D-statistic value of 0.05.

We now turn to the count and the nature of the tweets. For each tweet published by a journalist, an organization publishes three, on average. The difference in the means and the distributions are both high ($D = 0.27$) and significant ($p < 0.001$). Not just the average daily tweets is higher for the organizations, but also the number of URLs per tweet. In fact, an organization shares a link in every other tweet (perhaps from its own website). But a journalist shares only one link in six tweets. The differences in the means and the distributions are also high and significant ($D = 0.58$ and $p < 0.001$). Multiplying the average number of URLs and the average number of daily tweets indicates that while a journalist shares less than one link per day, an organization shares about nine. This suggests that organizational accounts perhaps play the role of "news broadcast" more than the personal ones. This is also confirmed by the use of hashtags. In fact, organizations try to reach out to users interested in the corresponding topics by using 0.42 hashtag per tweet on average, which is 35% higher than the average number of hashtags per tweet for journalists.

On the other hand, journalists have a more targeted communication behavior. In fact, they mention other users more

	Feature	EN Journalists		AR Journalists		EN vs. AR	
		Mean	SD	Mean	SD	KS	t-test
#	Followers	45K	410K	70K	240K	**0.24**	*
#	Lists	255	1,556	310	845	**0.13**	-
%	Verified	14.84	35.56	8.46	27.83	**0.06**	***
#	In-Retweets	2.92	20.70	12.19	46.06	**0.37**	***
#	In-Favorites	3.80	30.14	3.69	12.88	**0.16**	-
#	Daily Tweets	5.30	8.78	12.52	72.99	**0.20**	***
#	URLs	0.16	0.18	0.28	0.27	**0.21**	***
#	Hashtags	0.31	0.30	0.43	0.59	**0.16**	***
#	Mentions	0.90	0.39	0.51	0.35	**0.44**	***
%	Replies	44.21	19.96	32.9	23.64	**0.24**	***
%	Questions	13.50	7.12	7.58	6.82	**0.45**	***
%	Out-Retweets	22.29	15.61	27.4	20.23	**0.15**	***
#	Out-Favorites	942	3,462	1,157	3,800	**0.07**	-
%	Mobile	54.96	31.50	51.91	37.37	**0.12**	*
%	Desktop	29.05	26.79	28.41	31.6	**0.14**	-
%	App	6.18	16.75	1.25	8.23	**0.27**	***

Table 3: English vs. Arab journalists.

often (0.9 per tweet, vs. 0.5 for news organizations); dedicate more than 44% of their tweets for replying to other people, compared to less than 12% for news organizations; retweet other users' tweets (22.29% of their tweets, vs. 19.82% for news organizations); ask more questions (13.50% vs. 12.22% respectively); and favor others' tweets (942 vs. 753 respectively). Finally, by looking at the first-person singular and plural pronouns, it appears that the journalists' tweets are more personal than those of the organizations. They use the first-person singular pronouns seven times more than the organizations (4.01% vs. 0.54%). By contrast, the usage of the first-person plural pronouns is, respectively, 0.36% and 1.31%.

Journalists and organizations also differ in the medium used to publish their tweets. In fact, while they both use a desktop in about 30% of the time, mobile is the preferred medium for journalists (54.95%), and organizations tend to use special Twitter applications for posting more than 28% of the tweets.

In summary, journalists and organizations play different strategies in disseminating their tweets. Journalists prefer targeting their communication and maintaining a personal engagement with their audience. Their higher ratio of questions suggest that they also might be using Twitter to gather information. On the other hand, organizations appear to prefer broadcasting their posts and avoiding the personal style. Both of these strategies seem to reach the same degree of success, though. In fact, each tweet from these two groups receives, on average, about the same number of favorites. Although organizations receive 63% more retweets than journalists, this difference is, however, insignificant. Suspecting that this difference might be just an artifact of the difference in the number of followers, we compute the Pearson correlation between the number of followers and that of

received tweets for the union of the Twitter accounts. We find a moderate correlation of 0.57 ($p \ll 0.001$). This confirms that the high number of retweets that the news organizations receive can be partially explained by the number of their followers, rather than by the communication behavior.

Arab Journalists Broadcast, English Journalists Target

We compare 1,230 Arab against 2,648 English journalists. As Table 3 shows, Arab journalists appear to have more followers than the English ones. The difference of the means (70,050 and 45,037 respectively) is significant at $p < 0.05$. The KS test also shows a moderate, but significant, difference at $p \ll 0.001$, with $D = 0.24$. The difference in #*Lists* is small and insignificant, as the means are 310 and 255 respectively, with $p = 0.16$. On the other hand, a larger proportion of English journalists (14.84%) has the verified sign compared with the Arab ones (8.46%). The difference in the mean is statistically significant with $p \ll 0.001$. However, the distributions appear to be similar ($D = 0.06$).

The broadcast communication behavior is evident for Arab journalists. They tweet more than twice as much as the English ones, share 75% more links, and use 39% more hashtags. The differences are, at the same time, medium and significant at $p < 0.01$ for both Welch's t-test and Kolmogorov-Smirnov test. English journalists, on the other hand, appear to have a more engaging communication. They mention other users 76% more often than the Arab ones, are 34% more likely to reply to others' tweets, and ask 78% more questions. A small significant difference of 23% is, however, observed for retweets in favor of Arab journalists. Favorites appear to be similar for both the means and the distributions.

		AR Journalists		AR Consumers		EN Journalists		EN Consumers		AR J. vs. C.		EN J. vs. C.	
	Feature	Mean	SD	Mean	SD	Mean	SD	Mean	SD	KS	t-test	KS	t-test
#	Followers	70K	240K	9K	69K	45K	410K	11K	239K	**0.42**	***	**0.34**	***
#	Lists	310	845	16	152	255	1,556	85	1,119	**0.57**	***	**0.39**	***
%	Verified	8.46	27.83	0.35	5.94	14.84	35.56	2.96	16.96	**0.08**	***	**0.12**	***
#	In-Retweets	12.19	46.06	3.82	50.06	2.92	20.70	4.27	411.20	**0.42**	***	**0.31**	*
#	In-Favorites	3.69	12.88	0.94	11.78	3.80	30.14	4.35	195.01	**0.41**	***	**0.26**	-
#	Daily Tweets	12.52	72.99	26.16	724.72	5.30	8.78	13.06	505.29	**0.11**	***	**0.20**	***
#	URLs	0.28	0.27	0.19	0.27	0.16	0.18	0.23	0.27	**0.29**	***	**0.13**	***
#	Hashtags	0.43	0.59	0.25	0.43	0.31	0.30	0.37	0.47	**0.16**	***	**0.09**	***
#	Mentions	0.51	0.35	0.48	0.43	0.90	0.39	0.76	0.48	**0.11**	**	**0.19**	***
%	Replies	32.90	23.64	35.01	28.23	44.21	19.96	41.17	26.1	**0.08**	**	**0.13**	***
%	Questions	7.58	6.82	5.26	6.12	13.50	7.12	11.74	9.03	**0.24**	***	**0.18**	***
%	Out-Retweets	27.40	20.23	36.13	28.20	22.29	15.61	24.15	20.59	**0.17**	***	**0.08**	***
#	Out-Favorites	1,157	3,800	871	3,949	942	3,462	1,159	6,475	**0.06**	**	**0.09**	**
%	Mobile	51.91	37.37	67.45	39.20	54.96	31.50	55.51	38.38	**0.27**	***	**0.13**	-
%	Desktop	28.41	31.60	14.29	28.41	29.05	26.79	24.97	30.46	**0.38**	***	**0.17**	***
%	App	1.25	8.23	0.43	5.10	6.18	16.75	5.16	16.54	**0.06**	***	**0.13**	**

Table 4: Journalists vs. news consumers.

We see a small difference in the use of mobile and desktop devices for these two groups (the t-test is insignificant for the latter). However, English journalists have a higher adoption of third-party applications with a 5-to-1 ratio.

For each original tweet, on average, Arab journalists receive over four times more retweets than the English ones do. This difference is both high and significant at $p < 0.001$ for the two tests. The favorites have no statistically significant difference in the mean, and a small significant difference for the distributions.

In summary, Arab journalists appear to broadcast more tweets, and their audience seems to react positively. However, we need to analyze whether this difference is a characteristic of the practice of journalism in the two regions, or is simply due to the general usage patterns of Twitter from people in these regions. This is the subject of the next subsection.

Arab Journalists Are More Distinguishable

This section compares journalists with a control group of news consumers. Because of the potential cultural and regional differences between people speaking Arabic and English, we compare journalists and news consumers of each language independently. Table 4 summarizes these two comparisons.

We note first the unsurprising observation that journalists are more likely to have a verified account (8.46% vs. 0.35% for Arabic, and 14.84% vs. 2.96% for English). They also have more followers (respectively 70k vs. 9k, and 45K vs. 11K), and are included in more lists (310 and 255, vs. 16 and 85). In addition, the differences in the means and in the D-statistic show that journalists are more distinguishable than news consumers in the Arab world. For example, Arab journalists have verified accounts 22 times compared to Arab news consumers, while the chances for English journalists are only five times larger to have a verified account with comparison to English news consumers.

The reactions to the tweets are contrastive between the Arabic and English groups. Arab journalists receive much more reactions to their tweets compared to Arab news consumers for both retweets (12.19 vs. 3.82) and favorites (3.69 vs. 0.94). The English population experiences a different pattern, as the journalists receive less reactions than news consumers for retweets (2.92 vs. 4.27) and favorites (3.80 vs. 4.35).

For both of the Arab and English groups, journalists tweet, on average, less than half of what the control groups do (12.52 vs. 26.16, and 5.30 vs. 13.06, respectively). Arab journalists share more links (0.28 vs. 0.19) and use more hashtags (0.43 vs. 0.25). This is different from the pattern observed for English journalists who share less links per tweet (0.16 vs. 0.23) and use less hashtags (0.31 vs. 0.37). Thus, except for the tweeting frequency, there is no common pattern in the broadcast aspect that differentiates journalists from news consumers independently from the language.

Both of Arab and English journalists ask more questions, mention more users but retweet less often compared to the control groups. This suggests they might be using Twitter to gather information (with questions), from specific sources (using mentions), and are cautious about the validity of the information they convey (by minimizing retweets). Journalists have small differences with the control groups regarding replies and favorites, with opposite patterns between the Arabic and English groups. English journalists are more likely to reply to tweets, but less likely to favor them.

There is only small, and sometimes insignificant, differences between English journalists and their control group with respect to the tools used to disseminate tweets. How-

	Feature	EN Print		EN Radio		EN Television		Print vs. Radio		Radio vs. TV		Print vs. TV	
		Mean	SD	Mean	SD	Mean	SD	KS	t-test	KS	t-test	KS	t-test
#	Followers	26K	240K	43K	450K	68K	416K	**0.21**	-	**0.11**	-	0.11	-
#	Lists	321	920	200	1,785	353	1,156	**0.31**	*	**0.20**	*	0.13	-
%	Verified	24.63	43.13	9.35	29.12	24.07	42.80	**0.15**	***	**0.15**	***	0.01	-
#	In-Retweets	4.15	34.76	2.14	16.71	3.94	12.94	**0.30**	-	**0.25**	*	0.07	-
#	In-Favorites	3.61	44.29	3.43	26.86	4.51	22.07	0.09	-	**0.15**	-	*0.11*	-
#	Daily Tweets	6.26	9.97	4.99	7.85	4.82	10.50	*0.08*	*	0.06	-	**0.12**	*
#	URLs	0.26	0.21	0.14	0.16	0.14	0.15	**0.32**	***	0.04	-	**0.32**	***
#	Hashtags	0.26	0.28	0.30	0.28	0.41	0.37	**0.14**	**	**0.13**	***	**0.26**	***
#	Mentions	0.81	0.36	0.93	0.39	0.92	0.40	**0.15**	***	0.06	-	**0.12**	***
%	Replies	40.81	20.81	45.13	19.75	44.02	19.72	**0.13**	***	0.06	-	*0.10*	*
%	Questions	13.03	7.13	13.82	7.34	12.86	6.23	*0.08*	*	**0.09**	**	0.04	-
%	Out-Retweets	24.24	16.74	21.21	15.16	24.76	16.05	**0.10**	***	**0.12**	***	0.07	-
#	Out-Favorites	651	1,712	1,049	3,654	837	4,310	**0.15**	**	**0.12**	-	0.05	-
%	I	2.83	2.70	4.48	4.45	3.43	2.36	**0.32**	***	**0.18**	***	**0.19**	***
%	We	0.32	0.37	0.37	2.06	0.40	0.79	**0.10**	-	**0.11**	-	0.05	*
%	Mobile	41.86	28.77	56.93	31.87	60.99	29.11	**0.24**	***	*0.08*	*	**0.29**	***
%	Desktop	35.74	28.76	27.58	26.67	28.08	24.24	**0.15**	***	*0.08*	-	**0.15**	***
%	App	10.33	21.07	5.75	16.26	3.21	10.99	**0.11**	***	*0.09*	***	**0.15**	***

Table 5: Print, radio, and television journalists.

ever, we observe a stronger difference for the Arab groups, where journalists use less mobile platforms (51.91% vs. 67.45%) and more desktop ones (28.41% vs. 14.29%), compared to news consumers.

TV Journalists: Between Print and Radio

With the variety of media formats, journalists might gain diverse patterns in their use of Twitter. To study this effect, we compare three groups of English journalists: 1,690 who work at a radio station, 403 affiliated with a TV channel, and 479 who work for a print newspaper or magazine. Table 5 summarizes the statistical proprieties of these three groups, and the differences amongst them.

Despite the high difference in the average number of followers across the three media types, this difference is not significant. The KS test, however, shows that the difference in the distributions between print and radio journalists is at the same time the highest ($D = 0.21$), and the most significant ($p < 0.001$). This observation about the KS test holds also for the number of the lists containing these accounts, as well as for the ratio of verified accounts. Print and television journalists appear to be the most similar in these features as the t-test shows insignificant differences, and the D-statistic of the KS test is low. This indicates that the users following these accounts appear to cluster print and radio journalists far away from each other, and the television journalists in between them. The latter group is, however, closer to the print rather than to the radio journalists.

The differences in favorite tweets and retweets received by these groups of journalists are relatively small. But the reactions to the tweets of print and television journalists seem to be the most similar across the three comparisons, as the only significance we observe (at $p < 0.05$) is with the KS test for the favorites.

In terms of the broadcast behavior, each of these three groups uses less than two hashtags per day on average (i.e., the product of the number of hashtags and the number of tweets per day). However, print journalists do so through a lower number of hashtags (0.26) within a higher number of tweets (6.26). The opposite is observed for television journalists (0.41 and 4.82 respectively), and the radio journalists are positioned in the middle for these two features (0.30 and 4.99 respectively). Print journalists share 86% more links than the other ones. This difference is both high and significant at $p < 0.001$. In general, radio and television journalists are the most similar groups with respect to the broadcast features.

We do not observe a high difference in the means of the number of mentions and the ratios of replies, retweets and questions among the three categories of journalists. The similarity is the highest between print and television journalists, and the difference is most significant between print and radio ones. Regarding the technologies used to disseminate tweets, both television and radio journalists appear to prefer mobile devices, while print journalists use desktop and special applications more often.

As a summary across the various features, print and radio journalists are the most dissimilar groups. Television and radio journalists share properties related to the broadcast behavior and the medium for publishing tweets. television and print journalists share the characteristics of targeted communication, and the official presence in Twitter.

	Feature	UK		IE		UK vs. IE	
		Mean	SD	Mean	SD	KS	t-test
⧚	Followers	45K	341K	10K	26K	**0.17**	***
⧚	Lists	262	1,434	98	185	**0.10**	***
%	Verified	14.97	35.68	13.99	34.74	0.01	-
⧚	In-Retweets	2.96	20.16	1.32	4.14	**0.12**	***
⧚	In-Favorites	3.82	29.30	1.68	4.67	**0.23**	**
⧚	Daily Tweets	5.23	8.98	5.78	7.18	**0.11**	-
⧚	URLs	0.16	0.18	0.15	0.15	0.06	-
⧚	Hashtags	0.31	0.30	0.30	0.26	0.04	-
⧚	Mentions	0.90	0.39	0.94	0.37	0.07	*
%	Replies	43.83	20.01	46.37	19.41	*0.09*	*
%	Questions	13.75	7.29	11.96	5.67	**0.13**	***
%	Out-Retweets	21.90	15.49	24.31	15.81	0.07	**
⧚	Out-Favorites	862	3,443	1,442	3,427	**0.24**	**
%	I	4.19	4.10	2.87	2.25	**0.23**	***
%	We	0.37	1.81	0.30	0.40	0.06	-
%	Mobile	54.56	31.67	57.41	30.28	0.07	-
%	Desktop	29.45	26.82	26.47	26.20	0.06	-
%	App	6.10	16.80	6.78	16.63	0.04	-

Table 6: British vs. Irish journalists.

British and Irish Journalists Are Similar

In addition to clustering journalists by the type of the outlet they work for, or the language they speak, we can also compare journalists who speak the same language, but belong to different countries. Because the total number of journalists from Gibraltar is very small (i.e., 16), we limit our comparison to the 2,285 journalists from the United Kingdom (UK), and the 343 journalists from Ireland (IE). Table 6 summarizes our findings.

British journalists have more followers than the Irish ones (45K vs. 10K) and are included in more lists (262 vs. 98). They also receive more retweets (2.96 vs. 1.32) and favorites (3.82 vs. 1.68). But these facts are not surprising when we take in consideration the number of inhabitants in these two countries. Those groups of journalists are similar for the distributions of the verified accounts ($D = 0.01$). Irish journalists favor more tweets (1,442 vs. 862), ask less questions (11.96% vs. 13.75%), and have a low frequent use of first-person singular pronouns (2.87% vs. 4.19%). They mention other users and reply to their tweets about as much often as British journalists do. All of the other features experience a low or insignificant difference among the two groups. More specifically, we observe no statistically significant difference in the means of the broadcast features (i.e., ⧚*Daily Tweets*, ⧚*URLs*, and ⧚*Hashtags*), and only a small difference between their distributions (the D-statistic has a value of at most 0.11). There is also no statistically significant difference in the medium used to post tweets, for both of the statistical tests. Over all, British and Irish journalists appear to be similar to each other.

Conclusion

We conducted the largest study to date of the use of Twitter by news producers and consumers. We started from a public list of journalists and organizations from the European English speaking countries, and another one of Arab journalists. We crawled the tweets of these 5,000 Twitter accounts to extract eighteen summarizing features, as well as those of a large set of over a million of news consumers. We used Welch and Kolmogorov-Smirnov tests to quantify the statistical differences in eight comparisons. We found that the accounts of news outlets use an official style and share more links than journalists. These, instead, have a targeted communication that is engaging their audience. This suggests that they also use Twitter as a source of information. Arab journalists are less likely to have this communication pattern, though. They also are more distinguishable from news consumers than English journalists are. Journalists have differences across three media types. Print and radio journalists are the most dissimilar groups, and television journalists share some characteristics with each of them. For journalists speaking the same language but residing in different countries, only a small evidence of dissimilarities is observed.

Acknowledgment

References

Ahmad, A. N. 2010. Is Twitter a useful tool for journalists? *Journal of Media Practice* 11(2):145–155.

Bagdouri, M., and Oard, D. W. 2015. Profession-based person search in microblogs: Using seed sets to find journalists. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, CIKM '15, 593–602. New York, NY, USA: ACM.

Bowman, S., and Willis, C. 2003. We media: How audiences are shaping the future of news and information. `hypergene.net/wemedia`. Accessed: 2016-03-06.

Bruns, A. 2012. Journalists and Twitter : how australian news organisations adapt to a new medium. *Media International Australia Incorporating Culture and Policy* (144):97–107.

Castillo, C.; El-Haddad, M.; Pfeffer, J.; and Stempeck, M. 2014. Characterizing the life cycle of online news stories using social media reactions. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '14, 211–223. New York, NY, USA: ACM.

Chakravarty, I.; Roy, J.; and Laha, R. 1967. *Handbook of Methods of Applied Statistics*, volume 1. McGraw-Hill.

Dailey, D., and Starbird, K. 2014. Journalists as crowdsourcerers: Responding to crisis by reporting with a crowd. *Computer Supported Cooperative Work (CSCW)* 23(4-6):445–481.

De Choudhury, M.; Diakopoulos, N.; and Naaman, M. 2012. Unfolding the event landscape on Twitter: Classification and exploration of user categories. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, CSCW '12, 241–244. New York, NY, USA: ACM.

Diakopoulos, N., and Zubiaga, A. 2014. Newsworthiness and network gatekeeping on Twitter: The role of social deviance. In *Proceedings of the Eighth International Conference on Weblogs and Social Media*.

Hasanain, M.; Elsayed, T.; and Magdy, W. 2014. Identification of answer-seeking questions in Arabic microblogs. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, CIKM '14, 1839–1842. New York, NY, USA: ACM.

Hermida, A.; Lewis, S. C.; and Zamith, R. 2014. Sourcing the Arab Spring: A case study of Andy Carvin's sources on Twitter during the Tunisian and Egyptian revolutions. *Journal of Computer-Mediated Communication* 19(3):479–499.

Holton, A. E., and Lewis, S. C. 2011. Journalists, social media, and the use of humor on Twitter. *Electronic Journal of Communication* 21(1/2).

Hu, M.; Liu, S.; Wei, F.; Wu, Y.; Stasko, J.; and Ma, K.-L. 2012. Breaking news on Twitter. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, 2751–2754. New York, NY, USA: ACM.

Lasorsa, D. L.; Lewis, S. C.; and Holton, A. E. 2012. Normalizing Twitter. *Journalism Studies* 13(1):19–36.

Molyneux, L. 2014. What journalists retweet: Opinion, humor, and brand development on Twitter. *Journalism*.

Olteanu, A.; Castillo, C.; Diakopoulos, N.; and Aberer, K. 2015. Comparing events coverage in online news and social media: The case of climate change. In *Proceedings of the Ninth International Conference on Weblogs and Social Media*.

Parmelee, J. H. 2013. Political journalists and Twitter: Influences on norms and practices. *JMP* 14(4).

Phuvipadawat, S., and Murata, T. 2010. Breaking news detection and tracking in Twitter. In *Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on*, volume 3, 120–123.

Sheffer, M. L., and Schultz, B. 2010. Paradigm shift or passing fad? Twitter and sports journalism. *International Journal of Sport Communication* 3(4):472–484.

Verweij, P., and van Noort, E. 2014. Journalists' twitter networks, public debates and relationships in South Africa. *Digital Journalism* 2(1):98–114.

Vis, F. 2013. Twitter as a reporting tool for breaking news. *Digital Journalism* 1(1):27–47.

Welch, B. L. 1947. The generalization of 'Student's' problem when several different population variances are involved. *Biometrika* 34(1/2):pp. 28–35.

Zak, E. 2012. 4 questions with Liz Heron, the New York Times social media editor. `adweek.com/fishbowlny/-/250679`. Accessed: 2016-03-06.

Zhao, Z., and Mei, Q. 2013. Questions about questions: An empirical analysis of information needs on Twitter. In *Proceedings of the 22Nd International Conference on World Wide Web*, WWW '13, 1545–1556. New York, NY, USA: ACM.