Of Bots and Humans (on Twitter)

Zafar Gilani*, Reza Farahbakhsh†, Gareth Tyson‡, Liang Wang*, Jon Crowcroft*
*Computer Laboratory, University of Cambridge, Cambridge, UK, {szuhg2, lw525, jac22}@cam.ac.uk
†CNRS Lab UMR5157, Institut Mines Telecom, Paris, France, reza.farahbakhsh@it-sudparis.eu
‡Queen Mary University of London, London, UK, g.tyson@qmul.ac.uk

***Abstract*—Recent research has shown a substantial active presence of bots in online social networks (OSNs). In this paper we utilise our previous work (*Stweeler*) to comparatively analyse the usage and impact of bots and humans on Twitter, one of the largest OSNs in the world. We collect a large-scale Twitter dataset and define various metrics based on tweet metadata. Using a human annotation task we assign 'bot' and 'human' ground-truth labels to the dataset, and compare the annotations against an online bot detection tool for evaluation. We then ask a series of questions to discern important behavioural characteristics of bots and humans using metrics within and among four popularity groups. From the comparative analysis we draw differences and interesting similarities between the two entities, thus paving the way for reliable classification of bots, and studying automated political infiltration and advertisement campaigns.**

***Index Terms*—content propagation; social network analysis; bot characterisation; behavioural analysis**

I. Introduction

Bots (automated agents) exist in vast quantities in online social networks. They are created for a number of purposes, *e.g.* news, marketing, link farming,[1] political infiltration,[2] spamming and spreading malicious content.

The rise of bots on Twitter is evidenced by a number of studies [11], [15], [7], [5], and articles discussing bots.[3] This constitutes a radial shift in the nature of content production, which has traditionally been the realm of human creativity (or at least intervention). Although there have been past studies on bots (§II), we are particularly interested in exploring their role in the wider social ecosystem, and how their behavioural characteristics differ from humans. This is driven by many factors. The limited cognitive ability of bots clearly plays a major role, however, it is also driven by their diverse range of purposes, ranging from curating news to answering customer queries. This raises a number of interesting questions regarding how these bots operate, interact and affect online content production: What are the typical behaviours of humans and bots, in terms of their own activities as well as the reactions of others to them? What interactions between humans and bots occur? How do bots affect the overall social activities? The understanding of these questions can have deep implications in many fields such as social media analysis and systems engineering.

To answer these questions, we have performed a large-scale measurement and analysis campaign on Twitter (§III). We focus on bots in Twitter because it largely exposes public content, and past studies indicate a substantial presence of bots [4]. We offer a new and fundamental understanding of the characteristics of bots *vs.* humans, observing a number of clear differences (§IV). For example, we find that humans generate far more novel content, while bots rely more on retweeting. We also observe less intuitive trends, such as the propensity of bots to tweet more URLs, and upload bulkier media (*e.g.* images). We further analyse the social interconnectedness of bots and humans to characterise how they influence the wider Twittersphere. We observe that, although human contributions are generally considered more important via typical metrics (*e.g.* number of likes, retweets), bots still sustain significant influence over content production and propagation. As well as providing a powerful underpinning for future bot detection methods, our work makes contributions to the wider field of social content automation. Such understanding is critical for future studies of social media, which are often skewed by the presence of bots.

II. Related Work

Two main streams of research are relevant: (*i*) social, demographical and behavioural analyses of either bots or humans; and (*ii*) the impact of bots in social environments. Bot detection is not the focus of this paper, rather we characterise and compare the types of users for a broader understanding.

Social analysis of bots or humans. Most related to our work are behavioural studies of bots or humans. For example, [8] studied the *infiltration strategies* of social bots on Twitter using a manual approach. They use three metrics to quantify the infiltration of social bots: followers, popularity score, and message-based interaction (other users favouriting, retweeting, replying or mentioning the bot). They found that bots can successfully evade Twitter defences (only 38 out of their 120 bots got suspended over the course of 30 days).

Researchers have also inspected bot or human behaviour, though isolation. For example, [3] examined the retweet behaviour of people, focussing on *how people tweet*, as well

[1]Link farming – `http://bit.ly/2cXhfBv`
[2]Bots distort U.S. presidential election – `http://bit.ly/2l3VzGf`
[3]Bots in press and blogs – `http://bit.ly/2dBAIbB`

ASONAM '17, July 31 - August 03, 2017, Sydney, Australia

ACM ISBN 978-1-4503-4993-2/17/07. $15.00
http://dx.doi.org/10.1145/3110025.3110090

as *why and what people retweet*. The authors found that participants retweet using different styles, and for diverse reasons (*e.g.* for others or for social action). This is relevant to our own work, as we also study retweets. In contrast, our work provides further insights on important differences and striking similarities between bots and humans in terms of *retweet patterns*, *account lifetime*, *content creation*, *content popularity*, *entity interaction*, *content consumption*, *account reciprocity*, and *content propagation*. To the best of our knowledge, we are the first to perform this methodical comparison of representative metrics across these types of Twitter accounts.

Social influence of bots. In [2], authors use a bot on aNobii, a social networking site aimed at readers, to explore the *trust*, *popularity* and *influence* of bots. They show that gaining popularity does not require individualistic user features or actions, but rather simple social probing (*i.e.* bots following and sending messages to users randomly). The authors also found that an account can circumvent trust if it is popular (since popularity translates into influence). Closely related is [14], which develops models to identify users who are *susceptible* to social bots, *i.e.* likely to follow and interact with bots. The authors use a dataset from the Social Bot Challenge 2011, and make a number of interesting findings, *e.g.* that users who employ more negation words have a higher susceptibility level. Recent work [9] has also shown the impact of bots on Twitter activity using a non-infiltrating honeypot experiment.

In our work, we study the characteristics of existing bots in detail and argue that this provides far broader vantage into real bot activities. Hence, unlike studies that focus on the influence of individual bots (*e.g.* the Syrian Civil War [1]), we gain perspective on the wider spectrum of how bots and humans operate, and interact.

III. Methodology

We use and build upon our previous work ***Stweeler***[4] [10] for data collection, pre-processing, human annotation, and analysis. We define a 'bot' as any account that *consistently* involves automation over the observed period, *e.g.* use of the Twitter API or other third party tools, performing actions such as automated likes, tweets, retweets, *etc*. Note that a *tweet* is an original status and not a retweet, a *retweet* is a tweet which has 'RT' in text, and a *status* is either a tweet or a retweet. Also note that *content* on Twitter is limited to whatever is contained within a tweet: text, URL, image, and video.

A. Data Collection

We selected Twitter because it is open, large-scale and is known to contain a wide breath of bot activities. We collect data on bot and human behaviour for 30 days in April 2016 from the Streaming API. This resulted in approximately 65 million tweets, with approximately 2 to 2.5 million recorded per day. We then extracted the accounts and all associated metadata (*e.g.* account age) from tweets. In total, we recorded information on 2.9 million unique accounts. In this study, in addition to known metrics (age, tweets, retweets, favourites, replies and mentions, URL count, follower-friend ratio, *etc*), we also analyse a set of six novel metrics not explored in past bot research. These are: *likes per tweet*, *retweets per tweet*, *user replies and mentions*, *activity source count*, *type of activity sources*, and *size of content uploaded*. The selection of features is driven by [6] and, to our knowledge, this is the most comprehensive study to date.

B. Data Pre-Processing

Our data contains a range of accounts in terms of their popularity (*i.e.* number of followers). Hence, we partition profiles into four popularity groups to enable a deeper understanding. These are as follows:

$\mathbf{G_{10M+}}$**– celebrity status:** This is the subset of Twitter users with the highest number of followers, *i.e.* >9M followers. These are the most popular users, who hold celebrity status and are globally renowned. Popular and credible organisations (*e.g.* CNN, NetGeo) use these accounts for various purposes, which makes them free of spam, thus having high credibility and trustworthiness.

$\mathbf{G_{1M}}$**– very popular:** This subset of Twitter users is amongst the most popular on the platform, *i.e.* 900K to 1.1M followers. These users are close to celebrity status and global recognition (*e.g.* nytfood, pcgamer).

$\mathbf{G_{100k}}$**– mid-level recognition:** This subset represents popular accounts with mid-level recognition (*e.g.* CBSPhilly, DomusWeb), *i.e.* 90k to 110k followers.

$\mathbf{G_{1k}}$**– lower popularity:** This subset represents more ordinary users, *i.e.* 0.9k to 1.1k followers. These users (*e.g.* hope_bot, Taiwan_Agent) form a large base and, though they show lower individual and accumulated activity, they do form the all-important tail of the distribution.

Our dataset[5] is a representative sample of Twitter users, where each metric follows a Gaussian distribution. $\mathbf{G_{10M+}}$ and $\mathbf{G_{1M}}$ are similar in their characteristics (*cf.* §IV) and constitute 0.65% of the total 105k accounts we partitioned in the dataset. $\mathbf{G_{1k}}$ represents the bulk of Twitter, constituting 94.40% of the total partitioned accounts. $\mathbf{G_{100k}}$ bridges the gap between the most popular and least popular groups, constituting 4.93% of the total partitioned accounts. A possible $\mathbf{G_{10k}}$ would be similar to $\mathbf{G_{1k}}$, and a possible $\mathbf{G_{50k}}$ will be similar to $\mathbf{G_{100k}}$.

C. Bot Classification

To compare bots with humans, it is necessary to identify which accounts are operated by bots. We initially experimented with the updated release of BotOrNot [13], a state-of-the-art bot detection tool (to the best of our knowledge this is the only available online bot detection tool). However, inspection of the results indicated high inaccuracy with different thresholds (40% to 60%) to label an account as 'bot'. Hence, we chose to take a manual approach instead — we made this design choice to have a smaller but more

[4]***Stweeler***– https://github.com/zafargilani/stcs

[5]Datasets can be found here – https://goo.gl/SigsQB

TABLE I
SUMMARY OF TWITTER DATASET POST-ANNOTATION.

Group	#Bot accts	#Human accts	#Bot statuses	#Human statuses
$\mathbf{G_{10M+}}$	25	25	71303	79033
$\mathbf{G_{1M}}$	250	383	145568	157949
$\mathbf{G_{100k}}$	531	559	148015	82562
$\mathbf{G_{1k}}$	498	791	24328	13351
Total	1304	1758	389214	332895

reliable set of classifications. We employed human participants to perform a *human annotation task*[6] to identify bots and humans. We recruited four undergraduate students for the purposes of annotation. Each account was reviewed by all recruits independently, before being aggregated into a final judgement using a final collective review (via discussion if needed).

As well as providing the recruits with the Twitter profile, we also presented summary data to streamline the task. We further provide participants with a list of the 'sources' used by the account over the month, *e.g.* Twitter app, browser, *etc*. The human workers consider both the number of sources used, and the types of sources used. This is because sources can reveal traces of automation, *e.g.* use of the Twitter API. Additionally, the human worker would also visit a user's Twitter page and verify the content and URLs posted. Overall, we presented participants with randomised lists that fell into the four popularity groups containing roughly 25k accounts each. Human annotators were instructed to filter out any account that matched the following criteria: an account that does not exhibit activity (*i.e.* no tweet, retweet, reply, and mention), or an account that is suspended. In total, the volunteers successfully annotated 3062 active accounts: 1304 were classified as bots and 1758 as humans. Table I provides a summary of the data.

For context, we can cross validate by comparing the agreement of final annotations by the human workers to the BotOrNot annotation. The average inter-annotator agreement compares the pairs of labels by each human annotator to capture the percentage of accounts for which all four annotators unanimously agree. The average agreement is measured as a percentage of agreement: 0% shows lack of agreement and 100% shows perfect agreement. Our human annotation task shows very high unanimous agreement between human annotators for each popularity group: $\mathbf{G_{10M+}}$ (96.00%), $\mathbf{G_{1M}}$ (86.32%), $\mathbf{G_{100k}}$ (80.66%), and $\mathbf{G_{1k}}$ (93.35%). Whereas, BotOrNot shows lower than average agreement with the final labels assigned by the human annotators: $\mathbf{G_{10M+}}$ (46.00%), $\mathbf{G_{1M}}$ (58.58%), $\mathbf{G_{100k}}$ (42.98%), and $\mathbf{G_{1k}}$ (44.00%). Since, BotOrNot yields a lower accuracy, we restricted ourselves to the dataset of accounts that were manually annotated.

IV. WHICH MANNERS MAKETH THE BOT?

The purpose of this study is to discover the key account characteristics that are typical (or atypical) of bots and humans. Recall that we take a broad perspective on what a 'bot' is, *i.e.* any account that *consistently* involves automation over the observed period, but may involve human intervention. This

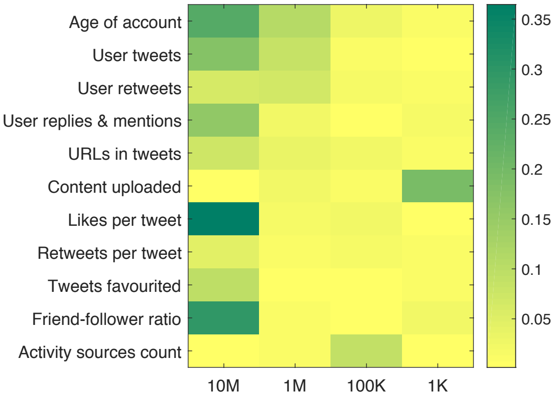

Fig. 1. Spearman's rank correlation coefficient (ρ) between bots and humans per measured metric. The figure shows none (0.0) to weak correlation (0.35) across all metrics, indicating clear distinction between the two entities.

definition is justified by the purpose of automation, *i.e.* humans act as *bot managers*, whereas bots are *workers*. To explore this, we use our data (§III) to empirically characterise bots (dashed lines in figures) and humans (solid lines in figures). To begin, we simply compute the correlation between each feature for bots and humans; Figure 1 presents the results as a heatmap (where perfect correlation is 1.0). Notice that most features exhibit very poor correlations (0.0 to 0.35), indicating significant discrepancies between bot and human behaviour — we therefore spend the remainder of this paper exploring these differences in depth.

A. Content Generation

We begin by asking *if bots generate more content on Twitter than humans?* We initially consider two forms of content creation: a *tweet*, which is an original status written by the account, and a *retweet*, which is repetition of an existing status. When using the term *status*, we are referring to the sum of both tweets and retweets. First, we inspect the amount of content shared by computing the number of statuses (*i.e.* tweets + retweets) generated by each account across the 30 days. As anticipated, humans post statuses less frequently than bots (monthly average of 192 for humans *vs.* 303 for bots), in all popularity groups except $\mathbf{G_{10M+}}$, where surprisingly humans post slightly more than bots. The sheer bulk of statuses generated by $\mathbf{G_{10M+}}$ (on average 2852 for bots, 3161 for humans in a month) is likely to acquire popularity and new followers. Overall, bots constitute 51.85% of all statuses in our dataset, even though they are only 43.14% of the accounts.

An obvious follow-up is *what do accounts tweet?* This is particularly pertinent as bots are often reputed to lack original content. To explore this, we inspect the number of *tweets vs. retweets* performed by each account. Figures 2(a) and 2(b) present the empirical distributions of tweets and retweets, respectively, over the 30 days. We see that the retweet distribution is rather different to tweets. Bots in $\mathbf{G_{1M}}$, $\mathbf{G_{100k}}$ and $\mathbf{G_{1k}}$ are far more aggressive in their retweeting; on average, bots generate 2.20× more retweets than humans. The only exception to this trend is $\mathbf{G_{10M+}}$ where humans retweet

[6] Human annotation task details – `http://bit.ly/2cH0YvA`

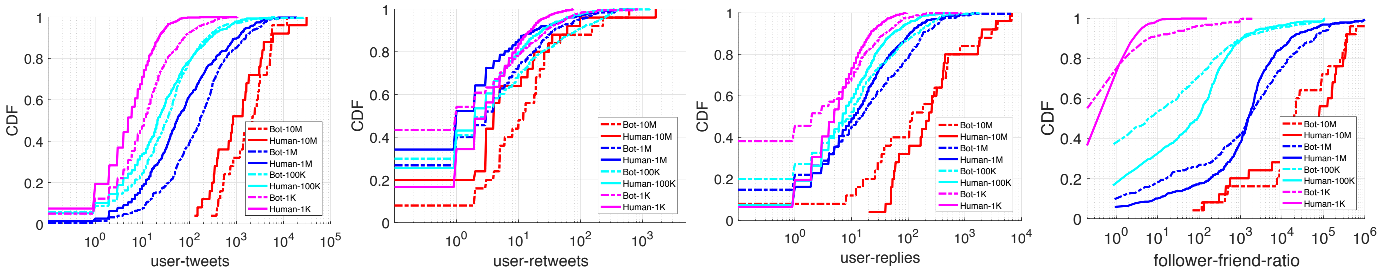

(a) Tweets issued by a user.

(b) Retweets issued by a user.

(c) Replies and mentions posted by a user.

(d) Follower-friend ratio of a user.

Fig. 2. Content Creation: Tweets issued, Retweets issued, Replies and mentions; Account Reciprocity: Follower-friend ratio.

1.54× more often than bots. This is likely driven by the large number of tweets generated by celebrity users. Typically, humans do generate *new* tweets more often, while bots rely more heavily on retweeting existing content. Generally, humans post 18 tweets for every retweet, whereas bots post 13 tweets for every retweet in all popularity groups except $\mathbf{G_{10M+}}$ (where both entities show similar trends).

Whereas tweets and retweets do not require one-to-one interaction, a further type of messaging on Twitter, via *replies*, does require one-to-one interaction. These are tweets that are created in response to a prior tweet (using the @ notation). Figure 2(c) presents the distribution of the number of replies issued by each account. We anticipate that bots post more replies and mentions given their automated capacity to do so. However, in $\mathbf{G_{10M+}}$ both bots and humans post a high number of replies, and bots post only marginally more than celebrities. While bot-masters in $\mathbf{G_{10M+}}$ deploy *chatbots* to address simple user queries, celebrities reply in order to engage with their fanbase. It also possible that celebrities employ managers as well as automation and scheduling tools (§IV-E) for such a purpose. Bots in the remaining popularity groups respond twice as frequently as their human counterparts. Again, this is driven by the ease by which bots can automatically generate replies: only the most dedicated human users can compete.

Finally, we briefly inspect the actual content of the tweets being generated by the accounts. We do this using two metrics: number of URLs posted by accounts, and the size of media (*e.g.* photos) uploaded. Figure 3(a) presents the scatter plot of the number of URLs (y-axis) and content uploaded in KB (x-axis). Bots place far more external URLs in their tweets than humans (see Table II): 162% in $\mathbf{G_{10M+}}$, 206% more in $\mathbf{G_{1M}}$, 333% more in $\mathbf{G_{100k}}$, and 485% more in $\mathbf{G_{1k}}$. Bots are a clear driving force for generating traffic to third party sites, and upload far more content on Twitter than humans. Figure 3(b) presents the distribution of the amount of content uploaded by accounts (*e.g.* photos). Account popularity has a major impact on this metric. Bots in $\mathbf{G_{10M+}}$ have a 102× lead over bots in other popularity groups. That said, humans in $\mathbf{G_{10M+}}$ have a 366× lead over humans in other popularity groups. Overall, bots upload substantially more bytes than humans do (see Table II): 141% in $\mathbf{G_{10M+}}$, 975% more in $\mathbf{G_{1M}}$, 376% more in $\mathbf{G_{100k}}$, and 328% more in $\mathbf{G_{1k}}$.

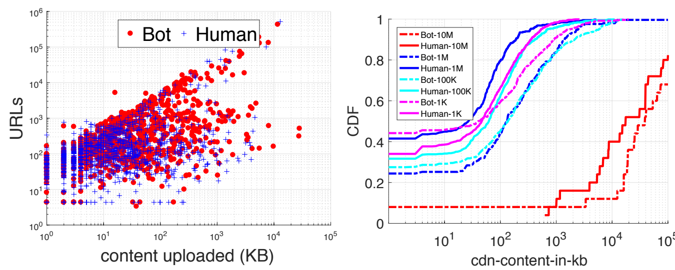

(a) URLs *vs.* Content uploaded.

(b) Size of CDN content in KByte uploaded by a user.

Fig. 3. Content Creation: URLs in tweets, Content uploaded on Twitter.

B. Content Popularity

The previous section has explored the amount of content generated by accounts, however, this does not preclude such content from being of a low quality. To investigate this, we compute standard popularity metrics for each user group.

First, we inspect the *number of favourites* or *likes* received for tweets generated by the accounts. This is a reasonable proxy for tweet quality. Figure 4(a) presents the empirical distribution of the number of favourites or likes received for all the tweets generated by the profiles in each group. A significant discrepancy can be observed. Humans receive *far* more favourites per tweet than bots across all popularity groups except $\mathbf{G_{1k}}$. Close inspection revealed that bots in $\mathbf{G_{1k}}$ are typically part of larger *social botnets* that try to promote each other systematically for purposes as outlined in §I. In contrast, human accounts are limited to their social peers and do not usually indulge in the 'influence' race. For $\mathbf{G_{10M+}}$, $\mathbf{G_{1M}}$ and $\mathbf{G_{100k}}$ popularity groups, humans receive an average of 27×, 3× and 2× more favourites per tweet than bots, respectively. $\mathbf{G_{1k}}$ bots are an exception that receive 1.5× more favourites per tweet than humans. These findings suggest that: (*i*) the term *popularity* may not be ideally defined by the number of followers, (*ii*) human content gathers greater engagement due to its personalised attributes.

A further *stronger* sign of content quality is another user retweeting content. Humans consistently receive more retweets for all popularity groups $\mathbf{G_{10M+}}$: 24-to-1, $\mathbf{G_{1M}}$ and $\mathbf{G_{100k}}$: 2-to-1, except $\mathbf{G_{1k}}$: 1-to-1. This difference, shown in Figure 4(b), is indicative of the fanbase loyalty, which is vastly higher for individual celebrities than reputable organisations. In other

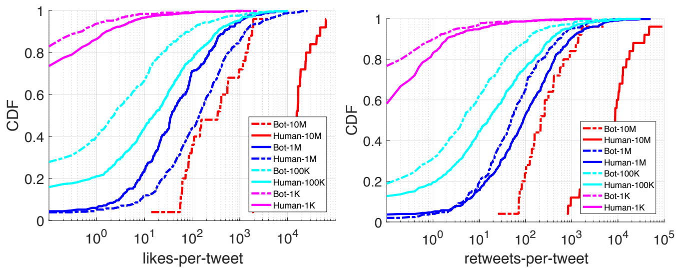

(a) Likes per tweet received by a user. (b) Retweets per tweet received by a user.

Fig. 4. Content Popularity: Likes per tweet, Retweets per tweet.

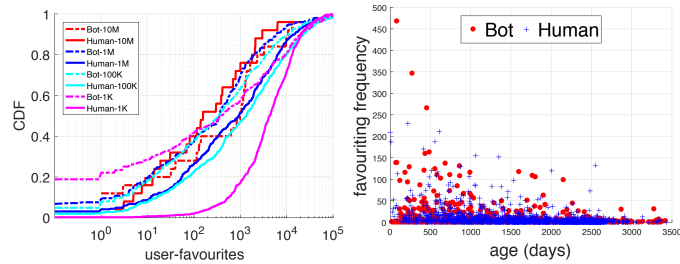

(a) Tweets favourited (liked) by a user. (b) Number of favourites performed *vs.* age of the account.

Fig. 5. Content Consumption: Likes performed, Favouriting behaviour.

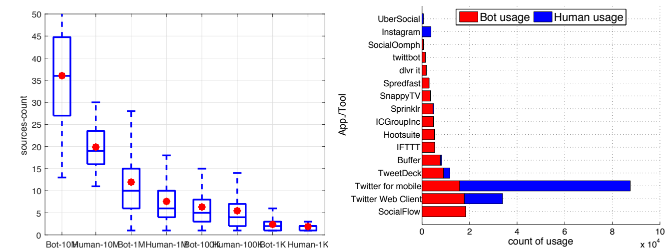

(a) Activity sources used by a user (Red dot is μ of the distribution). (b) Bar chart of accounts that use each type of Twitter source.

Fig. 6. Tweet Sources: Count of Activity Sources, Type of Activity Sources.

words, the *quality* of human content appears to be much higher. We then inspect *who* performs the retweets, *i.e.* do bots tend to retweet other bots or humans? We find that bots retweeting bots is over 3× greater than bots retweeting humans. Similarly, humans retweeting humans is over 2× greater than humans retweeting bots. Overall, bots are retweeted 1.5× more often than humans. This indicates a form of homophily and assortativity.

C. Content Consumption

Whereas the previous metrics have been based on content produced *by* the accounts under study, our dataset also includes the consumption preferences of the accounts themselves. Hence, we ask *how often do bots 'favourite' content from other users and how do they compare to humans?* Intuitively, bots would be able to perform far more likes than humans (who are physically constrained). Figure 5(a) shows the empirical distribution of the number of likes performed by each account. It can be seen that, actually, for most popularity groups ($\mathbf{G_{1M}}$, $\mathbf{G_{100k}}$, $\mathbf{G_{1k}}$), humans favourite tweets more often than bots (on average 8251 for humans *vs.* 5445 for bots across the entire account lifetimes). Linking into the previous discussion, it therefore seems that bots rely more heavily on retweeting to interact with content. In some cases, the difference is significant; *e.g.* humans in $\mathbf{G_{1M}}$ and $\mathbf{G_{100k}}$ place twice as many likes as bots do. $\mathbf{G_{10M+}}$, however, has an average of 1816 by humans compared to 2921 by bots.

We conjecture that there are several reasons for this trend: (*i*) humans appreciate what they like, (*ii*) bots are workers for their human managers and serve a purpose (*e.g.* promotion via tweets), (*iii*) humans have an incentive to like other tweets, potentially as a social practice (with friends) or in the hope of receiving likes in return [12]. To explore these strategies further, Figure 5(b) plots the number of favourites performed by an account *vs.* the age of the account. Firstly, bots are as old as humans: the oldest bot account is 3437 days old *vs.* 3429 days for the oldest human account. Secondly and more importantly, it can be seen that more recent (*i.e.* more modern) bots are significantly more aggressive in liking other tweets. Older bots, instead, use this feature less frequently; deeper inspection suggests this is driven by the trustworthy nature of older bots, which are largely run by major organisations.

D. Account Reciprocity

As well as content popularity, we can also measure reciprocity (*i.e.* friendship). Twitter classifies two kinds of relationships: reciprocal follower-relationship *i.e.* when two accounts follow each other, and non-reciprocal relationship *i.e.* an account has many followers who are not followed in return (this is often the case for celebrities). We measure this via the *Follower-Friend Ratio*. Figure 2(d) shows empirical distribution of the *Follower-Friend Ratio* for each group of accounts. Humans display higher levels of friendship ($\mathbf{G_{10M+}}$: 4.4×, $\mathbf{G_{1M}}$ and $\mathbf{G_{100k}}$: 1.33×, $\mathbf{G_{1k}}$: 15×) and thus a lower *Follower-Friend Ratio* than bots.

Previous research [4] argues that humans typically have a ratio close to 1, however, our analysis contradicts this assumption. For celebrities, very popular and mid-level recognition accounts this ratio is in the order of thousands-to-1, irrespective of whether an account is a bot or a human ($\mathbf{G_{10M+}}$: 629011-to-1 for bots *vs.* 144612-to-1 for humans, $\mathbf{G_{1M}}$: 33062-to-1 for bots *vs.* 24623-to-1 for humans, $\mathbf{G_{100k}}$: 2906-to-1 for bots *vs.* 2328-to-1 for humans). In fact, even the ratios for low popularity accounts are not 1, but consistently greater ($\mathbf{G_{1k}}$: 30-to-1 for bots *vs.* 2-to-1 for humans). This is caused by the human propensity to follow celebrity accounts (who may not follow in return), as well as the propensity of bots to indiscriminately follow large numbers of other accounts (largely in the hope of being followed in return).

E. Tweet Generation Sources

Finally, we inspect the tools used by bots and humans to interact with Twitter. This is possible because each tweet is tagged with the *source* that generated it; this might be the

TABLE II
FEATURE INCLINATION: $\mathcal{B}$ IS MORE INDICATIVE OF BOTS, WHEREAS $\mathcal{H}$ IS MORE INDICATIVE OF HUMAN BEHAVIOUR, AND ◯ IS NEUTRAL (*i.e.* BOTH EXHIBIT SIMILAR BEHAVIOUR). * REPRESENTS MAGNITUDE OF INCLINATION: * IS CONSIDERABLE DIFFERENCE, ** IS LARGE DIFFERENCE. *signif.* SHOWS STATISTICAL SIGNIFICANCE OF EACH FEATURE AS MEASURED BY *t-test*.

Feature & value	Fig.	10M+	1M	100K	1K	*signif.*
More user tweets	2(a)	◯	$\mathcal{B}$*	$\mathcal{B}$*	$\mathcal{B}$*	
Higher user retweets	2(b)	$\mathcal{H}$*	$\mathcal{B}$*	$\mathcal{B}$*	$\mathcal{B}$*	99%
More user replies and mentions	2(c)	◯	$\mathcal{B}$*	$\mathcal{B}$*	$\mathcal{B}$	99%
More URLs in tweets	3(a)	$\mathcal{B}$**	$\mathcal{B}$**	$\mathcal{B}$**	$\mathcal{B}$**	99%
More CDN content uploaded(KByte)	3(b)	$\mathcal{B}$**	$\mathcal{B}$**	$\mathcal{B}$**	$\mathcal{B}$**	95%
Higher likes received per tweet	4(a)	$\mathcal{H}$**	$\mathcal{H}$**	$\mathcal{H}$**	$\mathcal{B}$	99%
Higher retweets received per tweet	4(b)	$\mathcal{H}$**	$\mathcal{H}$**	$\mathcal{H}$**	$\mathcal{B}$	99%
More tweets favourited (liked)	5(a)	$\mathcal{B}$**	$\mathcal{H}$**	$\mathcal{H}$**	$\mathcal{H}$**	99%
More favourites by *younger* accounts	5(b)	$\mathcal{B}$	$\mathcal{H}$	$\mathcal{B}$	$\mathcal{B}$	
Higher follower-friend ratio	2(d)	$\mathcal{B}$**	$\mathcal{B}$*	$\mathcal{B}$*	$\mathcal{B}$**	
More activity sources	6(a)	$\mathcal{B}$*	$\mathcal{B}$	$\mathcal{B}$	$\mathcal{B}$	99%

website, third-party app or tools that employ the Twitter API. Figure 6(a) presents the number of sources used by human and bot accounts of varying popularities. Bots are expected to use a single source (*i.e.* an API or own tool) for tweeting. Surprisingly, it can be seen that bots actually inject tweets using more sources than humans (*cf.* Table II).

To explore this further, Figure 6(b) presents the number of accounts that use each source observed. It can be seen that bots use a multitude of third-party tools. Bot news services (especially from $\mathbf{G_{10M+}}$) are found to be the heaviest users of social media automation management and scheduling services *(SocialFlow, Hootsuite, Sprinklr, Spredfast)*, as well as a Cloud-based service that helps live video editing and sharing *(SnappyTV)*. Some simpler bots (from $\mathbf{G_{100k}}$ and $\mathbf{G_{1k}}$ groups) use basic automation services *(Dlvr.it, Twittbot)*, as well as services that post tweets by detecting activity on other platforms *(IFTTT)*. A social media dashboard management tool seems to be popular across most groups except $\mathbf{G_{1k}}$*(TweetDeck)*. Interestingly, it can also be seen that bot accounts regularly tweet using Web/mobile clients — pointing to the possibility of a *mix* of automated and human operation. In contrast, 91.77% of humans rely exclusively on the Web/mobile clients. That said, a small number (3.67%) also use a popular social media dashboard management tool *(TweetDeck)*, and automated scheduling services *(Buffer, Sprinklr)*. This is particularly the case for celebrities, who likely use the tools to maintain high activity and follower interaction — this helps explain the capacity of celebrities to so regularly reply to fans (§IV-A).

V. CONCLUSIONS & FUTURE WORK

Bots exercise a profound impact on Twitter. Our work confirms a number of noteworthy trends: (*i*) bots generally retweet more often, while some humans can exhibit bot-like activity ($\mathbf{G_{10M+}}$); (*ii*) bots can post up to 5× more URLs in their tweets (§IV-A); (*iii*) bots can upload 10× more content with their tweets; (*iv*) humans can receive as much as 27× more likes and 24× more retweets as bots (§IV-B); (*v*) bots retweeting other bots is over 3× more regular than bots retweeting humans, whereas humans retweeting other humans is over 2× greater, indicating homophily (§IV-B); (*vi*) humans favourite others' tweets much more often than bots do, though *newer* bots are far more aggressive in favouriting tweets to replicate human behaviour (§IV-C); (*vii*) humans enjoy higher levels of friendship and usually form reciprocal relationships (§IV-D); (*viii*) bots typically use many different sources for active participation on Twitter (up to 50 or more); and (*ix*) activity sources include basic automation and scheduling services (§IV-E) — used abundantly by bots and seldomly by humans.

We conjecture that these bot activities may lead to dramatic changes in social structures and interactions in the longterm (as the bot population increases). Thus, there is a wide array of problems to explore in future, such as: exploring influence botnets, analysing bot content, and developing accurate detection tools.

Acknowledgements: This work is partially funded by EU Metrics project (Grant EC607728) and EU H2020 ReThink.

REFERENCES

[1] N. Abokhodair, D. Yoo, and D. W. McDonald. Dissecting a social botnet: Growth, content and influence in twitter. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 839–851. ACM, 2015.

[2] L. M. Aiello, M. Deplano, R. Schifanella, and G. Ruffo. People are strange when you're a stranger: Impact and influence of bots on social networks. In *Proceedings of the 6th ICWSM*, 2012.

[3] D. Boyd, S. Golder, and G. Lotan. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In *System Sciences (HICSS), 2010 43rd Hawaii International Conference on*, pages 1–10, Jan 2010.

[4] Z. Chu, S. Gianvecchio, H. Wang, and S. Jajodia. Who is tweeting on twitter: Human, bot, or cyborg? In *ACSAC '10*. ACM, 2010.

[5] S. Cresci, R. D. Pietro, M. Petrocchi, A. Spognardi, and M. Tesconi. The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race. *CoRR*, abs/1701.03017, 2017.

[6] N. Dugué, A. Perez, M. Danisch, F. Bridoux, A. Daviau, T. Kolubako, S. Munier, and H. Durbano. A reliable and evolutive web application to detect social capitalists. In *Proceedings of the IEEE/ACM ASONAM*, pages 741–744. ACM, 2015.

[7] C. Edwards, A. Edwards, P. R. Spence, and A. K. Shelton. Is that a bot running the social media feed? testing the differences in perceptions of communication quality for a human agent and a bot agent on twitter. *Computers in Human Behavior*, 33:372 – 376, 2014.

[8] C. Freitas, F. Benevenuto, S. Ghosh, and A. Veloso. Reverse engineering socialbot infiltration strategies in twitter. In *Proceedings of the IEEE/ACM ASONAM*, pages 25–32, 2015.

[9] Z. Gilani, R. Farahbakhsh, and J. Crowcroft. Do bots impact twitter activity? In *Proceedings of the 26th International Conference on World Wide Web Companion*, WWW '17 Companion, pages 781–782, 2017.

[10] Z. Gilani, L. Wang, J. Crowcroft, M. Almeida, and R. Farahbakhsh. Stweeler: A framework for twitter bot analysis. In *Proceedings of the 25th International Conference Companion on World Wide Web*, WWW '16 Companion, pages 37–38, 2016.

[11] K. Lee, B. D. Eoff, and J. Caverlee. Seven months with the devils: A long-term study of content polluters on twitter. In *ICWSM*, 2011.

[12] L. Scissors, M. Burke, and S. Wengrovitz. What's in a like?: Attitudes and behaviors around receiving likes on facebook. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, CSCW '16, pages 1501–1510. ACM, 2016.

[13] O. Varol, E. Ferrara, C. A. Davis, F. Menczer, and A. Flammini. Online human-bot interactions: Detection, estimation, and characterization. *arXiv preprint arXiv:1703.03107*, 2017.

[14] C. Wagner, S. Mitter, C. Körner, and M. Strohmaier. When social bots attack: Modeling susceptibility of users in online social networks. *Making Sense of Microposts (# MSM2012)*, 2, 2012.

[15] X. Wu, Z. Feng, W. Fan, J. Gao, and Y. Yu. *Detecting Marionette Microblog Users for Improved Information Credibility*, pages 483–498. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.