Organizations are Users Too: Characterizing and Detecting the Presence of Organizations on Twitter

James McCorriston, David Jurgens, and Derek Ruths
james.mccorriston@mail.mcgill.ca, jurgens@cs.mcgill.ca, druths@networkdynamics.org
School of Computer Science
McGill University

Abstract

Much work on the demographics of social media platforms such as Twitter has focused on the properties of individuals, such as gender or age. However, because credible detectors for organization accounts do not exist, these and future large-scale studies of human behavior on social media can be contaminated by the presence of accounts belonging to organizations. We analyze organizations on Twitter to assess their distinct behavioral characteristics and determine what types of organizations are active. We first create a dataset of manually classified accounts from a representative sample of Twitter and then introduce a classifier to distinguish between organizational and personal accounts. In addition, we find that although organizations make up less than 10% of the accounts, they are significantly more connected, with an order of magnitude more friends and followers.

1 Introduction

The presence and activity of organizations on social media platforms confound the grand goal of using such platforms to learn about human behavior. Beyond corporate advertising and customer engagement initiatives, organizations like political parties, social groups, and local clubs also use these platforms for communication and coordination (Golbeck, Grimes, and Rogers 2010). The inability to distinguish between organizational and personal accounts can have significant ramifications for applications such as election predictions (Tumasjan et al. 2010), health monitoring (Schwartz et al. 2013), and crisis response (Saleem, Xu, and Ruths 2014) — all of which need to isolate signals specifically from individuals. Despite known organization activity on social media and its affects on large-scale measurements of human behavior, little is known about the scale of this presence.

A main cause for treating all accounts identically is the lack of a clear methodology for distinguishing between the two types. Only recently have works proposed classifying Twitter accounts as either personal or organizational (De Choudhury, Diakopoulos, and Naaman 2012; De Silva and Riloff 2014; Yin et al. 2014). However, prior approaches have taken overly-narrow definitions of organizational and personal accounts and sampled accounts of both types in biased manners, leading to unrealistic estimates of the presence of organizations on Twitter and the inability to automatically distinguish between both types accurately.

We propose a comprehensive study of organization demographics and behavior on Twitter. Using over 20,000 active, predominantly English-language Twitter accounts we assess the division, behavior, and social connectivity of organization and personal accounts, We show that organizations make up 9.4% of accounts on Twitter, and find evidence of organization behaviors that contradict findings of prior work using fewer Twitter accounts. From this analysis, we designed a novel classifier that is able to accurately distinguish between personal and organizational accounts, obtaining an F1 score of 95.5. The dataset and classifier are made publicly available.

2 Related Work

Distinguishing account types can be viewed as a type of latent attribute inference, which aims to infer various properties of online accounts. While only recently has latent attribute inference work begun to examine the organization-person distinction, much work has been done on other specific aspects such as political affiliation (Cohen and Ruths 2013), gender (Ciot, Sonderegger, and Ruths 2013; Alowibdi, Buy, and Yu 2013), age (Nguyen, Smith, and Rosé 2011; Nguyen et al. 2013), location (Jurgens 2013), or combinations thereof (Zamal, Liu, and Ruths 2012; Li, Ritter, and Hovy 2014). Our work is complementary and may offer an important benefit by removing noise from organizational accounts which do not have human attributes.

Three related studies have examined organizational presence in Twitter. While valuable exploratory works, these studies suffer from biased data collection and dataset size issues that restrict their findings to specialized kinds of organizational or personal accounts. First, De Silva and Riloff (2014) address the problem of detecting whether a tweet comes from an organization or personal account. Their dataset consists of tweets from 58 organizations, gathered from Twellow, and 600 personal accounts, identified by matching the account's profile description with a list of person names. Given the diversity of organization types and bias inherent to Twellow (Cohen and Ruths 2013), 58 organizational accounts offer a very sparse (and quite likely biased) sampling. Second, Yin et al. (2014) performed an analysis using a sample of 5000 accounts that posted at

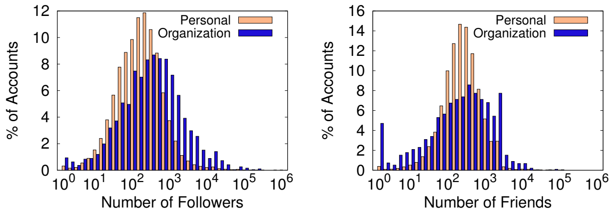

Figure 1: Degree distributions for an account's followers (left) and friends (right), with a log-scale x-axis.

least 200 geotagged tweets. As will be discussed, we determined that the majority of organizations do not produce geotagged tweets - thus this sample is biased towards a particular kind of organization. Moreover, in their manual analysis they could not recognize the type of 36% (1800 accounts) of the accounts they collected, a problem not seen in our annotation process. Third, De Choudhury, Diakopoulos, and Naaman (2012) classify accounts as organizational, journalist, or individual in order to assess who participates in online discussions. Their dataset was constructed by combining 4932 accounts from a sample of the public Twitter timeline (which overrepresents the most active accounts), organizations listed on Twellow (with its aforementioned bias), and journalists listed on the *Muckrack* directory.

3 Characterizing Organizational Presence

As no standard dataset exists, we first constructed a high-quality, unbiased dataset of classified Twitter accounts.

3.1 Data collection

To obtain an accurate representation of account type distributions, Twitter accounts were initially selected by uniformly sampling account numbers from within [10K, 1B]. The selection procedure avoids the frequency bias of sampling accounts from the Twitter gardenhose and the bias from sampling by walking the Twitter social network, which may not capture active accounts that do not participate in social behaviors. The upper limit of this account range was selected to limit the analysis to Twitter accounts created before November 2012, which ensures that the accounts have sufficient activity to identify their type. A follow-up replication of accounts in [1B, 3B] showed no significant difference in the account type distribution. The present study focuses on Twitter accounts that are actively communicating. Therefore, accounts were restricted to those posting at least 100 tweets. To focus on accounts that we can both manually and automatically analyze, the dataset was limited to those containing a non-empty English-language profile description. A total of 34,000 accounts were collected.

All accounts were annotated using Amazon Mechanical Turk (MTurk). Each MTurk task included fifteen questions that asked MTurk workers to classify the account as belonging to a person or organization using the following two definitions to differentiate between them: (1) a personal account is one controlled by an individual, and (2) an organizational account is one controlled by a group or an organization representing more than one person. Questions displayed

Label	# Accts	% Accts	% Volume	Est. Rate
Organization	1,911	9.4	5.4	0.58
Personal	18,362	90.6	94.6	0.43

Table 1: Organizational and personal accounts statistic. The last two columns report the % of total Twitter volume and hourly tweet rate, estimated from the 100 most-recent tweets per account.

an account's profile name, self-reported description, profile picture, and six recent tweets, which were filtered to omit retweets. Workers were paid $0.08USD per task.

To control for quality, two of the fifteen questions had known answers. All responses for the task were rejected if a worker answered either incorrectly, which led to the rejection of 34.3% of all responses. After removing these responses, only accounts receiving the same classification from all three workers were included. However, the vast majority of accounts had a unanimous labeling (90.7%) and workers achieved high inter-annotator agreement, having Cohen's κ=0.95. Ultimately, 20,273 accounts were labeled.

3.2 Account Presence and Activity

The presence of personal and organizational accounts was measured using the number of accounts and tweet volume.

Account composition Our analysis shown in Column 2 of Table 1 indicates that Twitter overwhelmingly consists of personal accounts, but that an active minority of organizational accounts exist. This percentage is larger than the estimate of 7.8% by Yin et al. (2014), who considered only accounts producing geotagged tweets; we attribute this difference to the inclusion of accounts that post from desktop devices and do not report their geolocation.

Tweet volume composition Given the differences in the composition of account types, we assessed what percentage of content overall was generated by each type. Total volume was measured by summing the number of lifetime posts of each account in the dataset and measuring their relative percentages of the total. As shown in Column 3 of Table 1, the vast majority of Twitter content (>94%) originates from personal accounts. Organizations generate substantially less volume than their percentage makeup of accounts would suggest (9.4% vs. 5.4%). However, examining the estimated frequency with which accounts post (Col. 4), organizational accounts are more likely to post messages more frequently than personal accounts.[1] Furthermore, because organizations also have more followers, as a result, their content receives disproportionately more visibility.

3.3 Account Characteristics

Social Behavior The Twitter social network is constructed from directional edges where account a_i can form a relationship with another account a_j; then, a_i is said to have a_j as

[1] Given the increasing attrition rate of Twitter users over time (Liu, Kliman-Silver, and Mislove 2014), we speculate that disinterested individuals slow their activity as their account ages before finally leaving the site, leading to a lower average tweet rate; in contrast, organizational accounts may be more consistent in their posting in order to maintain reputation.

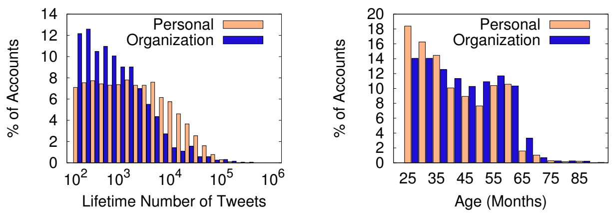

Figure 2: Distributions of tweet volume (left) and age (right).

a friend and a_j has account a_i as a follower. As shown in Figure 2 organizational accounts have both more followers and more friends, on average, than personal accounts. Note that the x-axis is in log-scale, thus the seemingly small shift represents nearly an order of magnitude difference. The significantly higher ratio of followers to friends for organizations highlights the impact of organizations: organization-produced content is on average seen by an audience nearly three times larger than that of a personal account.

However, we note that a sizable minority of organizational accounts have no friends, shown as the spike on the left of Figure 1b. A manual inspection of a sample of such accounts showed that these organizations are often followed by others and do actively tweet, in essence serving only to disseminate information to interested parties but not engaging in the social aspects of Twitter.

Tweet Volume Given the under-representation of organizations in the total volume of tweets, we examined how individual accounts of each type differ in the amount of content they produce. Figure 2 shows the distribution of tweet volume, with a logarithmic scale x-axis, revealing that personal accounts are far more likely to generate a higher number of tweets over their lifetime and conversely, the majority of organizational accounts generate few tweets.

Age Distribution Twitter experienced tremendous growth in its userbase and therefore we assess whether organizations and individuals joined at the same rate. The age distribution of both personal and organizational accounts, shown in Figure 2 is strikingly bimodal. Organizations have much more mass in this older mode (statistically significant at p $< 10^{-11}$), making their accounts older on average. One interpretation of this is that organizations have more staying power on Twitter, meaning that they do not abandon the platform as often as individuals. Individuals, on the other hand, experience more churn, either abandoning microblogging or migrating to other platforms over time. This is consistent with the notion that organizations are making a more calculated investment when setting up an account on the platform.

Geolocalized Activity Some Twitter posts are geotagged with the location where the post was generated. Yin et al. (2014) suggest that organizations tweet from a wider range of locations than individuals. However, we assessed the percentage of accounts creating geotagged tweets and the number of unique location names recorded for each type of account and found that organizations were far less likely to tweet from multiple locations than individuals. Further, the majority of accounts (78.8% of individuals, 87.9% of organizations) never produced a geotagged tweet, suggesting that the method by which Yin et al. (2014) selected organization accounts yielded non-representative organization behavior.

4 Classifying account types

Given workers's high agreement in distinguishing account type, we propose an automated method for the same task.

4.1 Experimental Setup

Because the task is a binary classification (personal or organizational), a support vector machine was used (Chang and Lin 2011). Borrowing from past studies on demographic inference, three types of features were used for distinguishing between account types: (1) post content features, (2) stylistic features, how the information is presented, and (3) structural and behavioral features based on how the account interacts with others. Four types of content features were chosen. To capture the language biases of each class, all open-class words were ranked according to their relative frequency between the two classes, selecting the k the most class-biased words. A similar procedure was used to select the k most-biased stemmed tokens and hashtags for each class. Finally, we include the k most-biased co-stems, which proved useful for inferring author attributes (Lipka and Stein 2011). The same value of k=13 was used for all four feature types after initial tuning showed it provided good performance; higher values of k did not noticeably improve performance.

Seven stylistic features were selected. Account vocabulary differences were measured by (a) the average word length and average number of words used per tweet and (b) the frequencies with which the tweet includes a hashtag or hyperlink. Furthermore, we hypothesize that organizations may be more constrained in the content they write about and therefore will have fewer unique words in their lexicon; therefore, we include the rate at which an account uses new words and new hashtags, measured as the number of unique terms (or hashtags) divided by the total number of tokens produced. Finally, we include the profile description length.

Five structural and behavioral features were incorporated using (a) the number of tweets and number of retweets, (b) percentages of tweets that include a mention of another user and those that are a retweet, (c) ratio of retweets to tweets, and (d) the ratio of the number of followers to friends.

Feature vectors were created using 200 tweets for each account and normalized such that value v_i for feature f_i was transformed to $\log(v_i - min_i + 1)$ where min_i is the minimum value observed for f_i. A LibSVM classifier was then trained using a radial basis function kernel. Classifier performance was measured using five-fold cross-validation. Performance is compared against the majority-class baseline, which labels all accounts as personal accounts. Given the skewed distribution of account types, two evaluations were performed, using datasets either (1) uniformly distributed between account types or (2) matched the natural distribution of types. Datasets for both evaluations were constructed to be the same size in order to make the results comparable.

4.2 Results

Both evaluation conditions produced highly accurate classifiers, achieving 88.8% accuracy in the uniformly-distributed

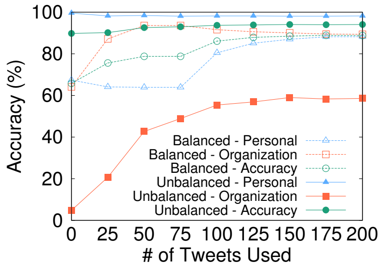

Figure 3: Performance when varying the number of tweets for classifying an account's type.

	Balanced			Natural		
	Org.	**Per.**	**% Acc.**	**Org.**	**Per.**	**% Acc.**
Our Method	89.4	88.2	88.8	58.6	98.2	94.1
Majority Class	100.0	0.0	50.0	0.0	100.0	90.6

Table 2: Classifier performance in two conditions.

setting and 94.1% in the naturally-distributed setting. Table 2 shows the type-specific accuracies and the performance of the majority-class baseline. Both classifiers attained a statistically significant performance improvement over the baseline, with $p < 10^{-6}$ for both using McNemar's test. The two classifiers illustrate the trade-off for end-users needing to maximize accuracy for discovering organization accounts (balanced condition) or for overall accuracy (natural condition). The full dataset is larger than that used by the classifier in the natural condition, which was limited to ensure comparability with the balanced condition. Therefore, we assessed whether using the full dataset of all 20,273 accounts would yield higher performance. Repeating the same setup with all data, The resulting classifier achieved an accuracy of 95.5, with a substantial boost in accuracy to 66.0 for classifying organizational accounts (vs. 58.6) and a small boost for personal accounts as well (98.6 vs. 98.2).

As a follow-up experiment, we analyzed classifier performance by varying the number of tweets used for testing an account's type using a fully-trained classifier, only varying the number of tweets used to generate the feature vectors for the test data. In the extreme case where an account has zero tweets, classifier performance is reliant entirely on structural features. Figure 3 shows the performance for both conditions and for the type-specific accuracies. In both conditions, accuracy is consistently high after 100 tweets are included. However, with fewer tweets, the two methods diverge in their accuracy: the balanced condition sees decreases performance for classifying personal accounts, while the natural condition sees a massive accuracy decrease for organizational accounts. Notably, when no tweets are used, the structural information alone still provides sufficient information to correctly classify some accounts.

5 Conclusion

This paper presents the first large-scale analysis of the presence and behavior of organizations on Twitter. Our work offers two main contributions. First, we created a high-quality dataset of over 20,000 active English-language accounts, revealing that organizations comprise 9.4% of Twitter accounts but are more heavily connected in the social network. Second, using this dataset, we created two highly accurate classifiers, one attaining 95.5% accuracy on all account types and a second designed for maximum accuracy at identifying organization accounts, both released open-source as `http://networkdynamics.org/software/`.

References

Alowibdi, J. S.; Buy, U. A.; and Yu, P. 2013. Empirical evaluation of profile characteristics for gender classification on twitter. In *Proc. of ICMLA*.

Chang, C.-C., and Lin, C.-J. 2011. Libsvm: A library for support vector machines. *ACM Trans. on Intelligent Systems and Technology* 2(3):27:1–27:27.

Ciot, M.; Sonderegger, M.; and Ruths, D. 2013. Gender inference on twitter in non-english contexts. In *Proc. of EMNLP*.

Cohen, R., and Ruths, D. 2013. Political orientation inference on twitter: It's not easy. In *Proc. of ICWSM*.

De Choudhury, M.; Diakopoulos, N.; and Naaman, M. 2012. Unfolding the event landscape on Twitter: classification and exploration of user categories. In *Proc. of CSCW*, 241–244.

De Silva, L., and Riloff, E. 2014. User type classification of tweets with implications for event recognition. In *Proc. of the Joint Workshop on Social Dynamics and Personal Attributes in Social Media*.

Golbeck, J.; Grimes, J. M.; and Rogers, A. 2010. Twitter use by the u.s. congress. *Journal of the American Society for Information Science and Technology* 61(8).

Jurgens, D. 2013. That's what friends are for: Inferring location in online social media platforms based on social relationships. In *Proc. of ICWSM*.

Li, J.; Ritter, A.; and Hovy, E. 2014. Weakly supervised user profile extraction from twitter. In *Proc. of ACL*.

Lipka, N., and Stein, B. 2011. Classifying with co-stems. In *Advances in Information Retrieval*. Springer. 307–313.

Liu, Y.; Kliman-Silver, C.; and Mislove, A. 2014. The Tweets They are a-Changin': Evolution of Twitter Users and Behavior. In *Proc. of ICWSM*.

Nguyen, D.; Gravel, R.; Trieschnigg, D.; and Meder, T. 2013. "How Old Do You Think I Am?" A Study of Language and Age in Twitter. In *Proc. of ICWSM*.

Nguyen, D.; Smith, N. A.; and Rosé, C. P. 2011. Author age prediction from text using linear regression. In *Proc. of the ACL-HLT Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*.

Saleem, H. M.; Xu, Y.; and Ruths, D. 2014. Novel situational information in mass emergencies: What does twitter provide? In *Proc. of HumTech*.

Schwartz, H. A.; Eichstaedt, J. C.; Kern, M. L.; Dziurzynski, L.; et al. 2013. Characterizing geographic variation in well-being using tweets. In *Proc. of ICWSM*.

Tumasjan, A.; Sprenger, T. O.; Sandner, P. G.; and Welpe, I. M. 2010. Predicting elections with twitter: What 140 characters reveal about political sentiment. In *Proc. of ICWSM*.

Yin, P.; Ram, N.; Lee, W.-C.; Tucker, C.; Khandelwal, S.; and Salathé, M. 2014. Two sides of a coin: Separating personal communication and public dissemination accounts in twitter. In *Advances in Knowledge Discovery and Data Mining*. Springer. 163–175.

Zamal, F.; Liu, W.; and Ruths, D. 2012. Homophily and latent attribute inference: Inferring latent attributes of Twitter users from neighbors. In *Proc. of ICWSM*.