Elites Tweet?

Characterizing Verified Twitter Users

Indraneil Paul (IIIT Hyderabad), Abhinav Khattar (IIIT Delhi), Shaan Chopra (IIIT Delhi), Ponnurangam Kumaraguru (IIIT Delhi), Manish Gupta (Microsoft India)

ICDE
MACAU 2019

Outline

A: PROBLEM AND MOTIVATION

- Characterizing verified Twitter users
- Understanding what sets them apart

B: DATASET DESCRIPTION

- Description of data collection
- Summary data statistics

C: NETWORK ANALYSIS

- Significance of centrality metrics
- Network structure findings

D: ACTIVITY ANALYSIS

- Changes of Tweeting patterns with real-world events

Motivation

Reasons to care and intended outcomes

Existing Literature

Previous human-annotated studies have demonstrated an authenticated status as one of the most robust predictors of **positive credibility** on Twitter.

This is backed up by subsequent findings:

1. Most **authentic** non-verified users on Twitter are within 7 degrees of separation of a verified user
2. A substantial majority of **spam** handles on Twitter are located within 7-10 degrees of separation from verified users

Thus, network distance from the core of verified users is also a reliable indicator of a non-verified user's credibility.

Visual Incentive

1. **Presence of authority and authenticity indicators:** Lends further credibility to the Tweets made by a user handle

2. **Presentation over relevance:** Psychological testing reveals that credibility evaluation of online content is influenced by its **presentation** rather than its relevance or apparent **credulity**

Attaining verified status might lead to a user's content being more frequently **liked** and **retweeted**.

Heuristic Models

The average user devotes only **three seconds** of attention per Tweet. This is symptomatic of users resorting to content evaluation heuristics.

One such relevant heuristic is the **Endorsement heuristic**, which is associated with credibility conferred to content by visual markers.

The presence of a marker such as a **verified badge** could hence, be the difference between a user reading a Tweet in a congested feed or completely ignoring it.

Heuristic Models

Another pertinent heuristic is the **Consistency heuristic**, which stems from endorsements by several authorities. This is important because a verified user on one social media platform is likelier to be verified on other platforms as well.

Hence, we posit that possessing a verified status can make a world of difference in the **outreach/influence** of a brand or individual in terms of the extent and quality.

Dataset

Collection sources, methods and summary

Collection Approach

We queried the Twitter REST API for the following:

1. The **@verified** handle on Twitter follows all accounts on the platform that are currently verified. We queried this handle on the **18th of July 2018** and extracted the user IDs.

2. We obtained the user objects for all verified users and subsetted for **English** speaking users.

3. For each verified user, we also queried the API in order to obtain the list of **outlinks** to other verified users.

Collected Metadata

For each verified member, we collected the following **metadata**:

1. Followers count
2. Friends count
3. Status count
4. Public list memberships
5. Tweet time series

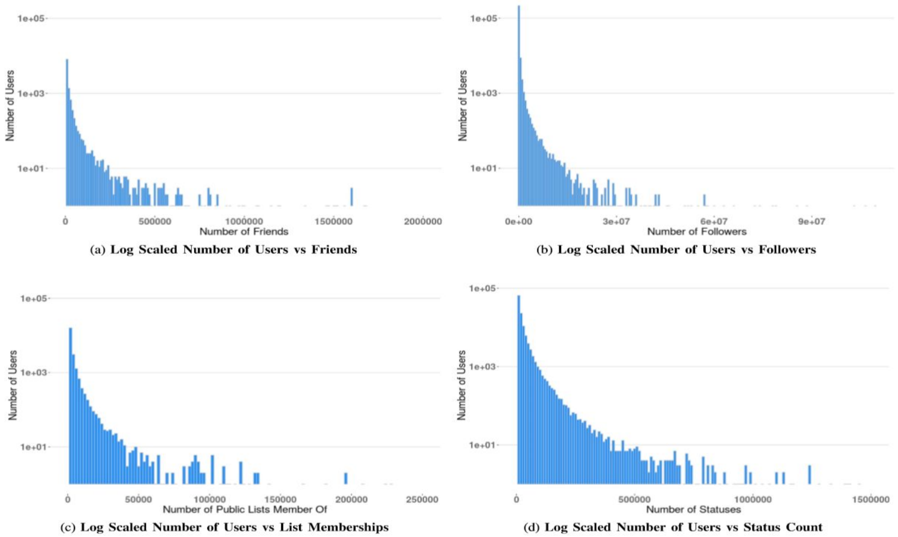

(a) Log Scaled Number of Users vs Friends

(b) Log Scaled Number of Users vs Followers

(c) Log Scaled Number of Users vs List Memberships

(d) Log Scaled Number of Users vs Status Count

Verified User Network

231,235

English language Twitter users

0.00148

Density

79,213,811

Network links

342.55

Average degree

6,027

Isolated users

6,251

Connected components

Miscellaneous Trivia

114,815

Most connected user:
Influencer **@6BillionPeople**

-0.04

Degree assortativity

0.1583

Low avg. clustering coefficient

Network Analysis

Delving into network centrality and connectivity

Attracting Components

Attracting components are components in a directed graph in which, if a random walk enters, it can never leave.

The acquired network consists of 6091 **attracting components**.

At the core of these components lie **famous** personalities (high in-degree users) who do not follow any other handle.

Power Law

Power-law is a key component in characterizing degree distribution of networks gathered from various **sources**. It refers to the presence of the following distributional property:

$$p(x) = Cx^{-\alpha}$$

This is closely related to the concept of the **Pareto distribution** or the 80-20 rule, where 20 percent of an entity is responsible for 80 percent of its characteristics.

We explore the presence of power laws in the network degree distribution and laplacian eigenvalue distribution.

Eigenvalue Distribution

We computed the 10,000 largest **eigenvalues** of the **Laplacian** matrix. The eigenvalues were computed using the **power iteration** method in existing solvers.

Inference of power-law parameters α and x_{min} is done using the continuous maximum-likelihood algorithm. Continuous MLE inference for the degree distribution yields parameter estimates of 3.18 for α and 9377.26 for x_{min} with a **p value** of 0.3

This is in keeping with earlier such findings in Laplacian eigenvalue distributions of synthetic and real world undirected social network datasets.

Degree Distribution

Further, we carry out a similar inference procedure for the out degree distribution of the nodes.

Inference of power-law parameters α and x_{min} is done using the discrete maximum-likelihood algorithm. Discrete MLE inference for the degree distribution yields parameter estimates of 3.24 for **α** and 1334 for **x_{min}** with a **p value** of 0.13

Our findings are in contrast with the absence of a power-law in the degree distribution when analyzing the whole Twitter network, as reported by existing work.

Reciprocity

The verified network has a **reciprocity rate** of 33.7%. This is lower than usually seen in other social networks such as Flickr (68%) due to the prevalence of brands and third-party sources of curated and crawled information, which typically do not **reciprocate engagements**.

This is higher than the previously reported reciprocity among the directed links in the entire Twitter network (22.1%). This is likely due to a larger core of publicly relevant and **consequential personalities** within this sub-graph of the Twitter network. This leads to a rarer occurrence of one sided follower-followee relationships.

Degrees of Separation

Existing work such as the **6 degrees of separation** and the **small-world model** after named after findings that many social and technological networks possessed small average path lengths.

The verified network is even more extreme in this aspect with an **average node distance** of 2.74 which is much lower than previous sampling estimates for all of Twitter (3.43, 4.12)

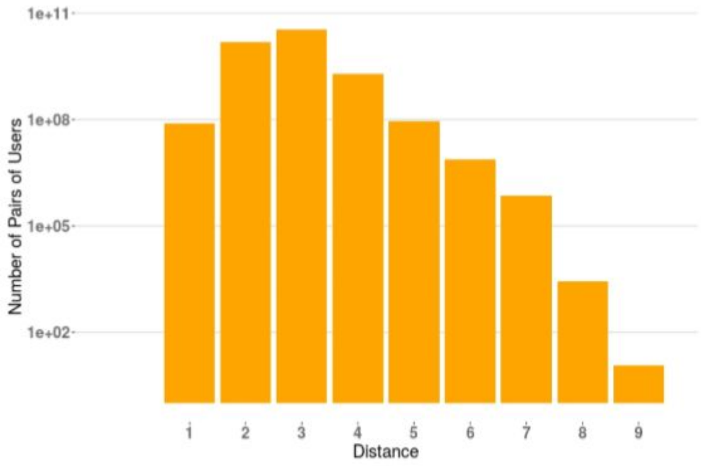

Bio Analysis

Each user on Twitter can have a **biography** (or bio) allowing him/her to describe themselves using a limited number of characters.

We attempt to gain insights from some of the most popular unigrams, bigrams and trigrams occurring in the bios of verified users.

We also filter out n-grams constituted largely of non-informative words.

A running theme common to all three cases is the dominance of **journalists** and news and weather outlets. Being a preeminent journalist in an English media outlet seems to be one of the surest ways to get verified on Twitter.

Bio Analysis

The most frequent **unigrams** portray several underlying themes such as:

1. They include cross-links to other social media handles (e.g. Instagram)
2. Personal descriptors (e.g. Father)
3. Professional descriptors (e.g. Tech)

Bigrams and trigrams reiterate a largely similar narrative, dominated by generic descriptors (e.g. Official Account) and business descriptors (e,g, Weather Alerts)

Bio Analysis

Bigram	Occurrences
Official Twitter	12166
Official Account	2788
Award Winning	2270
Follow Us	2268
Co Founder	1581
Husband Father	1540
Opinions Own	1222
New Album	1088
Singer Songwriter	1043
Co Host	933
Latest News	904
Breaking News	898
Anchor Reporter	855
Rugby Player	799
Managing Editor	769

Trigram	Occurrences
Official Twitter Account	5457
Official Twitter Page	1774
Weather Alerts EN	847
Emmy Award Winning	475
New York Times	464
Editor in Chief	461
Best Selling Author	296
Professional Rugby Player	253
Wall Street Journal	252
Professional Baseball Player	241
Report Crime Here	238
Award Winning Journalist	223
For Customer Service	174
Olympic Gold Medalist	174
Monday to Friday	174

Network Centrality

We delve into how a user's **centrality** in this network correlates with conventional metrics of reach such as **follower** and **list membership** count.

Public list membership has been shown to be a robust predictor of **influence** and **topical relevance** on Twitter.

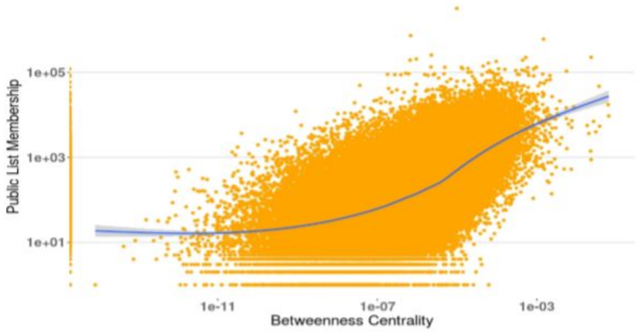

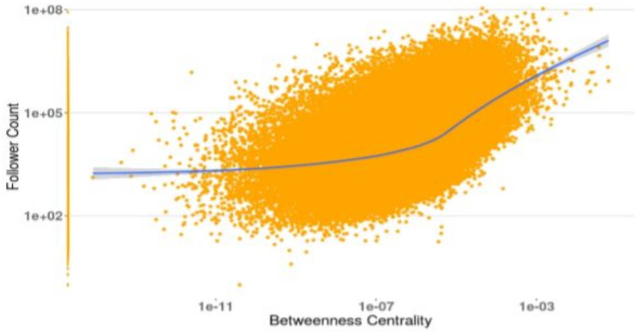

Network Centrality

We observe that public list membership and follower count in the entire Twitter network is **positively correlated** with **PageRank** and **Betweenness** centrality of that user in the English verified user sub-graph.

This backs up the general perception that a verified status is afforded, not just as a mark of **authenticity**, but also sufficient public interest.

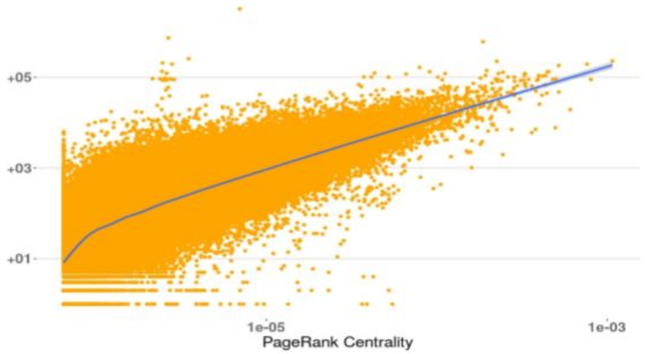

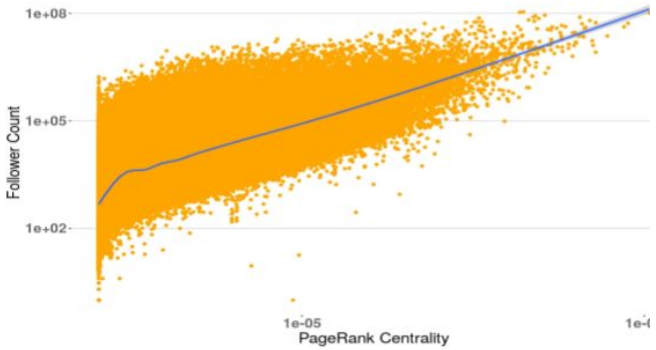

Activity Analysis

Digging into user activity patterns

Autocorrelation

We check for existing **auto correlations** in the time series using the **Ljung-Box** and the **Box-Pierce** portmanteau tests.

If the **p values** returned by the test are greater than 0.05, then the time-lagged correlation cannot be ruled out with a **95% significance** level. The Ljung-Box and Box-Pierce test results indicate a maximum p value of 3.81×10^{-38} and 7.57×10^{-38} respectively, thus strongly ruling out any lagged correlation.

This counters intuitive expectations that there would be a significant auto correlation in a week's lag given that **activity rates** on Sundays are reliably **lower** than those on weekdays.

Tweet Activity Pattern

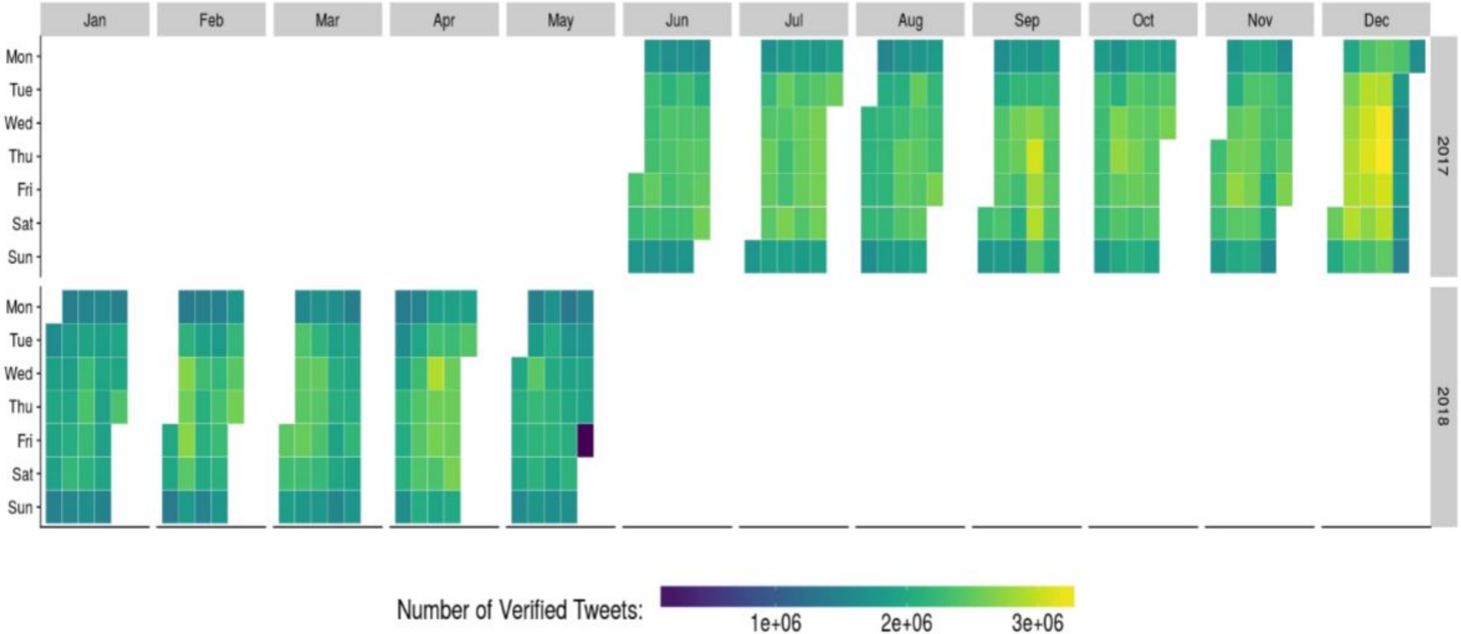

Stationarity

We next inquire whether the activity time series is stationary or not using a time series changepoint detection mechanism called **Pruned Exact Linear Time** (PELT). We assume that this time series is drawn from a normal distribution, with mean and variance that can change at a discrete number of **change-points**. We use the PELT algorithm to maximize the **log-likelihood** for the means and variances of the underlying distribution with a penalty for the number of change-points.

Results from several runs of the algorithm are recorded while cooling down the **penalty factor** and ramping up the number of change-points. Dates that fall in the change-point list in a significant number of runs of the algorithm are considered viable change-point candidates

We only find weak evidence for a changepoint around Christmas of 2017.

Stationarity

Existing work on smaller social networks, such as **Gab**, reveal that the activity time series drastically change in response to **socio-political events** occurring outside the network.

Hence, to investigate further, we employ an **Augmented Dickey-Fuller** test with both a **constant term** and a **trend term**. For upwards of 250 observations (we have 366) the critical value of the test is −3.42 when using a constant and a trend term at the **95% significance level**. If the test statistic value is more negative than the critical threshold, we reject the null hypothesis of a unit root and conclude the presence of stationarity.

Our test, returns a test statistic of −3.86 which is significantly **more negative** than the **critical threshold**, thus strongly suggesting stationarity

Key Contributions

Dataset

Released a fully featured dataset of 400k+ users, containing 79+ million edges and 494+ million Tweet time-stamps.

Characterization

We are the first study characterizing the connectivity and activity levels of verified users on Twitter.

Comparison

We compare the results to existing analytical results for the entire Twitter network.

Future Applications

1. **Superior verification heuristic**
   Aforementioned deviations likely constitute a unique fingerprint for verified users which can be leveraged gauge the strength of a user's case for such status

2. **Influence measure**
   Centrality and connectivity within the Twitter verified network may be utilized as a surrogate influence measure

3. **Realistic synthetic network generation**

Research Acknowledgements

IIIT Hyderabad

IIIT Delhi

Microsoft India

Thanks!

Any questions?

Find me at **ineil77.github.io**

Contact me at **indraneil.paul@research.iiit.ac.in**

ICDE MACAU 2019