On the Influence of Twitter Trolls during the 2016 US Presidential Election

Nikos Salamanos,[1] **Michael J. Jensen,**[2] **Xinlei He,**[3] **Yang Chen,**[4] **Michael Sirivianos**[5]

[1,5]Cyprus University of Technology {nik.salaman@cut.ac.cy; michael.sirivianos@cut.ac.cy}

[2]University of Canberra {Michael.Jensen@canberra.edu.au}

[3,4]Fudan University { xlhe17@fudan.edu.cn; chenyang@fudan.edu.cn }

Abstract

It is a widely accepted fact that state–sponsored Twitter accounts operated during the 2016 US presidential election spreading millions of tweets with misinformation and inflammatory political content. Whether these social media campaigns of the so–called "troll" accounts were able to manipulate public opinion is still in question. Here we aim to quantify the influence of troll accounts and the impact they had on Twitter by analyzing 152.5 million tweets from 9.9 million users, including 822 troll accounts. The data collected during the US election campaign, contain original troll tweets before they were deleted by Twitter. From these data, we constructed a very large interaction graph; a directed graph of 9.3 million nodes and 169.9 million edges. Recently, Twitter released datasets on the misinformation campaigns of 8,275 state–sponsored accounts linked to Russia, Iran and Venezuela as part of the investigation on the foreign interference in the 2016 US election. These data serve as ground–truth identifier of troll users in our dataset. Using graph analysis techniques we qualify the diffusion cascades of web and media context that have been shared by the troll accounts. We present strong evidence that authentic users were the source of the viral cascades. Although the trolls were participating in the viral cascades, they did not have a leading role in them and only four troll accounts were truly influential.

1 Introduction

The Russian efforts to interfere in and manipulate the outcome of the 2016 US presidential election were unprecedented in terms of the size and scope of the operations. Millions of posts across multiple social media platforms gave rise to hundreds of millions of impressions targeting specific segments of the population in an effort to mobilize, suppress, or shift votes [9]. Trolls were particularly focused on the promotion of identity narratives [10], though that does not distinguish them from many other actors during the election [17]. The Special Counsel's report described this interference as "sweeping and systematic" ([15], vol 1, 1). Russia demonstrated an impressive array of tactics for producing significant damage to the integrity of the communication spaces where Americans became informed and discussed their political choices during the election [14].

While Russia's efforts continue "unabated" [20], it is likely they and others will seek to target the American election in 2020 as well as continuing to target elections in Europe and elsewhere. It is important, therefore to identify the operational tactics of social media influence operations if we are to promulgate adequate defenses against them in the future.

There is a considerable debate as to whether state–sponsored disinformation campaigns that operated on social media were able to affect the outcome of the 2016 US Presidential election. While there is a large body of work which tried to address this question from distinct disciplinary angles [2, 9, 17] a conclusive result is still missing. There are several obstacles that any empirical study on this subject has dealt with: (i) the lack of complete and unbiased Twitter data – the Twitter API returns only a small sample of the users' daily activity; (ii) The tweets from deactivated profiles are not available; (iii) The followers and followees lists are not always accessible, hence the social graph is unknown. Moreover, the disinformation strategies that the operators of the state–sponsored accounts had employed is also unclear. A naive approach such as the one that flooded the network with fake–news is not obviously the case. A study of Russian social media activity has found that the majority of the communications are not obviously false [16]. It is equally possible that the operators had performed advanced manipulation techniques such as first building a reliable social profile aiming to engage a group of followers and then transmitting factually correct, but otherwise deceptive and manipulative claims advancing the political objectives of the disinformation campaign. Hence, text mining and machine learning techniques might not perform well under this scenario.

In this paper we measure the impact of troll activities on the virality of the ambiguous political information that had been shared on Twitter during the 2016 US Presidential election. For that purpose, a very large directed graph has been constructed by the interactions between the users (tweet replies and mentions). The graph consists of 9.3 million nodes and 169 million edges and it has been constructed based on two Twitter datasets: (i) A collection of 152.5 million tweets that was downloaded using the Twitter API during the US presidential election period (from September 21 to November 7, 2016). Hence, we have access to original troll tweets that have yet to be deleted by Twitter. (ii) A collection of original troll tweets which have been released by Twitter itself as part of the investigation on foreign interference in the 2016 US election – the misinformation cam-

paigns of 8,275 state–sponsored accounts linked to Russia, Iran and Venezuela states. Using graph analysis techniques and classification, we are able to identify the group of users that were most probably the driving force of the viral cascades that have been shared by the troll accounts.

Contribution: Our primary contributions are as follows:

- We construct one of the largest graphs in the literature, which represents the interactions between state–sponsored troll accounts and authentic users in Twitter during the period of 2016 US Presidential election. This is an approximation of the original followers–followees social graph.

- We present strong evidence that the trolls' activity was not the main cause that led to viral cascades of web and media material in Twitter. The experimental results clearly show that the authentic users who had close proximity with the trolls in the graph were the most active and influential part of the population and their activity was the driving force of the viral materials. A possible scenario is that instead of injecting new content, these trolls bandwagon or "resonate" with communities online with which they sought to form relationships. They do so, particularly by targeting opinion leaders in these communities. This is consistent with previous literature on information warfare tactics [4].

2 Related Work

2.1 Diffusion of disinformation

In [19] the authors investigated the diffusion cascades of true and false rumors distributed on Twitter from 2006 to 2017; approximately 126,000 rumor cascades which have been spread by 3 million people. The rumors had been verified as true or false by six fact–checking organizations. The main funding of this study is that false news diffused faster and more broadly than the truth and also human behavior contributes more to the spread of falsity than the trolls.

[3] examined 171 million tweets collected during five months prior to the 2016 US presidential election. From this collection, they analyzed 30 million tweets shared by 2.3 million users that contained at least one web–URL linking to a news outlet website (outside of Twitter). The 25% of these news were either fake or biased representing the spreading misinformation on Twitter. Then, in order to investigate the flow of information, the retweet networks are constructed for each category of news. Two users i and j are connected by the directed edge (j, i) if user i retweeted a tweet of user j. Hence, the edges represent the direction of information flow.

[8] investigated the extent to which Twitter users had been exposed to fake news during the 2016 US presidential election. The findings suggest that only a small fraction of 1% of the population was responsible for the diffusion of 80% of the fake news. Moreover, they proposed policies which had they been adopted by the social media platforms they would have reduced the spread of disinformation.

In [21, 22] the authors analyzed the characteristics and strategies of 5.5K Russian and Iranian troll accounts in Twitter and Reddit. Moreover, using the *Hawkes Process* they compute an overall statistical measure of influence that quantifies the effect these accounts had on social media platforms, such as Twitter, Reddit, 4chan and Gab.

[1] examined the Russian disinformation campaigns in the 2016 US election on Twitter. The study is based on 43 million posts shared on Twitter by 5.7 million distinct users (September 16 to November 9, 2016). The study focuses on the characteristics of the *spreaders*, namely the users that had been exposed and shared content produced by Russian trolls. They showed that existing techniques such as the *Botometer* [5] are able to effectively identify troll accounts.

An analysis of the role of Russian trolls on Twitter during the 2016 US election is presented in [11]. The *time-sensitive semantic edit distance* (a text distance metric) is proposed for the visualization and qualitative analysis of trolls' strategies such as *left–leaning* and *right–leaning*.

2.2 Identifying malicious activity

A well-known method for identifying troll accounts on Twitter is the *Botometer* (a.k.a. BotOrNot) introduced by [5]. The Botometer is a publicly available platform for estimating whether existing Twitter accounts have the characteristics of social bots. Finally, in [18] the *Rumor Gauge* is proposed; a method for predicting the veracity of rumors in Twitter during real–world events. It is a system for automatic verification of rumors in real–time, before verification by trusted channels such as governmental organizations is performed.

3 Datasets

3.1 Ground–truth Twitter data

Twitter has recently released a large collection of tweets of the state–sponsored troll accounts as part of Twitter's election integrity efforts[1]. We requested the unhashed version which consists of the tweets of Twitter accounts identified as Russian, Iranian and Venezuelan – 25 million tweets shared by 8,275 troll accounts (see Table 1). These troll IDs served as ground–truth identifiers of the troll users in our tweets collection.

Table 1: Ground–truth data

| Source: 8,275 troll accounts | | Target | |
Total tweets: 25,076,853		trolls	real users
Replies	1,549,742	2352	410,779
Retweets	8,617,208	3159	531,374
Mentions	10,641,427	2885	1,661,716

We observe that the majority of trolls' actions were retweets and mentions. Moreover, the target users were mostly real users, namely, Twitter accounts that were not in the ground truth troll IDs list. As we describe in the next section, we see the tweet–replies and tweet–mentions as *actions*, i.e. they represent the interactions from the users who performed these actions to the users who received them. For instance, when a user

---

[1]https://about.twitter.com/en_us/values/elections-integrity.html

i replies to a tweet of user j then the *source* is i while the *target* is j. Hence, in Table 1 with the term *source* we are referring to the 8,275 troll users.

3.2 Our Twitter dataset

The analysis in this paper is based on 152.5 million tweets from 9.9 million users. The tweets were downloaded using the Twitter streaming (1%) API in the period before and up to the 2016 US presidential election – from September 21 to November 7, 2016 (47 days). The tweets' track terms were related to political content such as "hillary2016", "clinton2016", "paul2016", "trump2016" and "donaldtrump2016"(see Appendix, Table 5 for the complete list of the track terms). Namely, a list of phrases used to determine which Tweets are delivered by the stream (see [2] for more details). These were collected using a Python script utilizing the Tweepy module. In addition to the tweet text, user screen name, and user ID, we also collected metadata including the hashtags and expanded URL data from Twitter, information on the account creation, user timezone, and user-supplied location and biographic information.

Table 2: The tweets collection

	real-users	trolls
user-IDs	9,939,698	822
Total tweets	152,479,440	35,489
Replies	12,942,628	160
Mentions	172,145,775	33,627
Retweets	N/A	N/A

Using the troll ground–truth IDs we identified 35.5K tweets from 822 troll accounts (see Table 2).

4 Methodology

4.1 The graph of interactions

In this paper, we followed a graph theoretical approach, namely we map users to nodes and we map the interactions between users to edges. We construct the graph based on the tweets collection we presented in the previous section – 152.5 million tweets collected during 47 days. The actions between the users are either replies or mentions (we had not collected the retweets). Each directed edge (i, j) corresponds to a tweet–action from user i to user j; either user i had replied to a tweet of user j or he had mentioned user j in his tweet, or both. So, we distinguish between replies and mentions. Both are actions from one user to another. Hence, we have the graph of users' interactions which is a directed multigraph (i.e., multiple edges are permitted between any pair of nodes) consisting of:

- nodes: (i) 821 trolls; (ii) 9,321,061 real users.

- edges: 169,921,921.

- 31,660 edges: from 659 trolls to 9,371 real users.

- 670,605 edges: from 121,924 real users to 285 trolls.

- ego–net nodes: 127,192 real users (based on both the in–edges and out–edges with the trolls)

4.2 The region of influence

The main goal of this study is to quantify the influence trolls had over real (authentic) users. A challenge we have to tackle is that only few trolls appear in our dataset and the actions/edges from trolls to real users are small in size. Moreover, we have 121,924 real-users which have at least one action with trolls (670,605 total actions). These users have either replied to troll tweets or they have mentioned a troll ID in their tweets. These real users are the troll ego-nets, namely the trolls' nearest neighbors – one–hop distance from the trolls in the multigraph. Yet, this number is not considerably large compared to the millions of total actions. Therefore, we have to expand the possible region of users by identifying a population that might have been influenced by troll activities. In order to achieve that, we use the web–URLs and media–URLs that have been shared by the trolls as an identifier of troll influence anywhere in the graph.

Specifically, we collect the *expanded URL*, *media expanded URL* and *media URL https* from all tweets. Then, we concentrate our attention to the URLs that have been shared by at least one troll user. For the rest of the paper we call this set *troll–URLs*. By that, we do not imply that these web contexts necessarily belong to trolls. They are just identifiers of the trolls' strategies, materials that the trolls wish to be viral. In other words, even if they are authentic web materials, they still reflect the political and strategic objectives of the trolls' disinformation campaigns.

We have 5,876,674 tweets from 620 trolls and 1,213,037 real users. These tweets contain: (i) 16,261 expanded URLs; (ii) 7,266 media expanded URLs; (iii) 7,259 media URL https.

We have three groups of users that shared troll-URLs: (i) the trolls; (ii) the real users that are part of the trolls' ego–nets; and (iii) the real users that are not directly connected with any of the trolls users. For the rest of the paper, we call all these users *spreaders*.

Then, we construct the directed induced subgraph based on this set of users, namely we select only the edges between the users that appear in the directed simple graph version of the multigraph. This subgraph represents the *region of influence* between trolls and real users, namely the largest region in the graph where the influence between trolls and real users could have taken place. In summary, we have: (i) nodes: 1,145,363; (ii) edges: 46,429,227; (iii) 560 trolls; (iv) 105,755 ego–net users; (v) 1,039,552 other users.

Hence, our goal is to compare the impact of these groups in the diffusion cascades of the trolls–URLs; the media and web context, which, as we mentioned previously, are a representation of the trolls disinformation objectives. The rationale of this approach is that if the trolls were the important factor in these diffusion cascades then their role during the 2016 US presidential election was substantial. Otherwise, the real users where
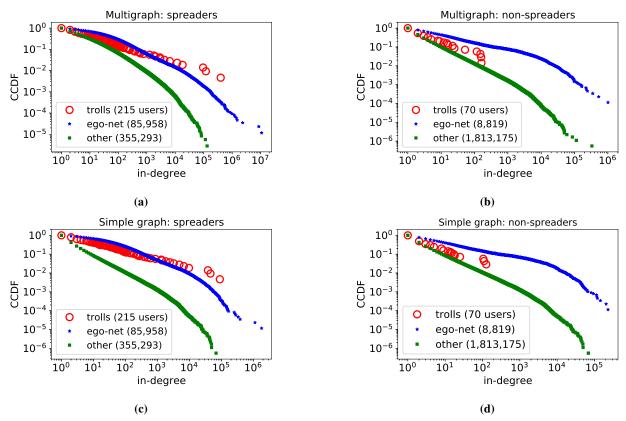
Figure 1: Multigraph & simple graph: CCDF of the in–degree for spreaders and non–spreaders.

responsible for the virality (if there was any) of the trolls' materials.

4.3 Diffusion Cascades

In order to compare the influence of trolls with that of real users, we have to reconstruct the path of information flow. Hence, for each troll–URL we construct the *time–inferred diffusion cascades* in a manner based on a causality assumption that has been previously used in the literature [6, 18, 19].

4.3.1 Causality hypothesis

We construct the *time–inferred diffusion cascades* for every troll–URL as follows:

Suppose that the URL x has been shared by user i on date t_i. First, we select the users with whom i was connected with in the graph before the date t_i, i.e., there were graph edges from i to them. In other words, we select the ego–net of i for $t < t_i$. If this set is not empty, then we examine whether any of these users had also shared the URL x on a date $t < t_i$. Suppose that two users j and k had shared URL x on dates $t_k < t_j$. Then, we assume that user i had been influenced by user j on adopting and diffusing the URL x and we update the cascade tree with the edge (j, i).

This approach has been previously used in the literature mainly for constructing retweet cascade trees based on followers–followees original graph. Since a user can retweet a tweet only once, there is no repetition in the user actions and the size of a cascade corresponds to the number of unique users involved in that cascade. In our case this is not true. The user actions may correspond to several tweets, hence repetitions are possible. In order to deal with this case, we assume that for a given URL, every user can be influenced only once by another user, the first time that he shared it.

4.4 Classification of users

We collect all the URLs and media URLs (the "troll–URLs") that were included in trolls tweets. Then, in our tweets collection, we identify all the tweets that have at least one troll–URL and have been shared either by trolls or by real users. So, our goal here is to compute a similarity measure between the troll accounts and the real users based on the context of their tweets. In this way, we extend the graph analysis in order to further characterize the interactions between the two groups of users, as well as the source of the viral cascades. For instance, if the very active and influential real users have high similarity with the trolls, then it is possible that these users were also trolls. To this end, we take the following steps:

1. We extract the textual data and train a *Doc2Vec* [13] model.

2. For each tweet, we get a fix-length vector to represent its content.

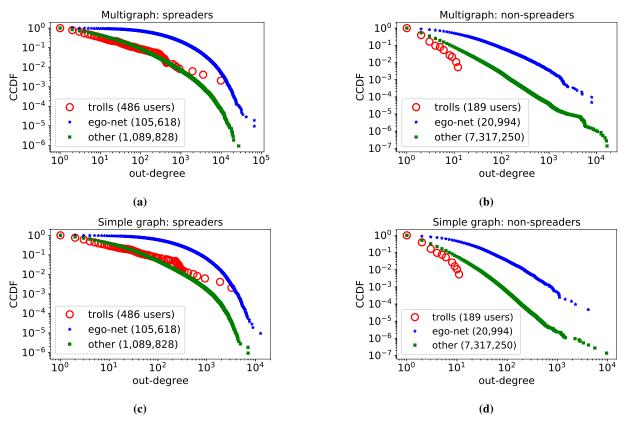3. For each user, we generate a sequence with time information and textual information.

4

(a)

(b)

(c)

(d)

Figure 2: Multigraph & simple graph: CCDF of the out–degree for spreaders and non–spreaders.

4. We feed sequences of users to a *Bi-LSTM* [7] model to classify them into different categories (3 categories in our setting: 1. Trolls; 2. The real-users that are connected with the trolls in the action graph; and 3. The rest of the real-users).

5. Due to the imbalance of size among the three groups, we randomly select 620 users from each group to form our dataset. 2/3 of them are used for training and the rest are used for testing.

Table 3: The classifier

class	precision	recall	f1-score	support
trolls (620)	0.8009	0.8204	0.8106	206
ego–net (105,755)	0.9259	0.6039	0.7310	207
others (1,107,282)	0.6533	0.8647	0.7443	207
micro avg	0.7629	0.7629	0.7629	620
macro avg	0.7934	0.7630	0.7619	620
weighted avg	0.7934	0.7629	0.7619	620

Table 3 shows the performance of our classifier. We can see that the F1-score of trolls detection reaches 0.8106 with both precision and recall higher than 0.8. This illustrates the usefulness of our classifier.

5 Results

The analysis is based on the comparison of the influence of three groups of users: (i) the trolls; (ii) the real-users that are directly connected with the trolls (the trolls' ego–net); (iii) the rest of the real users. As aforementioned, in order to identify a broad region of users that might have been influenced by the trolls, we use the web and media URLs as identifiers of trolls influence. We call *spreaders* the users that had shared at least one troll–URL. Hence, the three groups of users can be further divided in *spreaders* and *non–spreaders*.

The analysis we present in this section consists of the following steps:

1. We compute the degree distribution in the overall directed multigraph as well as in the corresponding directed simple graph (where only one edge is allowed between each pair of nodes).

2. We analyze the undirected version of the simple graph and we identify its connected components. Moreover, we compute the k-core values for the nodes in the largest connected component. We repeat this analysis for the region of influence.

3. We compute the classification scores of the users in the region of influence. Based on these scores, we estimate the similarity between trolls and real users.

4. We analyse the characteristics of the time inferred diffusion cascades of the trolls–URLs and we compute their
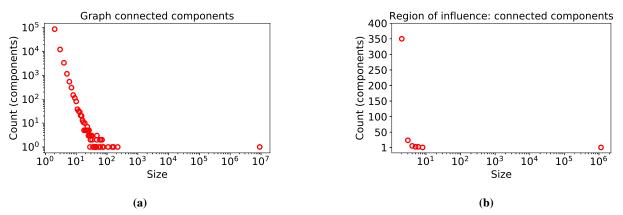
(a)

(b)

Figure 3: The connected components of the undirected versions: (a) simple graph; (b) region of influence
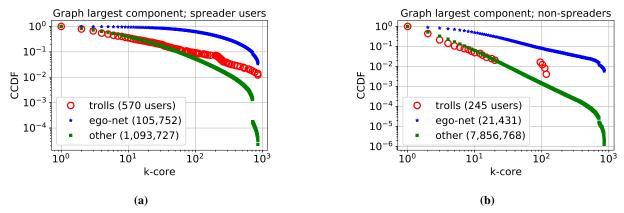

(a)

(b)

Figure 4: CCDF of the k-core values for the nodes in the largest connected component.
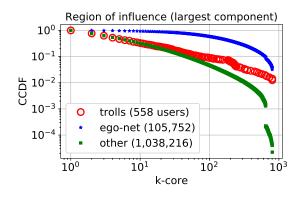

Figure 5: Region of influence, largest connected component: CCDF of k-core values

structural vitality.

5.1 Graph topology

5.1.1 Degree distribution

In Figures 1 and 2 we present the empirical complementary cumulative distribution (CCDF) of the in–degree and out–degree for each node/user in the directed multigraph as well as in the directed simple graph. We construct the simple graph by allowing only one edge for each pair of nodes that is already connected in the multigraph. Moreover, the multigraph has been

constructed based on the users' actions (replies and mentions) on other Twitter accounts and posts. Hence, the in–degree represents the nodes' popularity; that is, how many actions correspond to users interested in their tweets or their Twitter account in general. On the other hand, the out–degree is a measure of users' sociability/extroversion, i.e. how many actions a given user has performed on other Twitter accounts. Furthermore, it is important to compare the degree distributions in both graphs (multigraph and simple graph) because users with high degree in the multigraph do not necessarily have large degree in the simple graph. For instance, a given user might have a large in–degree in the multigraph only because he is popular to a small group of people which is highly engaged with his Twitter account – they perform a large number of actions on the tweets of the user in question. Hence, the user will have small in–degree in the simple graph.

As mentioned earlier, the real users can be divided in four groups based on the graph proximity they have with the trolls and the material they have shared. Hence, we have the ego–net and the rest of the users as well as the spreaders and non–spreaders. As can be seen in Figures 1 and 2, the spreaders that belong to the trolls' ego–net are the most active and the most popular ones. Moreover, the degree distributions between multigraph and simple graph are almost identical for the four groups of users.
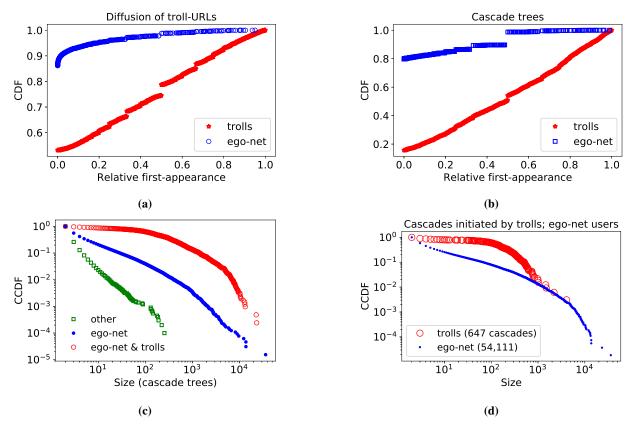
Figure 6: URL–diffusions with more than 100 distinct users: (a) CDF of the relative first–position in the URL–diffusions; (b) The relative first–position of the trolls as well as the ego–net users in the cascade trees; (c) CDF of the cascades size for the cascades where: i) at least one troll and one ego–net user are present, ii) at least one ego–net user and no trolls, iii) other cascades; (d) CCDF of the cascades size initiated by trolls and ego–net users.

5.1.2 Connected components

Here we examine the structure of the undirected version of the simple graph by computing the connected components. A connected component is a subgraph where for each pair of nodes i, j there is an undirected path – a graph traverse – from i to j. Since the subject of this study is the diffusion of information, the connectivity of a region implies that there is a possible path for information flow between the nodes that belong to this region.

The undirected simple graph consists of 9,321,882 nodes and 82,842,096 edges. We identify 104,954 connected components. Figure 3(a) presents the number of connected components for a given component size (i.e. number of nodes in the component) in a log–log plot. The largest part of the graph is well connected. The largest connected component – undirected subgraph – consists of 9,078,493 nodes and 82,698,678 edges while the second one has only 223 nodes. Moreover, 815 trolls and 127,183 ego–net users are in the largest connected component.

Regarding the region of influence, its undirected version consists of 388 connected components. The largest connected component has 1,144,526 nodes while the second largest has only 8 nodes (see Figure 3(b)).

5.1.3 k–core decomposition

We compute the k–*core decomposition* of the nodes in the largest connected component of the undirected versions of the overall graph and the region of influence. The k–core values is one of the most effective centrality measures for identifying the influential spreaders in a complex network [12].

In Figures 4(a) and 4(b), we present the empirical complementary cumulative distribution (CCDF) of the k–core values for spreaders and non–spreaders, respectively. The ego–net users have the largest k–core values in general, a strong evidence that they were the most influential part of the population. Moreover, Figure 5 presents the k-core values of the nodes in the largest connected component of the region of influence. Again, the ego–net users have in general larger k-core values than the trolls.

5.2 Cascade trees

We now turn our attention to the diffusion cascades of the troll–URLs. First, in Figure 6(a) we present a general result, the *relative first–appearance* of trolls and the ego–net users. For each URL, we rank the user IDs in descending order based on the date that they post their first tweet which contain the URL in question. This list is in fact the history of the diffusion of a given URL – a series of consecutive instances of sharing a given URL through tweeting. It is just the list of user IDs that
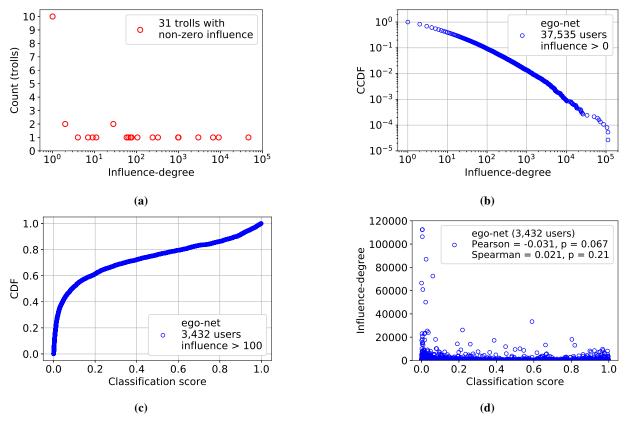
Figure 7: Influence: (a) Number of trolls with a given influence–degree value (for the troll users with with influence–degree > 0) ; (b) CCDF of the influence–degree for the ego–net users with non–zero influence values; (c) CCDF of the classification scores for the ego–net users with influence–degree > 100; (d) Scatter plot of the classification scores and the influence–degree for the ego–net users with influence–degree > 100.
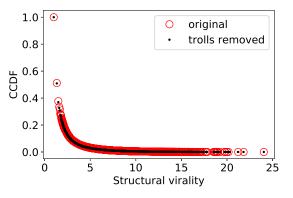

Figure 8: The structural virality of the original cascade trees as well as the cascade trees when the trolls have been removed from the multi-graph graph.

shared the URL, in chronological order. Recall that in our data, any URL might have been shared multiple times by the same user. The size of a given URL–diffusion is simply the number of user IDs belonging in the list (including the repetitions of IDs). In other words, the *relative first–appearance* of a user in the URL–diffusion is the index of the user in the list divided by the length of this list. In Figure 6(a), we show the empirical cumulative distribution (CDF) for the *relative position* of trolls and ego–net spreaders. Clearly, the largest part of the ego–net users appear before the trolls in the *diffusion–lists*. Almost 85%

of the ego–net users appear first in the URL–diffusions.

Next, we examine the URLs that have been shared by more than 100 distinct users. This selection led to 5,092 URL–diffusions. Then, based on the method we described in the previous section, we construct the cascade trees for each URL–diffusion. In this way, we obtain 88,714 cascade trees for 5,084 URLs that have at least one non–empty cascade tree. In summary: (i) 4,125 cascades have at least one troll user and all of them have at least one ego–net user; (ii) 64,525 cascades have at least one ego–net user and zero troll users. In the cascade trees each user appears only once, hence the size of a cascade is equal to the number of distinct IDs belonging in the tree. In Figure 6(b), we present the empirical complementary cumulative distribution (CCDF) of the *relative first–appearance* of trolls and the ego–net users in the cascade trees. We observe that 80% of the ego–net users appear very early in the cascades.

The most viral cascades are those with both trolls and ego–net users (see Figure 6(c)). At the same time, only 647 cascades had been initiated by troll users versus 54,111 by ego–net users. Roughly, 10% of the cascades that had been initiated by trolls have a larger size, yet the truly viral cascades had been initiated mainly by the ego–net users (Figure 6(d)). This evidence that although the trolls participated in the viral cascades, they did not have a leading role in them. Instead, the primary source of the viral cascades was the real users.

5.2.1 Influence

Did the trolls have an influential role in the diffusion cascades? We address this question by computing for each user i the *influence–degree*, namely the number of real users that have been influenced by i in the cascade trees he participates. Only 31 trolls have non–zero influence–degree. From them, 6 trolls have influence–degrees 978, 998, 2,961, 6,636, 9,040, and 46,224, respectively (Figure 7(a)). Hence, only four trolls where truly influential. Regarding the ego–net users: 37,535 have non–zero influence; 3,453 users, 511 users and 34 users have influence–degree larger than 10^2, 10^3 and 10^4, respectively (Figure 7(b)).

How similar with the troll accounts were the ego–net users? In order to estimate this, we use the classifications scores of the troll–category for the ego–net users with influence–degree > 100. We have 3,453 ego–net users with influence–degree > 100, where 21 users were in the training set and 3,432 in the test set. Hence, we present the classification scores for the 3,432 users, only. The 80% of the users have classification score smaller than 0.6 (see Figure 7(c)). Moreover, in Figure 7(d) we show the scatter plot of the classification scores and the influence-degree for the ego–net real users with influence larger than 100. The classification scores – the similarity of the real users with the trolls in terms of tweets content – represent the independent variable while the influence–degree the dependent one. The most influential users have low classification scores. Furthermore, the Pearson and Spearman correlation coefficients are ($r = -0.031$, p-value $= 0.067$) and ($r_s = 0.021$, p-value $= 0.21$), respectively. Hence, there is neither a linear nor a monotonic relationship between the two variables.

5.2.2 Structural virality

We conclude the analysis with the computation of the *structural virality* for each cascade tree. The structural virality of a cascade tree T with $n > 1$ nodes is the average distance between all pairs of nodes in a cascade [6]. That is:

$$\nu(T) = \frac{1}{n(n-1)} \sum_{i=1}^{n} \sum_{j=1}^{n} d_{ij} \qquad (1)$$

where T is the cascade tree with n nodes and d_{ij} is the shortest path between the nodes i and j. The intuition is that $\nu(T)$ is the average depth of nodes when we consider all nodes as the root of the cascade.

We compute the structural virality for the 88,714 cascade trees of 5,084 URLs that were discussed previously. We examined two cases: (i) the original cascades (88,714); (ii) the cascade trees that we have reconstructed for the case where the trolls have been removed from both the graph and the 5,084 URL–diffusions. The exclusion of the trolls results in 95,539 cascade trees. This indicates that a portion of the original trees broke in smaller trees, and probably some trolls act as "bridges". That is, they connect paths of the overall diffusion flow. On the other hand, it seems that their role was not substantial. The exclusion of the trolls did not affect the distribution of the virality values (see Figure 8).

5.3 Summary of the results

For all the metrics we have applied, the trolls as a whole have an inferior role compared to the ego–net users. The ego–net users, namely the real users that had at least one interaction with the trolls, were the most influential population and the source of the viral cascades, as well. Table 4 provides the specifics.

Table 4: Top–k results: trolls vs ego–net *spreaders*

Metrics	trolls	ego–net
Popularity: in–degree $> 10^3$	12	5,223
Sociability: out–degree $> 10^3$	3	21,887
Nodes in the largest k-core	7	3,552
Source node ("patient–zero"): Number of cascades	647	54,111
Source node: number of cascades with cascade size $> 10^3$	12	725
influence–degree $> 10^3$	4	511

Furthermore, we have derived the following statistics:

1. Only 12 trolls but 5,223 ego–net users have in–degree larger than 10^3 (multigraph).

2. Only 3 trolls but 21,887 ego–net users have out–degree larger than 10^3 (multigraph).

3. Only 7 trolls but 3,552 ego–net users belong to the largest k-core region of the graph (max k-core value = 854).

4. Only 647 cascades had been initiated by trolls but 54,111 by ego–net users.

5. Only 12 viral cascades (size larger than 10^3) had been initiated by trolls but 725 by ego–net users.

6. Only 4 trolls had influence–degree larger than 10^3 but 511 for the ego–net users.

7. The most influential ego–net users have low similarity with the trolls.

Finally, although the trolls participated in viral cascades, their role was not substantial. Removing the trolls from the graph as well as from the URL–diffusions did not affect the distribution of the structural virality values.

6 Conclusion

In this paper, we have extensively studied the influence that state–sponsored troll Twitter accounts had during the 2016 US presidential election. We analyzed a very large graph which represents the interactions between trolls and real users and we concentrated our attention to the region of influence, namely the well-connected part of the graph where trolls could have influenced real users. We present strong evidence that the trolls' activity was not the source of the viral cascades. The authentic users who had close proximity with the trolls were the most influential part of the population and their activity was the driving force of the viral cascades.

7   Acknowledgments

8   Appendix

Table 5 presents the 77 track terms used in Twitter API.

Table 5: The track terms used in Twitter API

ben%20carson	bencarson
bernie%20sanders	bernie2016
bettercandidatethanhillary	carlyfiorina2016
carson2016	clinton
clinton2016	cruz2016
cruzcrew	cruzintocaucus
donaldtrump	donaldtrump2016
dumptrump	election2016
feelthebern	fiorina
fiorina2016	fitn
heswithher	hilaryclinton
hillary2016	hillaryclinton
hrc	huckabee
huckabee2016	imwithher
iwearebernie	jill%20stein
jillstein	johnkasich
kasich	kasich2016
kasich4us	letsmakeamericagreatagain
makeamericagreatagain	makeamericawhiteagain
marco%20rubio	marcorubio
martinomally	mikehuckabee
nevertrump	newyorkvalues
nhpolitics	nhpsc
omalley	paul2016
primaryday	randpaul2016
readldonaldtrump	realdonaldtrump
redstate	rick%20santorum
ricksantorum	rubio2016
rubiowa	sensanders
sentedcruz	stein2016
teamKasich	teamcarly
teamclinton	teamcruz
teamhillary	teammarco
teamrubio	teamtrump
ted%20cruz	tedcruz
the%20donald	thedonald
therealdonaldtrump	trump
trump2016	trumptrain
unitedblue	

References

[1] A. Badawy, K. Lerman, and E. Ferrara. Who falls for online political manipulation? In *Companion Proceedings of The 2019 World Wide Web Conference*, WWW '19, pages 162–168, New York, NY, USA, 2019. ACM.

[2] Y. Benkler, R. Faris, and H. Roberts. *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics*. Oxford University Press, 2018.

[3] A. Bovet and H. A. Makse. Influence of fake news in twitter during the 2016 us presidential election. *Nature Communications*, 10(7), 2019.

[4] H. G. Clark. *Information Warfare: The Lost Tradecraft*. Washington, DC: Narrative Strategies, 2017.

[5] C. A. Davis, O. Varol, E. Ferrara, A. Flammini, and F. Menczer. Botornot: A system to evaluate social bots. In *Proceedings of the 25th International Conference Companion on World Wide Web*, WWW '16 Companion, pages 273–274, Geneva, Switzerland, 2016. International World Wide Web Conferences Steering Committee.

[6] S. Goel, A. Anderson, J. Hofman, and D. J. Watts. The structural virality of online diffusion. *Management Science*, 62(1), 2015.

[7] A. Graves and J. Schmidhuber. Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Networks*, 18(5-6):602–610, 2005.

[8] N. Grinberg, K. Joseph, L. Friedland, B. Swire-Thompson, and D. Lazer. Fake news on twitter during the 2016 u.s. presidential election. *Science*, 363(6425):374–378, 2019.

[9] K. H. Jamieson. *Cyberwar: How Russian Hackers and Trolls Helped Elect a President What We Don't, Can't, and Do Know*. Oxford University Press, 2018.

[10] M. Jensen. Russian trolls and fake news: Information or identity logics? *Journal of International Affairs*, 71(1.5):115–124, 2018.

[11] D. Kim, T. Graham, Z. Wan, and M. Rizoiu. Tracking the digital traces of russian trolls: Distinguishing the roles and strategy of trolls on twitter. *CoRR*, abs/1901.05228, 2019.

[12] M. Kitsak, L. K. Gallos, S. Havlin, F. Liljerosand, L. Muchnik, H. E. Stanley, and H. A. Makse. Identification of influential spreaders in complex networks. *Nature Physics*, 2010.

[13] Q. V. Le and T. Mikolov. Distributed representations of sentences and documents. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pages 1188–1196, 2014.

[14] M. Mazarr, A. Casey, A. Demus, S. Harold, L. Matthews, N. Beauchamp-Mustafaga, and J. Sladden. *Hostile Social Manipulation: Present Realities and Emerging Trends*. Santa Monica: Rand Corporation, 01 2019. Available at https://www.rand.org/pubs/research_reports/RR2713.html.

[15] R. S. Mueller. *Report on the Investigation into Russian Interference in the 2016 Presidential Election*. Washington, DC: Department of Justice, 2019. Available at https://www.justice.gov/storage/report.pdf.

[16] B. Schafer. A view from the digital trenches - lessons from year one of hamilton 68, November 26 2018. Washington, DC: German Marshall Fund. Available at https://securingdemocracy.gmfus.org/a-view-from-the-digital-trenches-lessons-from-year-one-of-hamilton-68.

[17] J. Sides, M. Tesler, and L. Vavreck. *Identity Crisis: The 2016 Presidential Campaign and the Battle for the Meaning of America*. Princeton University Press, 2018.

[18] S. Vosoughi, M. N. Mohsenvand, and D. Roy. Rumor gauge: Predicting the veracity of rumors on twitter. *ACM Trans. Knowl. Discov. Data*, 11(4):50:1–50:36, July 2017.

[19] S. Vosoughi, D. Roy, and S. Aral. The spread of true and false news online. *Science*, 359(6380):1146–1151, 2018.

[20] C. Wray. A conversation with christopher wray, 09 2019. Council on Foreign Relations. Available at https://www.cfr.org/event/conversation-christopher-wray-0.

[21] S. Zannettou, T. Caulfield, E. De Cristofaro, M. Sirivianos, G. Stringhini, and J. Blackburn. Disinformation warfare: Understanding state-sponsored trolls on twitter and their influence on the web. In *Companion Proceedings of The 2019 World Wide Web Conference*, WWW '19, pages 218–226, New York, NY, USA, 2019. ACM.

[22] S. Zannettou, T. Caulfield, W. Setzer, M. Sirivianos, G. Stringhini, and J. Blackburn. Who let the trolls out?: Towards understanding state-sponsored trolls. In *Proceedings of the 10th ACM Conference on Web Science*, WebSci '19, pages 353–362, New York, NY, USA, 2019. ACM.