Introduction

Networks appear all around us in different forms. And in a world dominated by social media today, the impacts of networks can have huge implications in the way news spreads and influences people on the platform. Interestingly, both real and fake news spread rapidly on the platform, necessitating the need for understanding what drives such cascades in the first place.

Research has demonstrated that people who are like-minded tend to form connections with each other in online social networks. The paper, *Birds of a Feather: Homophily in Social Networks*[1], goes on to explain the demographic, social, geographical and behavioral characteristics that help explain this phenomenon, while highlighting that “ties between nonsimilar individuals also dissolve at a higher rate, which sets the stage for the formation of niches (localized positions) within social space.”

Likewise studies [2] have also shows the effect of group pressure on individual judgements, indicating how people tend to follow the herd even if they were intentionally misled with falsities. This can prove to have negative consequences in social media where people often turn to for help in order to validate their opinions and beliefs in the absence of proof by direct observation. For example, a tweet falsely claiming that Snapchat was planning to revert to its old design was retweeted 1.5 million times before it was deleted [3]. Even more problematically, Facebook has come under harsh criticism for serving as a breeding ground for Russian propaganda and allowing fake news to flourish on the site.

Even more recently, an analysis of approximately 126,000 stories tweeted by 3 million people more than 4.5 million times between 2006 and 2017 [4] found that “falsehood diffused significantly farther, faster, deeper, and more broadly than the truth in all categories of information, and the effects were more pronounced for false political news than for false news about terrorism, natural disasters, science, urban legends, or financial information.” It also found that false news was more novel than real news, suggesting that people were more likely to share novel information, and that “bots accelerated the spread of true and false news at the same rate, implying that false news spreads more than the truth because humans, not robots, are more likely to spread it.”

Twitter, Bots and Russia

On December 20, 2017, Twitter admitted that Russian trolls actively meddled in national conversations during the U.S. Presidential Elections [5]. It reported that it had identified over 3,814 accounts related to Kremlin-based propaganda outfit called Internet Research Agency [6] and several thousand bots, totally up to over 200,000 tweets [7].

This Twitter data is available to download via [7] and this paper explores the Neo4j dataset that’s available to explore the different connections between the tweets and users themselves.

<u>Hashtags:</u>

Exploring the tweets, it reveals some interesting patterns in that maga (Make America Great Again), tcot (Top Conservatives on Twitter), trump, neverhillary, trumppence16, ccot (Conservative Christians on Twitter), crookedhillary turn out to be some of the most used hashtags, signalling that these troll accounts tweeted content that favored Donald Trump over Hillary Clinton:

hashtag	count
"politics"	3442
"maga"	2286
"tcot"	2048
"trump"	1922
"news"	1757
"pjnet"	968
"neverhillary"	967
"trump2016"	801
"merkelmussbleiben"	796
"trumppence16"	744
"rejecteddebatetopics"	620
"wakeupamerica"	617
"hillary"	609
"hillaryclinton"	582
"trumptrain"	555
"blacklivesmatter"	542

While there are 12919 unique hashtags used by the troll accounts, checking which two hashtags were most frequently used together reveals the below result.

hashtag_1	hashtag_2	count
"maga"	"tcot"	311
"pjnet"	"tcot"	286
"maga"	"trump"	281
"maga"	"trumppence16"	263
"ccot"	"tcot"	231
"p2"	"tcot"	231
"maga"	"trumptrain"	213
"gop"	"tcot"	191
"news"	"world"	187
"hillary"	"trump"	169
"maga"	"neverhillary"	162
"maga"	"trump2016"	158
"cleveland"	"politics"	145
"tcot"	"trump"	134
"trump2016"	"trumptrain"	126
"ccot"	"gop"	120

Tweets:

A check for the kinds of tweets posted by these troll accounts showed that these accounts were spreading a mix of original as well as retweeted content. For example, the user "AmelieBaldwin" made a total of 6092 retweets of a total of 9262 tweets available in her timeline. Likewise, the user "GiselleEvns" had the most number of original tweets (6007), a significantly higher number than the user who comes second in terms of posting original content: "TheFoundingSon" with 3377 original tweets.

screen_name	totalTweets	totalRetweets	totalReplies	originalContent
"AmelieBaldwin"	9262	6092	0	3170
"hyddrox"	6810	3634	0	3176
"GiselleEvns"	6634	626	1	6007
"PatriotBlake"	4150	1013	0	3137
"TheFoundingSon"	3586	157	52	3377
"MelvinSRoberts"	3366	231	0	3135
"mrclydepratt"	3262	82	0	3180
"brianaregland"	3260	258	0	3002
"LeroyLovesUSA"	3239	483	1	2755
"BaoBaeHam"	3215	87	0	3128
"melanymelanin"	3214	196	0	3018
"GarrettSimpson_"	3213	410	0	2803
"LauraBaeley"	3211	155	1	3055
"JeffreyKahunas"	3211	466	1	2744
"EmileeWaren"	3198	469	0	2729
"DatWiseNigga"	3196	53	0	3143

Checking the tweets posted by GiselleEvns, it appears to be mostly random, with new tweets made in every 4-5 second interval. This in itself lends credence to the fact that this account could have been a bot.

"GiselleEvns"	"for sure https://t.co/bl7EbSYJry"	"2017-01-25 15:11:42"
"GiselleEvns"	"RT @DocDarnell: #ItsRiskyTo be a risk analyst 👈"	"2017-01-25 15:11:36"
"GiselleEvns"	"RT @Danzig303: #ItsRiskyTo\r\nPost anything no matter how funny related to politics unless you are ready to argue with idiots"	"2017-01-25 15:11:31"
"GiselleEvns"	"RT @RossMoorhouse: #ItsRiskyTo\r\nIgnore the sign https://t.co/Eu6BgXi87q"	"2017-01-25 15:11:24"
"GiselleEvns"	"#ItsRiskyTo steal my tweets"	"2017-01-25 15:11:17"

Upon further checking the most retweeted tweets with a retweet count of more than 1000, it appears that there are very few accounts whose tweets get retweeted a lot. The top five include: TEN_GOP, Pamela_Moore12, Crytal1Johnson, SouthLoneStar and gloed_up. TEN_GOP had the most, with 271 tweets posted by the account retweeted more than 1000 times.

name	count
"TEN_GOP"	271
"Pamela_Moore13"	61
"Crystal1Johnson"	36
"SouthLoneStar"	8
"gloed_up"	7
"USA_Gunslinger"	6
"TrayneshaCole"	5
"TheFoundingSon"	5
"Jenn_Abrams"	3
"tpartynews"	3
"Luke_Jones13"	2
"Blk_Voice"	2
"BleepThePolice"	1
"redlanews"	1
"BlackNewsOutlet"	1
"RealTEN_GOP"	1

The next check was performed to see when these tweets were posted. Of a total of 15765 tweets made by accounts with over 15,000 followers, it can be gleaned that 11,115 tweets were made during 12:00 PM to 08:00 PM, with the tweet activity peaking around 03:00-04:00 PM. Around 6,246 tweets (out of the toal 11,115 tweets) were posted just in that span of two hours.

<u>What are these accounts tweeting?</u>

Another area of interest is to see what the troll accounts are tweeting. If they are including a URL in their tweets, the domain names of these URLs can be indicator as to whether they are tweeting true information or merely spreading propaganda. Querying the data set for this information, some of the most frequent domains appear to be the following: twitter.com, bit.ly, twibble.io, youtu.be (YouTube), wapo.st (Washington Post) and www.breitbart.com.

The domains can then also be used to check if they are intentionally spreading misinformation by tweeting links to articles that are nothing but “fake news”.

As bit.ly accounts for a significant number (4442), the next step undertaken was to expand the shortened URLs and see what domains they resolved to. The results of these steps are depicted as follows:

domain	count
"twitter.com"	7298
""	4533
"bit.ly"	4442
"twibble.io"	945
"youtu.be"	472
"wapo.st"	465
"ln.is"	459
"www.breitbart.com"	377
"dlvr.it"	335
"www.youtube.com"	305
"sh.st"	297
"fb.me"	250
"ow.ly"	242
"goo.gl"	214
"www.wcvb.com"	210
"buff.ly"	198

Upon checking further the tweets coming from the domain “twitter.com”, it appears that a majority of them (1379) come from accounts that have been suspended by Twitter, followed by Donald Trump (124), Hillary Clinton (95), and www.thehill.com (70).

Twitter Handle	Count
<tweets from suspended accounts>	1379
realDonaldTrump	124
HillaryClinton	95
thehill	70
politico	28
wikileaks	25
gloed_up	24
mitchellvii	23
FoxNews	23
TEN_GOP	17
PrisonPlanet	15
washingtonpost	15
dcexaminer	14
foxnews	14
CNN	13

Tracing back the domain names of URLs which had been shortened using different URL shorteners like “bit.ly”, "ln.is", "sh.st", "dlvr.it", "ow.ly", "buff.ly", "goo.gl" and "ift.tt", URLs from the following domains were found to be the most commonly shared:

Domain	Count
http://www.libertywritersnews.com/	102
http://truthfeed.com/	94
https://www.youtube.com/	75
http://www.breitbart.com/	69
https://www.reddit.com/	49
https://www.rawstory.com/	39
http://dailycaller.com/	30
http://newsninja2012.com/	27
http://www.nydailynews.com/	27
http://www.foxnews.com/	27
http://www.thegatewaypundit.com/	27
http://therealstrategy.com	24
https://townhall.com/	21
https://www.bizpacreview.com/	17
http://money.cnn.com/	17
http://freebeacon.com/	16
https://www.redstate.com/	16
https://www.conservativereview.com/	15
https://twitter.com/	14
https://www.washingtontimes.com/	14
https://www.americanthinker.com/	13
http://conservativebyte.com/	13
https://truepundit.com/	13
https://www.counterpunch.org/	12
http://www.freerepublic.com/	11
http://www.bipartisanreport.com/	10

A check of most of these domains like Liberty Writers News, TruthFeed, The Gateway Pundit reveal themselves as associated with extreme right and known to spread propaganda and fake news.

Interaction between trolls:

As we have seen above, most of the tweets made by the troll accounts seem to be automated in that they are posted in intervals of 4-5 seconds. The fact that such bots are present in a platform like Twitter gives us an indication that these bots can be part of a bigger bot network, tweeting similar content or retweeting each other, thus simulating interactions between real users on the social media network.

It's therefore useful to infer relationships between these troll accounts see how many of these accounts are retweeting other trolls. In other words, the amplification of messages by trolls. A sample of such relationships is depicted below.

The first figure shows a retweet made by troll account NoJonathanNo (Jonathan Norwood) that was originally tweeted by another troll account called BlackToLive: “RT @BlackToLive: A unique picture of a black man supporting Trump and police https://t.co/UETXNOgdHV”

The second figure explores a more intricate web of such retweet relationships between troll accounts. It shows the extent to which tweets posted by troll accounts were retweeted among other troll accounts, adding credence to the fact that there exist some troll accounts that were more popular than others and were actively tweeting propaganda.

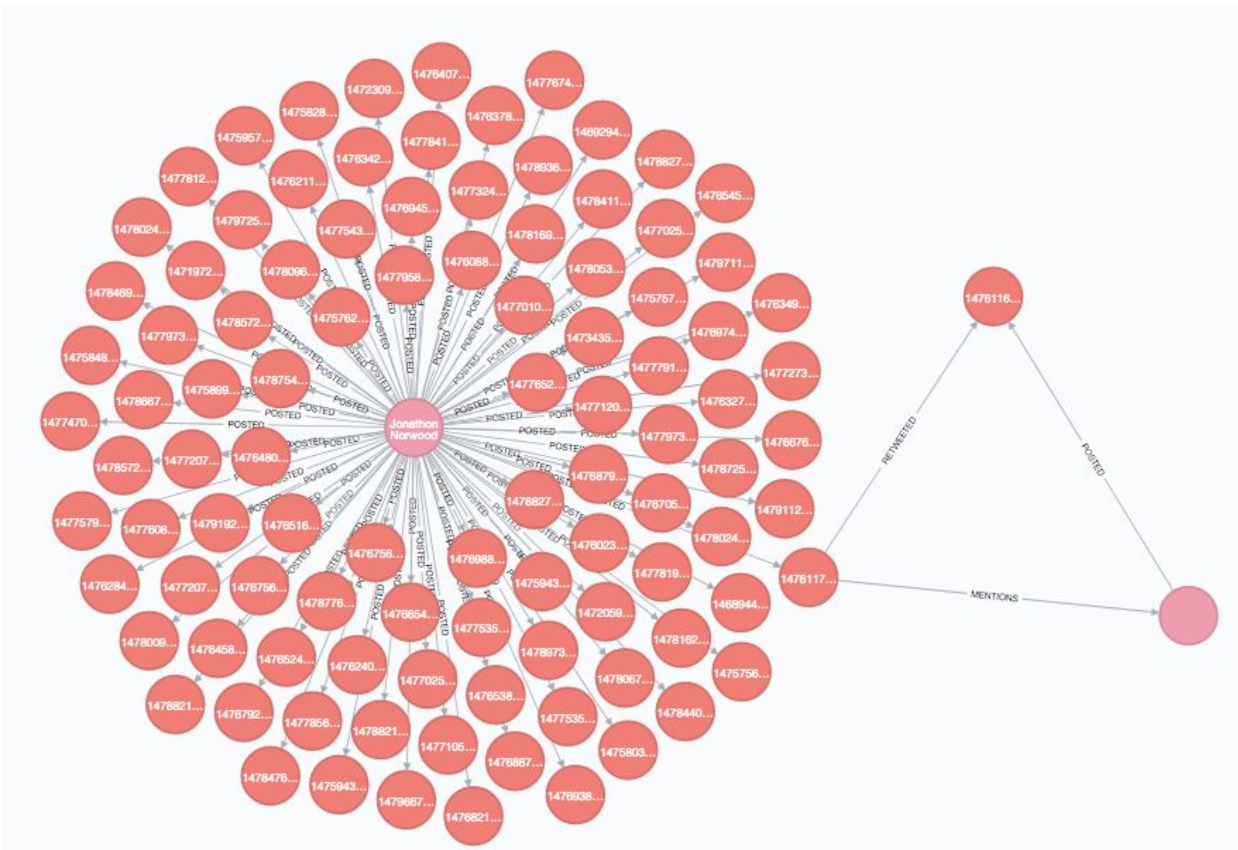

A tweet posted by user BlackToLive (in pink, right) is retweeted by Jonathan Norwood, a troll account suspended by Twitter

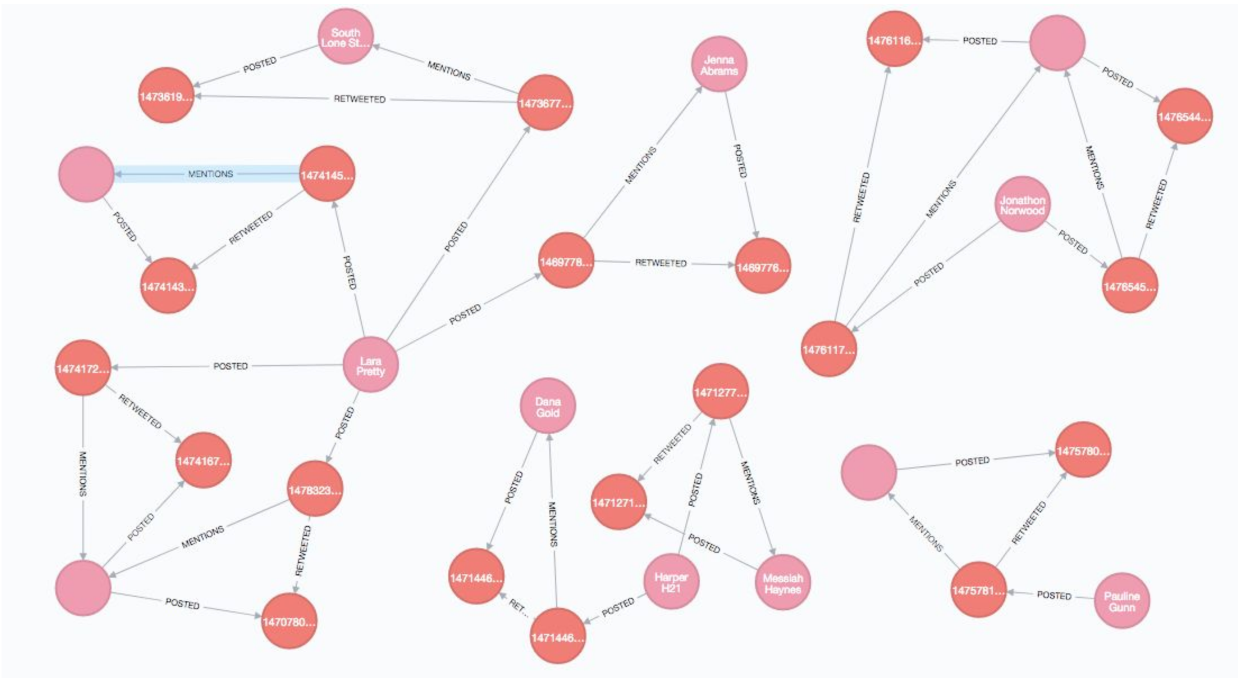

A pictorial representation that shows a sample of tweets that were retweeted among the troll accounts

An extension to this approach also involves finding the most important nodes in the network, which can done by making use of Neo4j's PageRank algorithm which recursively goes over the graph to identify how well a node is connected to others. Querying the database for the troll accounts with the highest pagerank scores reveals TEN_GOP to the top account, followed by TheFoundingSon, GiselleEvns and others.

troll	pagerank
"TEN_GOP"	10.3859635
"TheFoundingSon"	8.281644000000002
"GiselleEvns"	6.4624315
"tpartynews"	6.3289815
"ChrixMorgan"	4.231436500000001
"NotRitaHart"	3.9303325000000005
"BlackToLive"	3.9081304999999995
"DanaGeezus"	3.310725
"gloed_up"	2.7065620000000004
"DaileyJadon"	2.1636159999999998
"Pamela_Moore13"	2.0134805
"RealTEN_GOP"	1.8268460000000004
"March_for_Trump"	1.8126679999999995
"DominicValent"	1.7664449999999998
"MelvinSRoberts"	1.7552079999999999
"CassieWeltch"	1.7074124999999996

But identifying bot networks also requires adding a community property that allows us to partition one node as belonging to one community over the other. This can be achieved in Neo4j using labelPropagation algorithm using the retweet count as a metric to do the partitioning. The result of this step can be seen below:

members	community
["PrettyLaraPlace", "cassishere", "CarrieThornthon", "CynthiaMHunter", "USA_Gunslinger", "TEN_GOP", "KateRitterrrr", "deusXYX", "hyddrox", "AmelieBaldwin", "J0hnLarsen", "Pamela_Moore13", "DonnaBRivera", "DrMichaelGarcia", "Boryabuchin", "wadeharriot", "razvedchica_", "ErofeenkoAnton", "PaolaKinck", "happkendrahappy", "AlexsVladimirov", "SovetnikStatski", "Bocharnikov_V", "finley1589", "PatriotBlake", "Hollydler", "tpartynews", "SouthLoneStar", "BigSeanBeast", "NotBallinYet", "ImaSwerve", "KarenParker93", "RH0lbr00k", "WokeFromDay1", "QueennArielle", "ChesPlaysChess", "MissouriNewsUS"]	210
["CharlesJHarper", "IlikeBIGbuttand", "Adrienne_GG", "CurtisBigMan", "Gab1Aldana", "MelvinSRoberts", "BrucieDublin", "heyheyhailey", "CalebPaar", "dannythehappies", "holycrapchrix", "LoraGreeen", "Aiden7757", "hipppo_", "AmandaVGreen", "pureDavie", "brightandglory", "BGarner2107", "cascaseyp", "GiselleEvns", "DanaGeezus", "ChrixMorgan", "WorldOfHashtags", "CatelineWatkins", "HellieEdwards", "KenCannone", "abigailssilk", "DominicValent", "JmsCoxxx", "DickyIrwin", "MeggieONeil", "DonnieLMiller"]	188
["JeffreyKahunas", "WesternWindWes", "_NickLuna_", "JeanneMccarthy0", "LauraBaeley", "GarrettSimpson_", "EmileeWaren", "MichelleArry", "JacquelinIsBest", "CooknCooks", "mrclydepratt", "DorothieBell", "RyanMaxwell_1", "EvaGreen69", "heyits_toby", "LeroyLovesUSA", "JudeLambertUSA", "PatriotRaphael", "Aldrich420", "PriceForPierce", "Mil0Blake", "hollandpatrickk", "LazyKStafford", "_Billy_Moyer_", "c__wells", "evewebster373"]	47
["DaileyJadon", "traceyhappymom", "KathieMrr", "Jasper_Fly", "mr_clampin", "queenofthewo", "CassieWeltch", "NotRitaHart"]	268
["NoJonathonNo", "HilmKhloe", "melanymelanin", "BlackToLive", "JohnBranchh", "BrianTheLifter", "WillisBonnerr"]	42
["BleepThePolice", "gloed_up", "TrayneshaCole", "internalmemer"]	206
["PaulineTT", "RamonasNails", "Crystal1Johnson", "BLMSoldier"]	43
["russilanrogov", "LavrovMuesli", "WhiteHouseCards"]	342
["JarrardNorman", "Jenn_Abrams", "TheFoundingSon"]	191

While this only shows the extent of interaction of interaction between different troll accounts, a sample of some of the tweets from a couple of accounts show that they tend to retweet the same content, adding credence to the belief that they could possibly belong to a larger bot network.

source	target	count
"GiselleEvns"	"GiselleEvns"	78
"gloed_up"	"gloed_up"	56
"TheFoundingSon"	"TheFoundingSon"	52
"tpartynews"	"TEN_GOP"	47
"RealTEN_GOP"	"TEN_GOP"	40
"rightnpr"	"TEN_GOP"	37
"MelvinSRoberts"	"GiselleEvns"	37
"TrayneshaCole"	"gloed_up"	36
"TEN_GOP"	"tpartynews"	32
"RealTEN_GOP"	"tpartynews"	28
"rightnpr"	"tpartynews"	28
"queenofthewo"	"NotRitaHart"	25
"GiselleEvns"	"DaileyJadon"	24
"AmelieBaldwin"	"TEN_GOP"	21
"NotRitaHart"	"NotRitaHart"	21
"KathieMrr"	"NotRitaHart"	19

Conclusion

The internet has long been used as a tool for political activism and social control. The findings above illustrate that political conversation on Twitter was by and large dominated by bots and right wing accounts, and the number of tweets favouring Donald Trump was more than that of Hillary Clinton.

The term “fake news” is difficult to operationalize, but a sample of tweets posted by bots and other suspended accounts reveals that social media users shared many links to political news and information, but junk news, characterized by ideological extremism, misinformation and the intention to persuade readers to respect or hate a candidate or policy based on emotional appeals, was just as, if not more, prevalent than the amount of information produced by professional news organizations.

The findings also show that troll accounts tend to retweet each other’s tweets frequently and that there are a few accounts who are more widely followed than the others. These accounts were also found to be mostly right wing, mostly sharing content that was propaganda and fake news, lending credence to the belief that these accounts were actively working towards meddling in the U.S. presidential elections.

Queries

Here is a snapshot of Neo4j queries run on the dataset for reference:

1) Hashtag co-occurrence:

MATCH (h1:Hashtag)<-[:HAS_TAG]-(t:Tweet)-[:HAS_TAG]->(h2:Hashtag)
WHERE h1.tag < h2.tag
RETURN h1.tag as hashtag_1, h2.tag as hashtag_2, count(t) AS count
ORDER BY count DESC LIMIT 25

2) Total number of Hashtags in the dataset:

MATCH (h:Hashtag) return count(*)

3) Most used hashtags:

```
MATCH (:Troll)-[:POSTED]->(:Tweet)-[:HAS_TAG]->(ht:Hashtag)
RETURN ht.tag AS hashtag, COUNT(*) AS count
ORDER BY count DESC LIMIT 100;
```

4) The locations where the tweets containing most used hashtags are tweeted from:

```
MATCH (tr:Troll)-[:POSTED]->(t:Tweet)-[:HAS_TAG]->(ht:Hashtag)
with ht.tag AS hashtag, tr.location as location, COUNT(*) AS count
where count >= 100 and tr.location <> ""
return hashtag, location, count
order by count desc
```

5) The most popular Troll Tweets:

```
MATCH (tr:Troll)-[:POSTED]->(t:Tweet)
WHERE EXISTS(t.retweet_count)
RETURN tr.screen_name, tr.verified, tr.friends_count, t.text, t.retweet_count
ORDER BY t.retweet_count DESC LIMIT 100
```

6) The most unique trolls with the most retweeted tweets above 1000:

```
MATCH (tr:Troll)-[:POSTED]->(t:Tweet)
WHERE EXISTS(t.retweet_count) and t.retweet_count > 1000
with tr.screen_name as name, count(t) as count
return name, count
ORDER BY count DESC LIMIT 100
```

7) The years when these troll accounts were created:

```
MATCH (t:Troll) WHERE NOT t.created_at = ""
RETURN substring(t.created_at, 26) AS year, COUNT(*) AS num
ORDER BY year
```

8) The times these tweets were posted:

```
MATCH (t:Tweet)<-[:POSTED]-(u) WHERE exists(t.text) AND u.followers_count > 15000
RETURN u.user_key as user, u.followers_count as followers, substring(t.created_str,0,4) as year,
toInteger(substring(t.created_str,11,2)) as timeofday, size(split(t.text," ")) as wordcount,
size((t)-[:HAS_TAG]->()) as hashtagcount
order by timeofday, user desc
```

9) The kinds of contents posted by troll accounts - original tweets, replies and retweets:

```
MATCH (tr:Troll)-[:POSTED]->(tw:Tweet) WITH tr, tw
OPTIONAL MATCH (tw)-[:RETWEETED]-(rt:Tweet)
OPTIONAL MATCH (tw)-[:IN_REPLY_TO]-(irp:Tweet)
RETURN distinct tr.screen_name as screen_name, count(tw) as totalTweets, count(rt) as
totalRetweets, count(irp) as totalReplies, (count(tw) - (count(rt) + count(irp))) as originalContent
ORDER BY totalTweets DESC;
```

10) The tweets of GiselleEvns

```
MATCH (tr:Troll)-[:POSTED]->(tw:Tweet) WITH tr, tw
where tr.screen_name = "GiselleEvns"
```

```
RETURN tr.screen_name, tw.text, tw.created_str
order by tw.created_str desc
```

11) The medium through which these tweets were posted:

```
MATCH (:Troll)-[:POSTED]->(tw:Tweet)-[:POSTED_VIA]->(s:Source)
RETURN s.name AS source, COUNT(*) AS num
ORDER BY num DESC
```

12) Which trolls are interacting with others:

```
MATCH p=
(:Troll)-[:POSTED]->(:Tweet)<-[:RETWEETED]-(:Tweet)<-[:POSTED]-(:Troll)
RETURN p LIMIT 10
```

13) What sources are they tweeting?

```
MATCH (tr:Troll)-[:POSTED]->(t:Tweet)-[:HAS_LINK]-(u:URL)
RETURN distinct u.expanded_url as domain, count(u) as count
ORDER BY count desc

MATCH (tr:Troll)-[:POSTED]->(t:Tweet)-[:HAS_LINK]-(u:URL)
WITH t, replace(replace(u.expanded_url, "http://", '' ), "https://", '') AS url
RETURN head(split(url, "/")) as domain, count(t) as count
ORDER BY count desc
```

13) Calculate pagerank:

```
CALL algo.pageRank(
"MATCH (t:Troll) RETURN id(t) AS id",
"MATCH (r1:Troll)-[:POSTED]->(:Tweet)<-[:RETWEETED]-(:Tweet)<-[:POSTED]-(r2:Troll)
RETURN id(r2) as source, id(r1) as target",
{graph:'cypher'})
```

14) Most important nodes:

```
MATCH (t:Troll) WHERE EXISTS(t.pagerank)
RETURN t.screen_name AS troll, t.pagerank AS pagerank
ORDER BY pagerank DESC LIMIT 25
```

15) Community detection:

```
CALL algo.labelPropagation(
"MATCH (t:Troll) RETURN id(t) AS id",
"MATCH (r1:Troll)-[:POSTED]->(t:Tweet)<-[:RETWEETED]-(:Tweet)<-[:POSTED]-(r2:Troll)
RETURN id(r2) AS source, id(r1) AS target, COUNT(t) AS weight","OUTGOING",
{graph:'cypher', write: true, iterations: 200})
```

16) Troll Communities:

```
MATCH (t:Troll) WHERE EXISTS(t.partition)
RETURN COLLECT(t.screen_name) AS members, t.partition AS community
ORDER BY SIZE(members) DESC LIMIT 20
```

17) Cross-tweets across two different troll accounts:

```
MATCH (r1:Troll)-[:POSTED]->(t:Tweet)<-[:RETWEETED]-(:Tweet)<-[:POSTED]-(r2:Troll)
RETURN r2.screen_name AS source, r1.screen_name AS target, COUNT(t) AS count
ORDER BY count DESC LIMIT 50
```

Neo4j Data Model of Troll Dataset

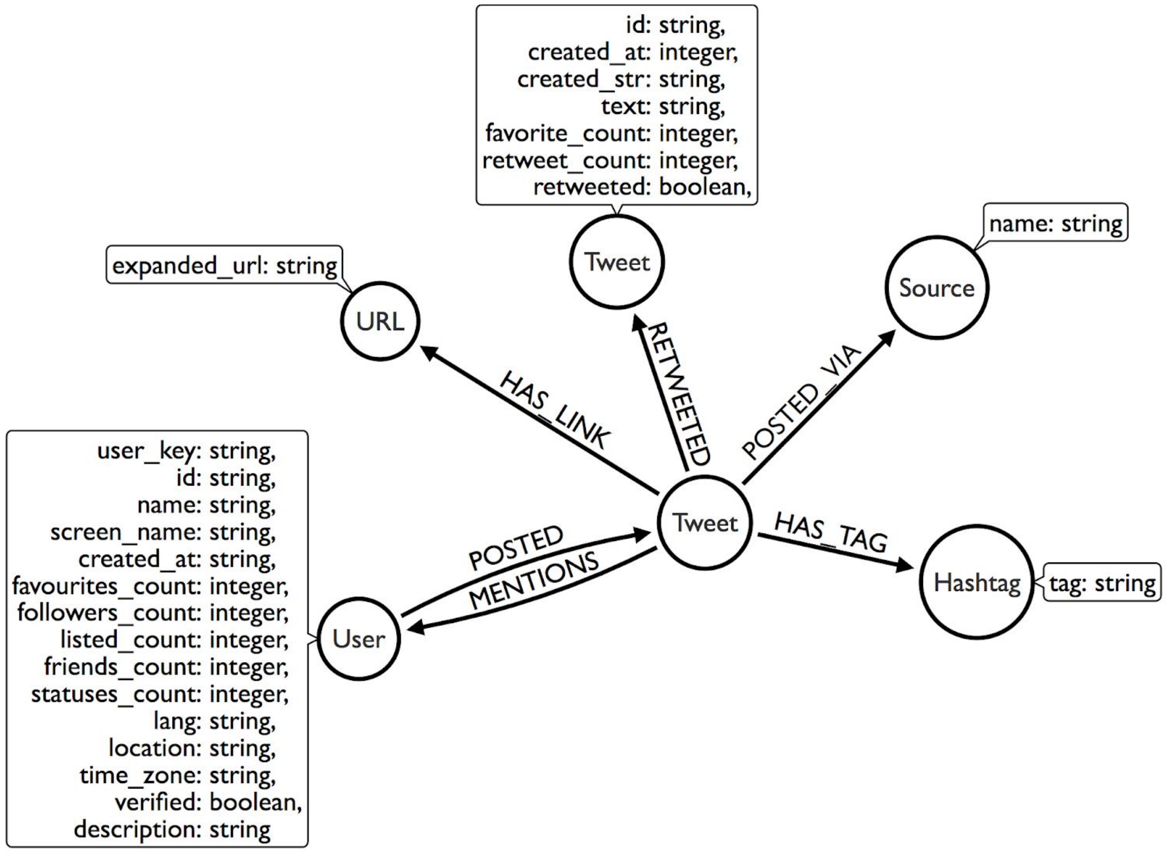

References

[1] Birds of a Feather: Homophily in Social Networks http://aris.ss.uci.edu/~lin/52.pdf

[2] Effects of group pressure upon the modification and distortion of judgments. https://www.gwern.net/docs/psychology/1952-asch.pdf

[3] No, Snapchat won't reverse its update for 50,000 retweets https://mashable.com/2018/02/09/snapchat-redesign-50000-retweets-fake-photoshop/#0KPsYWy4Kqqg

[4] The spread of true and false news online http://science.sciencemag.org/content/sci/359/6380/1146.full.pdf

[5] Russian trolls went on attack during key election moments https://www.nbcnews.com/tech/social-media/russian-trolls-went-attack-during-key-election-moments-n827176

[6] Twitter admits there were many more Russian trolls on its site during the 2016 U.S. presidential election https://www.recode.net/2018/1/19/16911538/twitter-russia-congress-2016-election-trump

[7] Twitter deleted 200,000 Russian troll tweets. Read them here.
https://www.nbcnews.com/tech/social-media/now-available-more-200-000-deleted-russian-troll-tweets-n844731

[8] Neo4j Explorer
https://neo4j.com/sandbox-v2/

[9] Media Bias/Fact Check
https://mediabiasfactcheck.com/