
Network Centrality

Part 3 – Case Studies

Saptarshi Ghosh

Department of CSE, IIT Kharagpur

Social Computing course, CS60017

Includes material borrowed from various online sources, including slides by Lada Adamic

Case Study 1

Measuring User Influence in Twitter: The Million Follower Fallacy, Cha et al., ICWSM 2010

Different influence measures for OSN

- Compared different influence measures for the Twitter social network
 - Network structure based:
 - In-degree (number of followers)
 - Activity based:
 - Number of times a user is retweeted
 - Number of times a user is mentioned
 - Two measures compared using Spearman's rank correlation coefficient
-

Results of comparison

- Across all three measures, top influentials were public figures (politicians, celebrities, ...) and websites (news media sites)
- But top influentials according to indegree have low overlap with top influentials according to activity

Table 1: Spearman's rank correlation coefficients

Correlation	All	Top 10%	Top 1%
Indegree vs retweets	0.549	0.122	0.109
Indegree vs mentions	0.638	0.286	0.309
Retweets vs mentions	0.580	0.638	0.605

Case Study 2

Understanding and Combating Link Farming in the
Twitter Social Network, Ghosh et al., WWW 2012

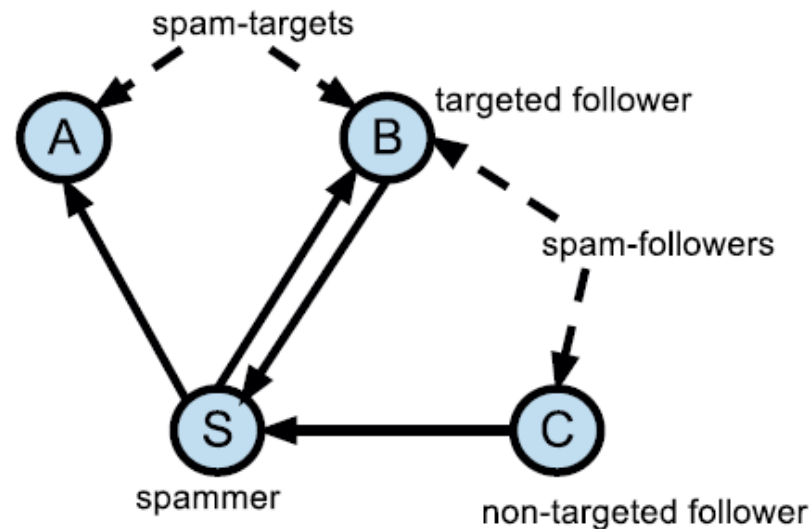
Why link farming in Twitter?

- Twitter has become a Web within the Web
 - Vast amounts of information and real-time news
 - Twitter search becoming more and more common
 - Search engines rank users by follower-rank, Pagerank to decide whose tweets to return as search results
 - High indegree (#followers) seen as a metric of influence
 - Link farming in Twitter
 - Spammers follow other users and attempt to get them to follow back
-

Started by identifying spammers

- Identified 41,352 spammers in Twitter
 - Accounts suspended by Twitter
 - Had posted blacklisted URLs
 - Many of the spammer accounts had high number of followers (in-degree)
 - Average in-degree for random user: 36
 - Average in-degree for spammer: 234
-


Terminology for spammers' links



- Spam-targets: users followed by spammers
- Spam-followers: users who follow spammers
 - Targeted: spam-target and spam-follower
 - Non-targeted: follow spammers without being targeted

Link farming by spammers

- Spammers farm links at large scale
 - Over 15 million users (27% of total) targeted by 41,352 spammers (0.08% of total)
- 1.3 million spam-followers
 - 82% are targeted → spammers get most links by reciprocation



Who are the spam-followers?

- Non-targeted spam-followers
 - Mostly sybils / hired helps of spammers
 - Most have now been suspended by Twitter
 - Targeted spam-followers
 - Ranked on the basis of number of links to spammers
 - 60% of follow-links acquired by spammers come from the top 100,000 targeted followers – LINK FARMERS
 - Are the link farmers themselves spammers?
-

Are link farmers themselves spammers?

- No, over 80% are **real, popular, active users**
 - Most of them are marketers trying to promote their business or some product
 - Includes some **verified accounts**
 - Many link farmers within the top 5% users according to PageRank
-

Top link-farmers: examples

Top 5 link farmers according to	
#links to spammers	Pagerank
Larry Wentz: Internet, Affiliate Marketing	Barack Obama: campaign staff
Judy Rey Wasserman: Artist, founder	Britney Spears: It's Britney
Chris Latko: Interested in tech. Will follow back	NPR Politics: Political coverage and conversation
Paul Merriwether: helping others, let's talk soon	UK Prime Minister: PM's office
Aaron Lee: Social Media Manager	JetBlue Airways: Follow us and let us help

Why are popular users link farming?

- Social etiquette – you follow me, I follow you
 - Amass social capital in the network
-

Is it easy to farm links in Twitter?

- We created a Twitter account and followed some of the top targeted spam-followers
 - Followed 500 randomly selected link farmers
 - Within 3 days, 65 reciprocated by following back
 - Our account ranked within the top 9% of all users in Twitter in 3 days !!!



The problem with link farming

- Existence of a set of users from whom social links (hence social influence) can be farmed easily
 - Spammers easily gain links from popular users
 - Increases the PageRank of spammers as well
 - Leads to increases spam in Twitter search results
-

Combating link farming in Twitter

- Key challenges
 - Real, popular users engaged in link farming
 - Detecting and suspending spammers alone will not help
 - Discourage users from following others carelessly
 - Penalize users for following someone bad – lower the influence scores of users who follow spammers
-

CollusionRank

- Identify a seed set of known spammers
 - In PageRank style
 - Negatively bias initial scores towards the known spammers
 - Iteratively penalize users who follow spammers, or those who follow spam-followers
-

CollusionRank

Input: network, G ; set of known spammers, S ; decay factor for biased Pagerank, α

Output: Collusionrank scores, c
initialize score vector d for all nodes n in G

$$d(n) \leftarrow \begin{cases} \frac{-1}{|S|} & \text{if } n \in S \\ 0 & \text{otherwise} \end{cases}$$

/ compute Collusionrank scores */*

$c \leftarrow d$

while c not converged **do**

for all nodes n in G **do**

$$tmp \leftarrow \sum_{nbr \in followings(n)} \frac{c(nbr)}{|followers(nbr)|}$$

$$c(n) \leftarrow \alpha \times tmp + (1 - \alpha) \times d(n)$$


end for

end while


return c

How effective is CollusionRank?

- Compare ranks of **spammers** and **link farmers**
 - PageRank
 - CollusionRank
 - PageRank + CollusionRank



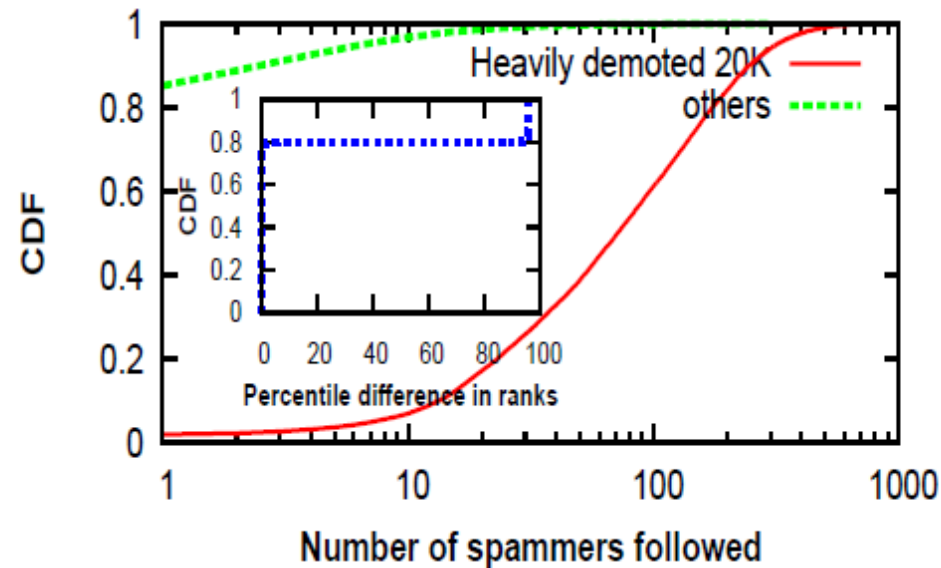
(a) Rankings of all 41,352 spammers



(b) Rankings of Top 100,000 capitalists

Pagerank + Collusionrank

- Selectively penalizes spammers & link-farmers
 - Out of top 100K according to Pagerank, 20K demoted heavily, rest 80% not affected much (inset)
 - The heavily demoted 20K follow many more spammers than the rest (main figure)



Case Study 3

Cognos: Crowdsourcing Search for Topic Experts
in Microblogs, Ghosh et al., SIGIR 2012

Topical search on Twitter

- Twitter has emerged as an important source of information & real-time news
 - Most common search in Twitter: search for trending topics and breaking news
 - Topical search
 - Identifying topical attributes / expertise of users
 - Searching for topical experts
 - Searching for information on specific topics
-

Prior approaches to find topic experts

- ❑ Research studies
 - ❑ Pal et. al. (WSDM 2011) uses 15 features from tweets, network, to identify topical experts
 - ❑ Weng et. al. (WSDM 2010) uses ML approach
 - ❑ Application systems
 - ❑ Twitter Who To Follow (WTF), Wefollow, ...
 - ❑ Methodology not fully public, but reported to utilize several features
-

Prior approaches use features extracted from

- ❑ User profiles
 - ❑ Screen-name, bio, ...
 - ❑ Tweets posted by a user
 - ❑ Hashtags, others retweeting a given user, ...
 - ❑ Social graph of a user
 - ❑ #followers, PageRank, ...
-

Problems with prior approaches

- User profiles – screen-name, bio, ...
 - Bio often does not give meaningful information
 - Information in users profiles mostly unvetted
 - Tweets posted by a user
 - Tweets mostly contain day-to-day conversation
 - Social graph of a user – #followers, PageRank
 - Does not provide topical information
-

We proposed

- Use crowdsourcing
 - How does the Twitter crowd describe a user?
 - Social annotations
 - Crowdsourced information collected using a feature called Twitter Lists
-



Pete Cashmore

@mashable NYC / SF

Breaking social media, tech and digital news and analysis from Mashable.com, the top resource and guide for all things web.

Updates from @mashable staff.

<http://mashable.com>

Tweets

Favorites

Following ▾

Followers

Lists ▾

mashable's lists



[@mashable/news](#)

A curated list of news organization's Twitter accounts.



[@mashable/tech](#)

Experts and sources to keep up with the latest in tech.



[@mashable/design](#)

Tweets and tips from designers.



[@mashable/food](#)

Love food? Here are chef's, cooks and others in food to follow.



[@mashable/celebrity](#)

Celebrities on Twitter.



[@mashable/journalism](#)

Journalists interested in the future of news media.



[@mashable/music](#)

Musicians on Twitter.



Pete Cashmore

@mashable NYC / SF

Breaking social media, tech and digital news and analysis from Mashable.com, the top resource and guide for all things web.

Updates from @mashable staff.

<http://mashable.com>

Tweets Favorites Following Followers Lists



nytimes The New York Times

Where the Conversation Begins. Follow breaking news, NYTimes.com home page articles, special features and more.



BBCNews BBC News

The latest stories, features and updates from BBC News



WSJ Wall Street Journal

Breaking news, investigative reporting, business coverage and features from The Wall Street Journal.



cnnbrk CNN Breaking News

CNN.com is among the world's leaders in online news and information delivery.

mashable's lists



@mashable/news

A curated list of news organization's Twitter accounts.



@mashable/tech

Experts and sources to keep up with the latest in tech.



@mashable/design

Tweets and tips from designers.



@mashable/food

Love food? Here are chef's, cooks and others in food to follow.



@mashable/celebrity

Celebrities on Twitter.



@mashable/journalism

Journalists interested in the future of news media.



@mashable/music

Musicians on Twitter.



Pete Cashmore

@mashable NYC / SF

Breaking social media, tech and digital news and analysis from Mashable.com, the top resource and guide for all things web.

Updates from @mashable staff.

<http://mashable.com>

Tweets Favorites Following Followers Lists

mashable's lists



@mashable/news

A curated list of news organization's Twitter accounts.



@mashable/tech

Experts and sources to keep up with the latest in tech.



@mashable/design

Tweets and tips from designers.



@mashable/food

Love food? Here are chef's, cooks and others in food to follow



@mashable/celebrity

Celebrities on Twitter.



@mashable/journalism

Journalists interested in the future of news media.



@mashable/music

Musicians on Twitter.



101Cookbooks 101 Cookbooks

Heidi Swanson from 101Cookbooks.com - Healthy, vegetarian recipes made from natural foods and seasonal produce.



epicurious epicurious

Written by Tanya Steel and the Epicurious editorial staff



LATimesfood LA Times Food

News, recipes + reviews from the LA Times Food staff, test kitchen + Daily Dish blog, by @renelynch.



TylerFlorence Tyler Florence 

Chef, Restaurateur, Wine Maker, Cookbook Writer, Shop Keep, Product Designer, Dad.

Using Lists to infer topics for users

- If U is an expert / authority in a certain topic
 - U likely to be included in several Lists
 - List names / descriptions provide valuable semantic cues to the topics of expertise of U



Barack Obama ✓

@BarackObama Washington, DC

2,513	683,900	11,864,332	162,383
Tweets	Following	Followers	Listed

Lists following BarackObama



@jyo_0827/news-and-politician
ニュースや政治家



@docsports_shs/national-government
national government tweets



@gazhazDCFC/politics



@GOPSackSuckers/truepatriots
Warriors for truth, justice, and equality.



@WendyMagley/interests



@DorianaLeonardo/internazionale



Amitabh Bachchan ✓

@SrBachchan Mumbai, India

11,314	146	1,920,019	19,416
Tweets	Following	Followers	Listed

Lists following SrBachchan



@sukhdam1999/famous



@raxus1/bollywood
Bolly actors, actresses, authors



@yadav657/actors
old is gold



@umairsaeed/celebrities



@Littl_Rock/bollywood-celebrities



@umairsaeed/entertainment

Mining Lists to infer expertise

ashton kutcher 
@aplusk

Movies TV by Patty Holmes
59 members

Hollywood by mpc2000
1 members

Stars by Nadine Schultz
Meine Lieblingsstars
4 members

Entertainment by Al Royce
54 members

Celebrities by Ben Elcomb
142 members

Celebs by KING5 Photog Jim
126 members

Hollywood by Praesidian
17 members

- Collect Lists containing a given user U
- Merge List meta-data to get a 'topic document' T_U for U
- Identify U 's topics from T_U
 - Basic IR techniques: case-folding, remove domain-specific stopwords
 - Extract nouns and adjectives using part-of-speech tagger
 - Topics for U : the extracted words along with their frequencies

Mining Lists to infer expertise

ashton kutcher 
@aplusk

Movies TV by Patty Holmes
59 members

Hollywood by mpc2000
1 members

Stars by Nadine Schultz
Meine Lieblingsstars
4 members

Entertainment by Al Royce
54 members

Celebrities by Ben Elcomb
142 members

Celebs by KING5 Photog Jim
126 members

Hollywood by Praesidian
17 members

- Collect Lists containing a given user U
- Merge List meta-data to get a 'topic document' T_U for U
- Identify U's topics from T_U
 - Basic IR techniques: case-folding, remove domain-specific stopwords
 - Extract nouns and adjectives using part-of-speech tagger
 - Topics for U: the extracted words along with their frequencies

Dataset

- Crawled the List-data for all users in our Twitter dataset, in November 2011
 - 1.3 million users are included in 10 or more Lists
 - Includes a large majority of the most popular users
 - Our studies focus on this set of users
-

Topics inferred from Lists



ChuckGrassley ✓
@ChuckGrassley Iowa
U.S. Senator born, raised and still living in New Hartford, IA.
<http://facebook.com/grassley>
<http://www.youtube.com/SenChuckGrassley>
<http://grassley.senate.gov>

politics, senator, congress, government,
republicans, iowa, gop, conservative



Claire McCaskill ✓
@clairecmc Missouri/ Washington DC
<http://twitter.com/clairecmc>

politics, senate, government, congress,
democrats, Missouri, progressive, women



The Linux Foundation
@linuxfoundation San Francisco, CA
A nonprofit consortium dedicated to fostering the growth of Linux.
<http://www.linux-foundation.org/>

linux, tech, open, software, libre, gnu,
computer, developer, ubuntu, unix

Evaluating the List-based methodology

- Are the inferred topics (i) accurate (ii) informative?
 - Evaluated using feedback through a user-survey
 - More than 93% evaluators judged the topics to be both accurate and informative
 - The few negative judgments were a result of subjectivity
-

Lists work better than other features



ashton kutcher ✓

@aplusk

I make stuff, actually I make up stuff, stories mostly, collaborations of thoughts, dreams, and actions. Thats me.

Los Angeles, California · <http://www.facebook.com/Ashton>

Profile bio

love, daily, people, time, GUI, movie,
video, life, happy, game, cool

Most common
words from tweets

celeb, actor, famous, movie, stars,
comedy, music, Hollywood, pop culture

Most common
words from Lists


Who-is-who service

- Developed a Who-is-Who service for Twitter
 - Shows word-cloud for major topics for a given user
 - <http://twitter-app.mpi-sws.org/who-is-who/>



ladamic : Lada Adamic

Associate Professor, School of Information (+ Complex Systems and EECS), University of Michigan



Search system for topic experts

- Given a query (topic)
 - Identify users related to the topic using Lists
 - Rank identified users



Ranking experts

- ❑ Used a ranking scheme solely based on Lists
- ❑ Two components of ranking user U w.r.t. query Q
 - ❑ Relevance of user to query – cover density ranking between topic document T_U of user and Q
 - ❑ Popularity of user – number of Lists including the user

Topic relevance(T_U , Q) \times log(#Lists including U)

Search system for topic experts

Cognos, a search system for topic experts

<http://twitter-app.mpi-sws.org/whom-to-follow/>

Location Filter : Worldwide

Search results for "politics"

**BarackObama** : Barack Obama ✓

This account is run by #Obama2012 campaign staff. Tweets from the President are signed -BO.

**whitehouse** : The White House ✓

Official WH twitter account. Comments & messages received through official WH pages are subject to the PRA and may be archived. Learn more <http://wh.gov/privacy>

**algore** : Al Gore ✓**maddow** : Rachel Maddow MSNBC ✓

I see political people...

**billmaher** : Bill Maher ✓

The Official Bill Maher Twitter

**nprpolitics** : NPR Politics

Political coverage and conversation from @NPRnews

Cognos
results for
"politics"


Location Filter :

Search results for "stem cell"



StemCellTracker : Stem Cell Tracker

The #1 resource for up-to-date stem cell news.



CellStemCell : Cell Stem Cell

A monthly journal at the forefront of stem cell science and policy. Editor: Dr. Deborah Sweet



StemCellNetwork : Stem Cell Network

SCN supports cutting-edge projects that translate Canadian stem cell research discoveries into new and better treatments.



iPSCellNews : Stem Cell (iPS) News

Latest updates about reprogramming & induced pluripotent stem cell research. -Jordan Kho, PhD student at Baylor College of Medicine-



ATStemCell : All Things Stem Cell

Blog discusses stem cells in multifaceted manner: history, apps, probs, news, & more. Run by PhD stem cell grad student/science writer- Need something written?

Cognos
results for
"stem cell"



User-evaluation of Cognos

Please rate the top 10 results for "science news"



bbcscitech : BBC SciTech



Set up by @mario, supported by backstage.bbc.co.uk

Relevant  



Reuters_Science : Reuters Science News



From newly charted astronomical anomalies at the far reaches of the universe to the rise of nanotechnology, nobody covers science like Reuters.com.

Relevant  



newscientist : New Scientist



New Scientist is the world's leading science and technology weekly

Relevant  



NatGeo : National Geographic


Since 1888, we've traveled the Earth, sharing its amazing stories with new generations. Official Twitter account of National Geographic.

Relevant  





science : science

Science news from ScienceNewsBlog.com.

Relevant  



NASA : NASA

Relevant  

Sample queries for evaluation

Enter your query

Search

Location Filter : Worldwide

Sample Queries

News : politics sports entertainment science technology business

Journalists : politics sports entertainment science technology business

Politics : conservative news liberal politicians USA / German / Brazilian / Indian politicians

Sports : F1 baseball soccer poker tennis NFL NBA Bundesliga LA Lakers

Entertainment : celebrities movie reviews theater music

Hobbies : hiking cooking chefs traveling photography

Lifestyle : wine dining book clubs health fashion

Science : biology astronomy computer science complex networks

Technology : iPhone mac linux cloud computing

Business : markets finance energy

Evaluation results

- ❑ Overall 2136 relevance judgments
 - ❑ 1680 said relevant (78.7%)
 - ❑ Large amount of subjectivity in evaluations
 - ❑ Same result for same query received both relevant and non-relevant judgments
 - ❑ E.g., for query “cloud computing”, Werner Vogels got 4 relevant judgments, 6 non-relevant judgments
-

Cognos vs. Twitter Who-To-Follow

Please compare the anonymized search results for "iphone"

A is better B is better Both are equally good Both are equally bad

Search Engine A



iPhone_News : iPhone News

iPhone news and notes from around the web.



AppStore : App Store

Follow us for official App Store tweets including our featured apps, exclusive offers, and more.



iphone_dev : iphone_dev

The iPhone Dev Team. We liberate your iPhones!



freeiphoneapps : Free iPhone Apps

We're currently giving away the new iPad! We also give away lots of cool paid iPhone, iPad, and Mac applications on a daily basis. FOLLOW US.



iphonehackx : iPhone Hacks

Unlock the potential of your iPhone, iPad and iPod Touch

Search Engine B



p0sixninja : Joshua Hill

iPhone Hacker



iH8sn0w : iH8sn0w

I'm the guy that made f0recast, iREB, iFaith, and sn0wbreeze. I also make #pie and dj on the side.



chronicdevteam : Chronic Dev

Hax



MuscleNerd : MuscleNerd

iPhone hacker



iPhone_News : iPhone News

iPhone news and notes from around the web.

Cognos vs. Twitter Who-To-Follow

- ❑ Considering 27 distinct queries asked at least twice
 - ❑ Judgment by majority voting

 - ❑ Cognos judged better on 12 queries
 - ❑ Computer science, Linux, Mac, Apple, Ipad, Internet, Windows phone, photography, political journalist, ...

 - ❑ Twitter Who-To-Follow judged better on 11 queries
 - ❑ Music, Sachin Tendulkar, Anjelina Jolie, Harry Potter, metallica, cloud computing, IIT Kharagpur, ...
-

Results for query music

De-anonymized search results for "music"

Search Engine A : Our results



katyperry : Katy Perry

i kissed a girl AND diddled her skittle.



ladygaga : Lady Gaga

mother mons+er



taylorswift13 : taylorswift13



jtimberlake : Justin Timberlake

Official Justin Timberlake Twitter.



Pink : P!nk

it`s all happening

Search Engine B : Twitter results



iTunesMusic : iTunes Music

Official music updates for the U.S. iTunes Store including new releases, pre-orders, iTunes LP, exclusive offers and more.



guardianmusic : Guardian music

Squashing music into 140 characters since 2008



yahoo_music : Yahoo! Music

The official Twitter account of Yahoo! Music. We tweet about music news, concerts, performances, videos, and all the things that make us yodel!



SonyMusicGlobal : Sony Music Global

The home of Sony Music on Twitter!



CountryMusic : Country Music Associ

Official Tweet of the Country Music Association (CMA) managed by @bennett49r & @chappedman.

Scalability problem

- ❑ Twitter now has around 500 million users
 - ❑ 740K new users join daily

 - ❑ How to keep the system up-to-date by discovering newly joining experts?

 - ❑ Twitter restricts crawling through API
 - ❑ Brute-force crawl of all users is infeasible
-

Solution

- ❑ Only 1.1% users are listed 10 or more times
 - ❑ If experts can be identified efficiently, possible to crawl their Lists
- ❑ Used hubs to identify authorities / experts
 - ❑ Hubs – users who selectively List many experts
 - ❑ Identify hubs using HITS, crawl Lists created by top hubs
 - ❑ 50% of users listed by top 2% hubs listed 10 or more times

Details in paper
