New Data Mining Opportunities of Social Media

Huan Liu

Joint Work with

Reza Zafarani

Geoffrey Barbier

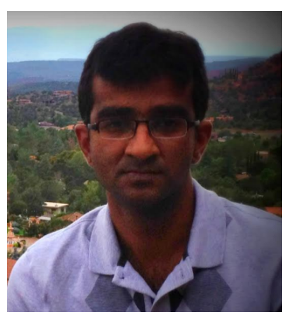

Pritam Gundecha

Xia Hu

Shamanth Kumar

Fred Morstatter

Social Media Mining by Cambridge University Press

Home | Book | Errata | Slides | Table of Contents | Tutorials

Social Media Mining

An Introduction

A Textbook by Cambridge University Press

Reza Zafarani	*Arizona State University*
Mohammad Ali Abbasi	*Arizona State University*
Huan Liu	*Arizona State University*

Download

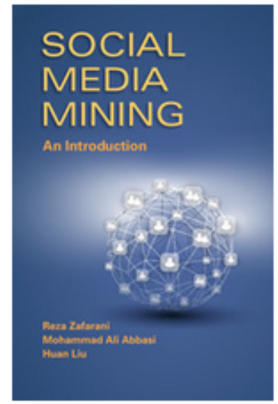

CAMBRIDGE UNIVERSITY PRESS

amazon.com

BARNES&NOBLE BOOKSELLERS

eBooks.com

The growth of social media over the last decade has revolutionized the way individuals interact and industries conduct business. Individuals produce data at an unprecedented rate by interacting, sharing, and consuming content through social media. Understanding and processing this new type of data to glean actionable patterns presents challenges and opportunities for interdisciplinary research, novel algorithms, and tool development. Social Media Mining integrates social media, social network analysis, and data mining to provide a convenient and coherent platform for students, practitioners, researchers, and project managers to understand the basics and potentials of social media mining. It introduces the unique problems arising from social media data and presents fundamental concepts, emerging issues, and effective algorithms for network analysis and data mining. Suitable for use in advanced undergraduate and beginning graduate courses as well as professional short courses, the text contains exercises of different degrees of difficulty that improve understanding and help apply concepts, principles, and methods in various scenarios of social media mining.

http://dmml.asu.edu/smm/

Social Media Data is Big

A Big-Data Paradox

- Collectively, social media data is indeed big
- For an individual, there is little data on a site
  - How much activity data do we generate daily?
  - How many posts did we post this week?
  - How many friends do we have?
- Often, we use different social media services for varied purposes
  - Facebook, Twitter, Instagram, YouTube, …
- “Big” social media data often may not be big
  - Searching for more data with limited data

An Example

Little data about an individual

Many social media sites

Partial Information

Complementary Information

Better User Profiles

Reza Zafarani

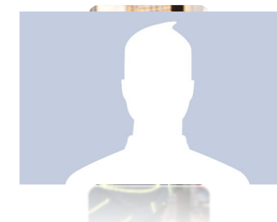

	LinkedIn	Twitter
Age	N/A	N/A
Location	Phoenix Area	Tempe, AZ
Education	ASU (2014)	ASU

Connectivity is not available

Consistency in Information Availability

Can we connect individuals across sites?

Searching for More Data with Limited Data

- Each social media site can have varied amount of user information
- What is guaranteed to exist for the joint set of these sites?
  - Usernames
- A user's usernames on different sites can be different
- We set out to verify that the information provided across sites belong to the same individual

Our Behavior Generates Information Redundancy

- Information shared across sites provides a behavioral fingerprint

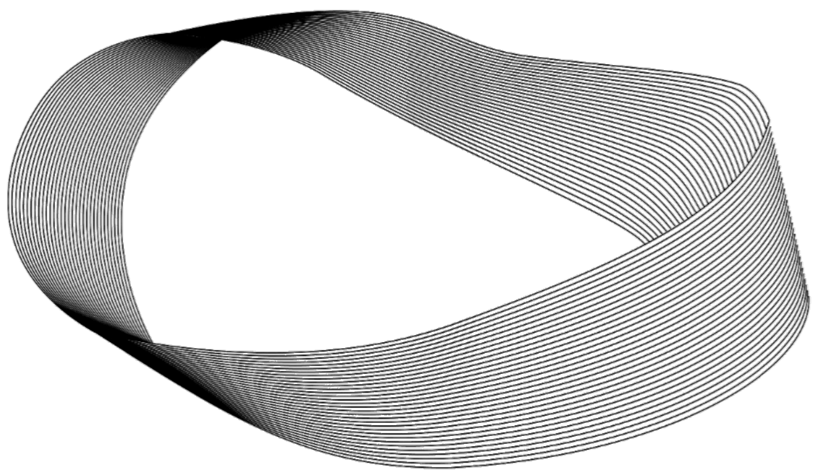

MOBIUS

- **Behavioral Modeling**
- **Minimum Information**

MOdeling **B**ehavior for **I**dentifying **U**sers across **S**ites

Starting with Minimum Information of a User

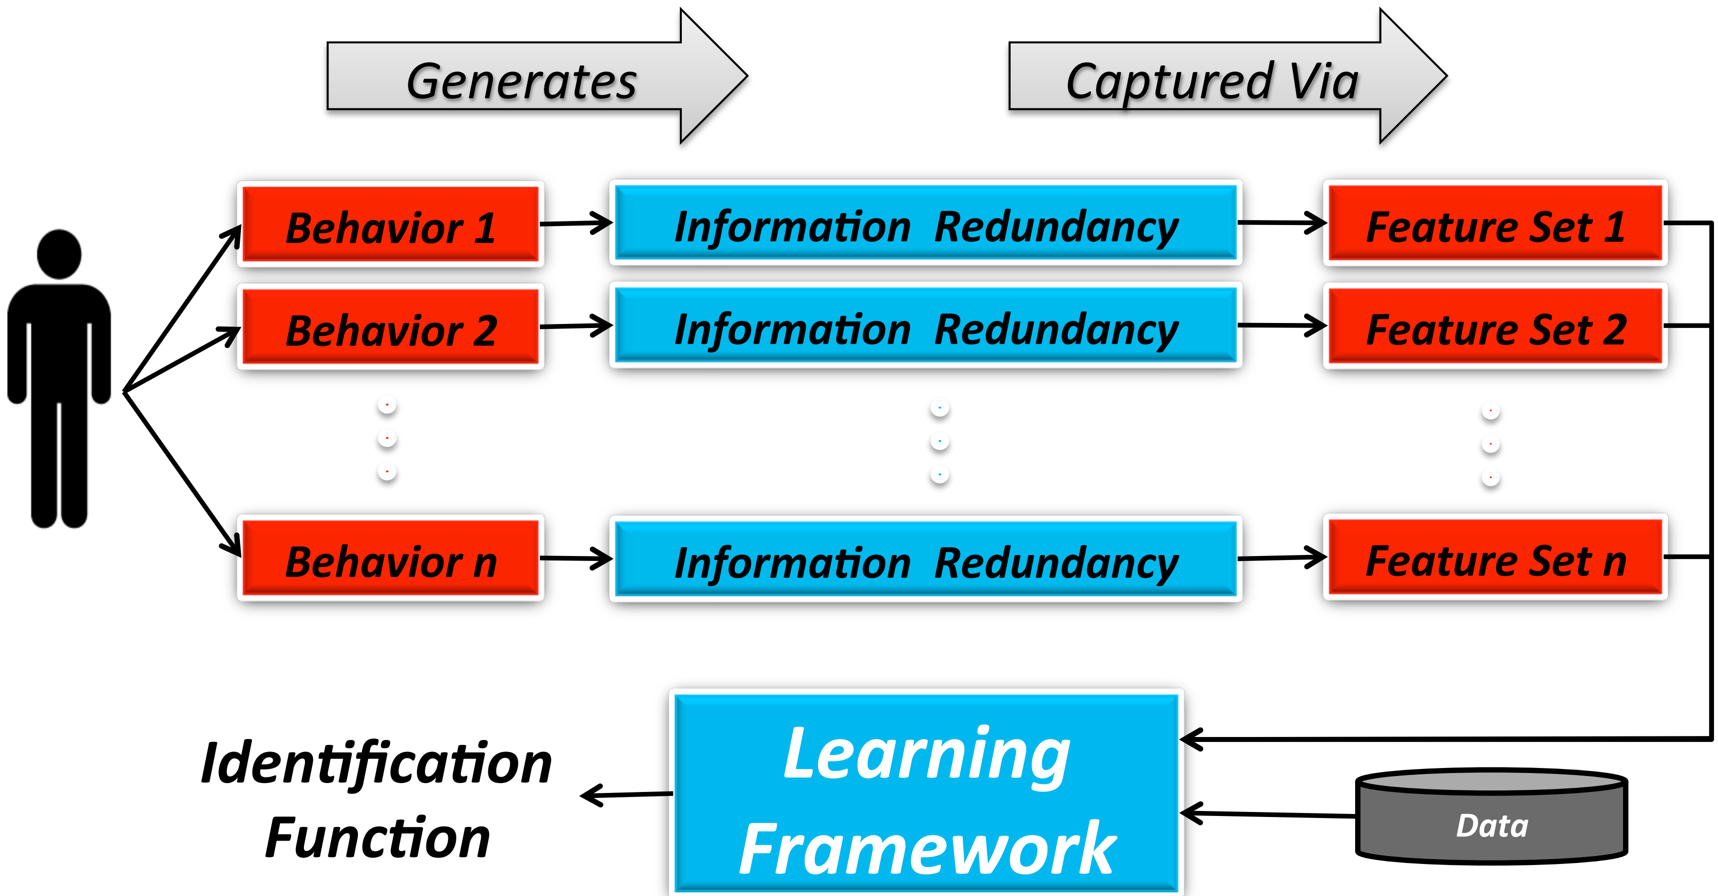

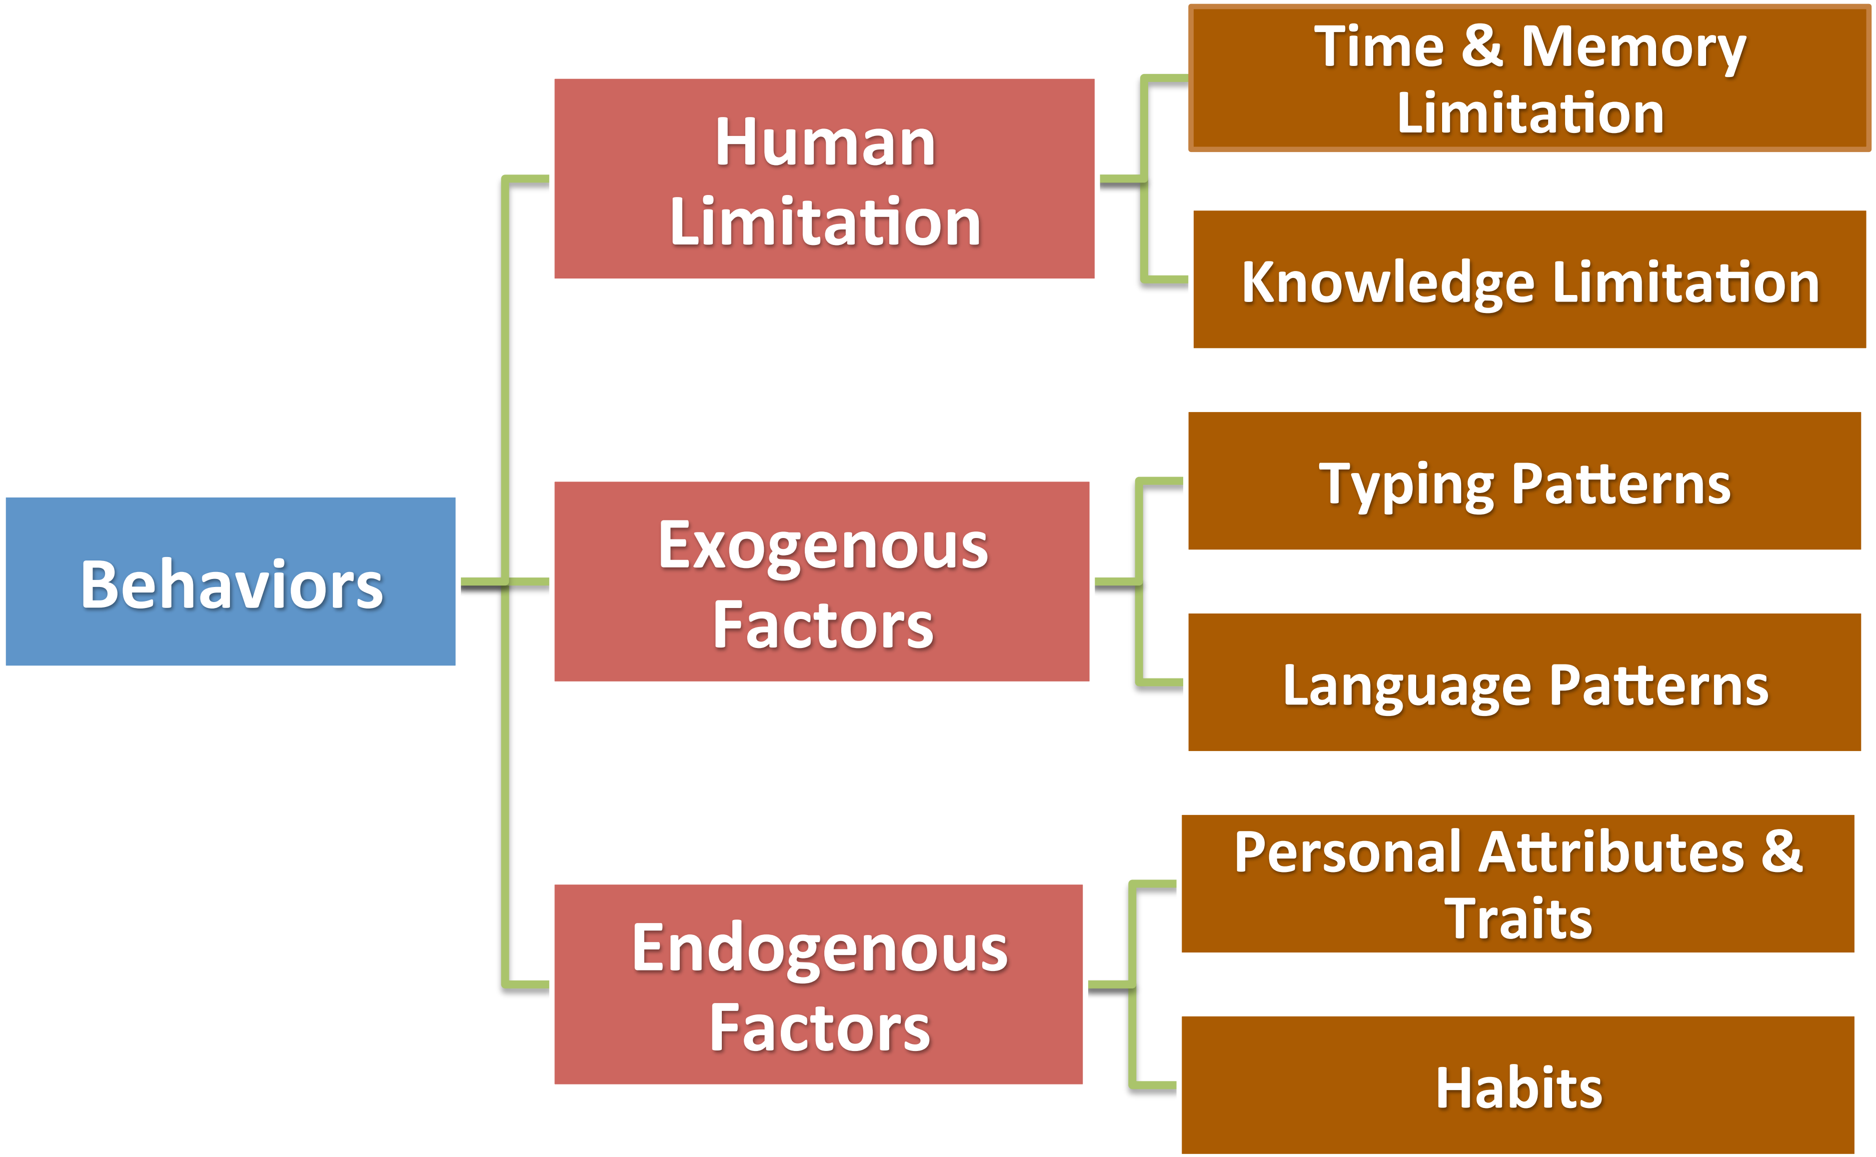
Behaviors
Human Limitation
Exogenous Factors
Endogenous Factors
Time & Memory Limitation
Knowledge Limitation
Typing Patterns
Language Patterns
Personal Attributes & Traits
Habits

Time and Memory Limitation

Using Same Usernames

59% of individuals use the same username

Username Length Likelihood

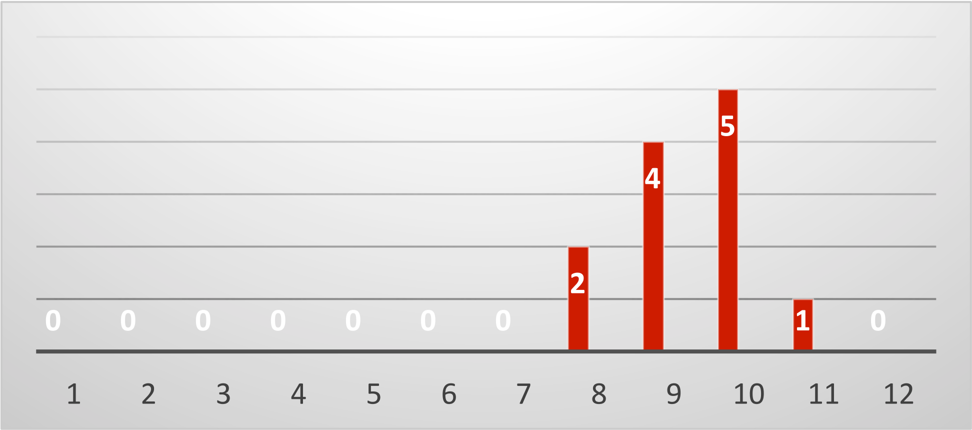

Knowledge Limitation

Limited Vocabulary

Identifying individuals by their vocabulary size

Limited Alphabet

Alphabet Size is correlated to language:

शमंत कुमार -> Shamanth Kumar

Typing Patterns

QWER1234

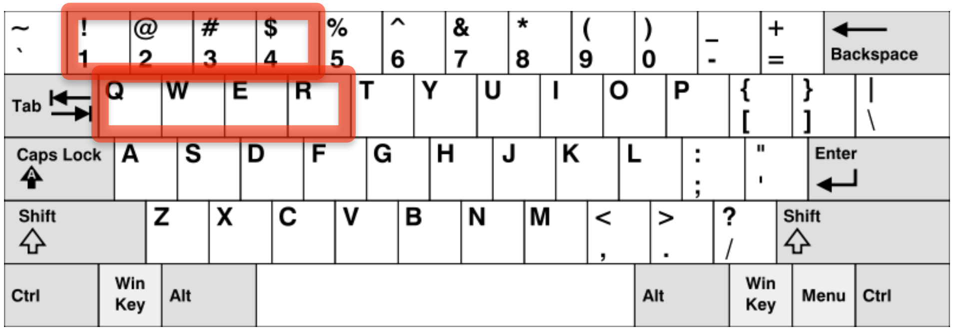

AOEUISNTH

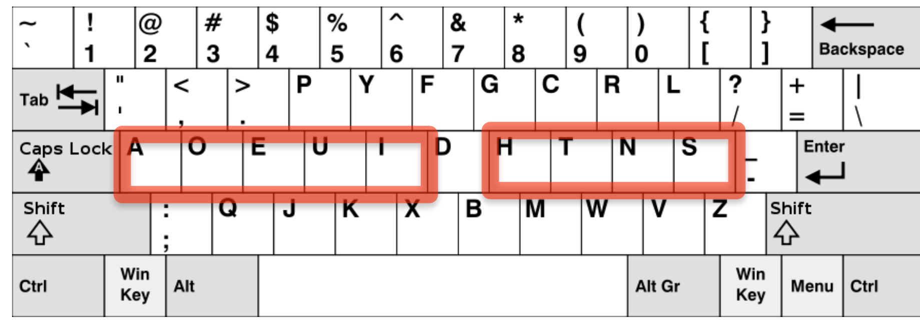

QWERTY Keyboard
Variants: AZERTY, QWERTZ

DVORAK Keyboard

Keyboard type impacts your usernames

**We compute features that capture typing patterns:
the distance you travel for typing the username,
the number of times you change hands when typing it, etc.**

Habits - old habits die hard

Modifying Previous Usernames

Adding Prefixes/Suffixes, Abbreviating, Swapping or Adding/ Removing Characters

Creating Similar Usernames

Nametag and Gateman

Username Observation Likelihood

Usernames come from a language model

Obtaining Features from Usernames

For each username:

414 Features

Similar Previous Methods:

1) Zafarani and Liu, 2009
2) Perito et al., 2011

Baselines:

1) Exact Username Match
2) Substring Match
3) Patterns in Letters

MOBIUS Performance

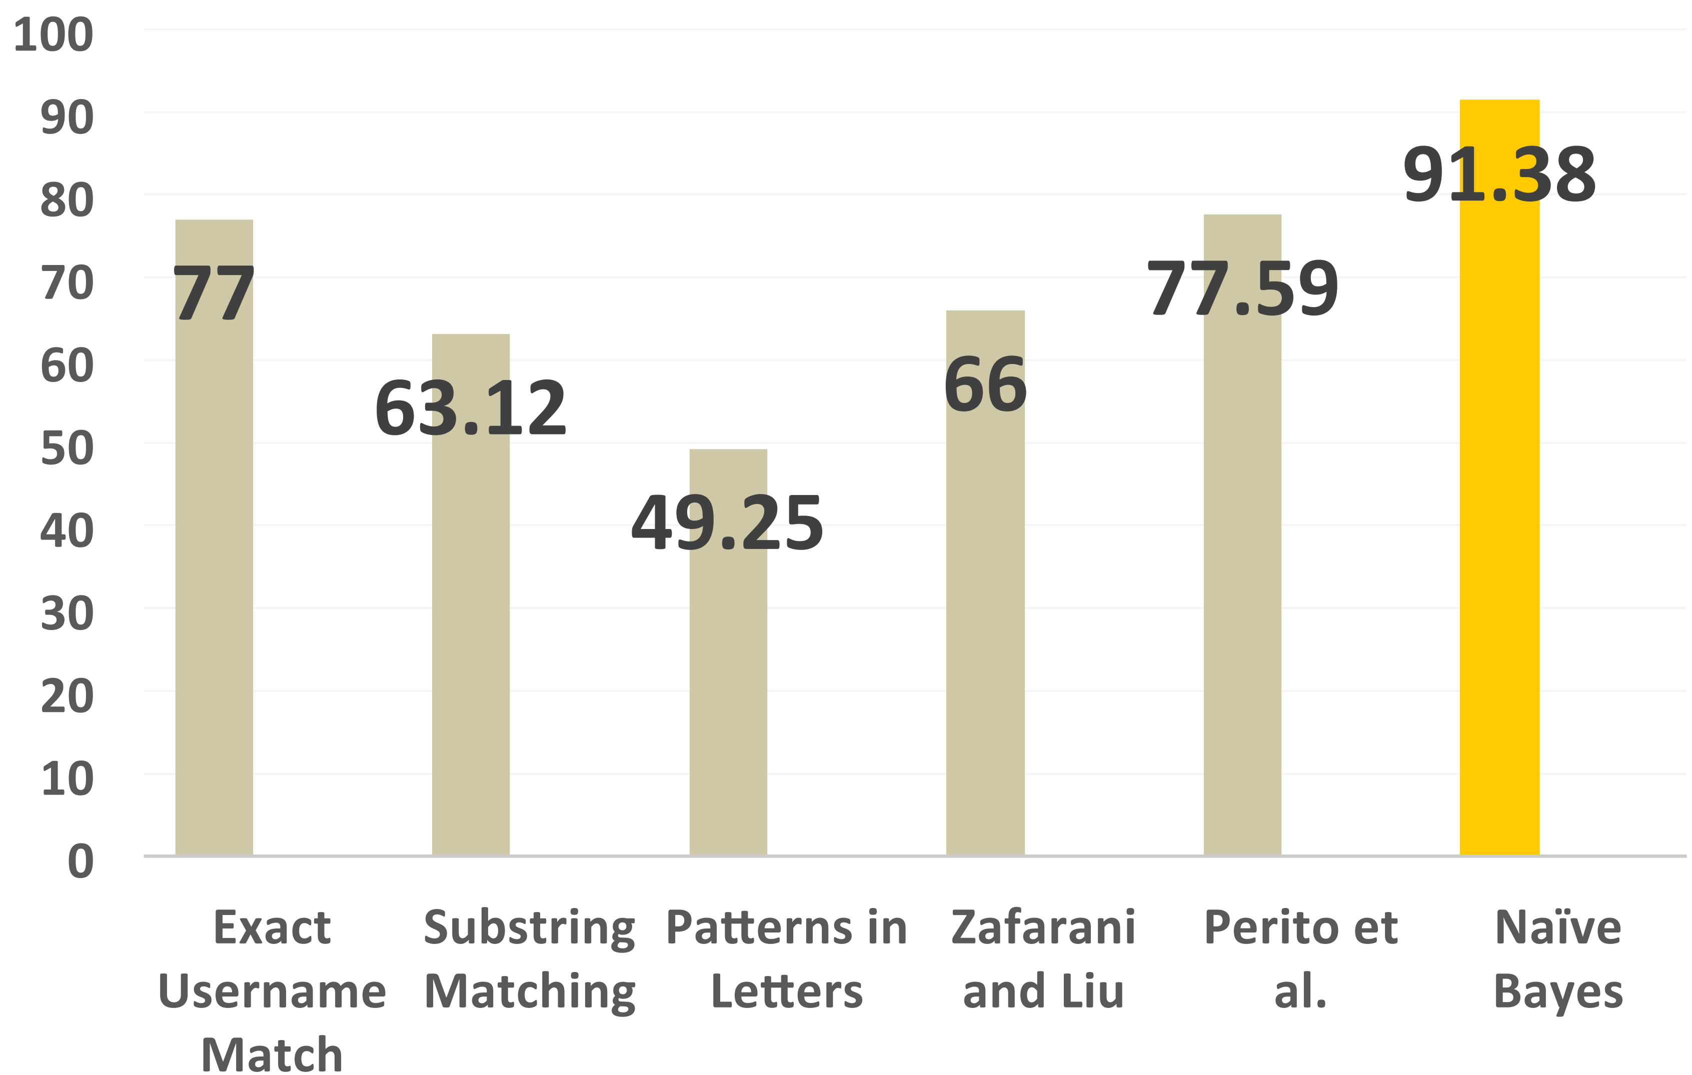

Choice of Learning Algorithm

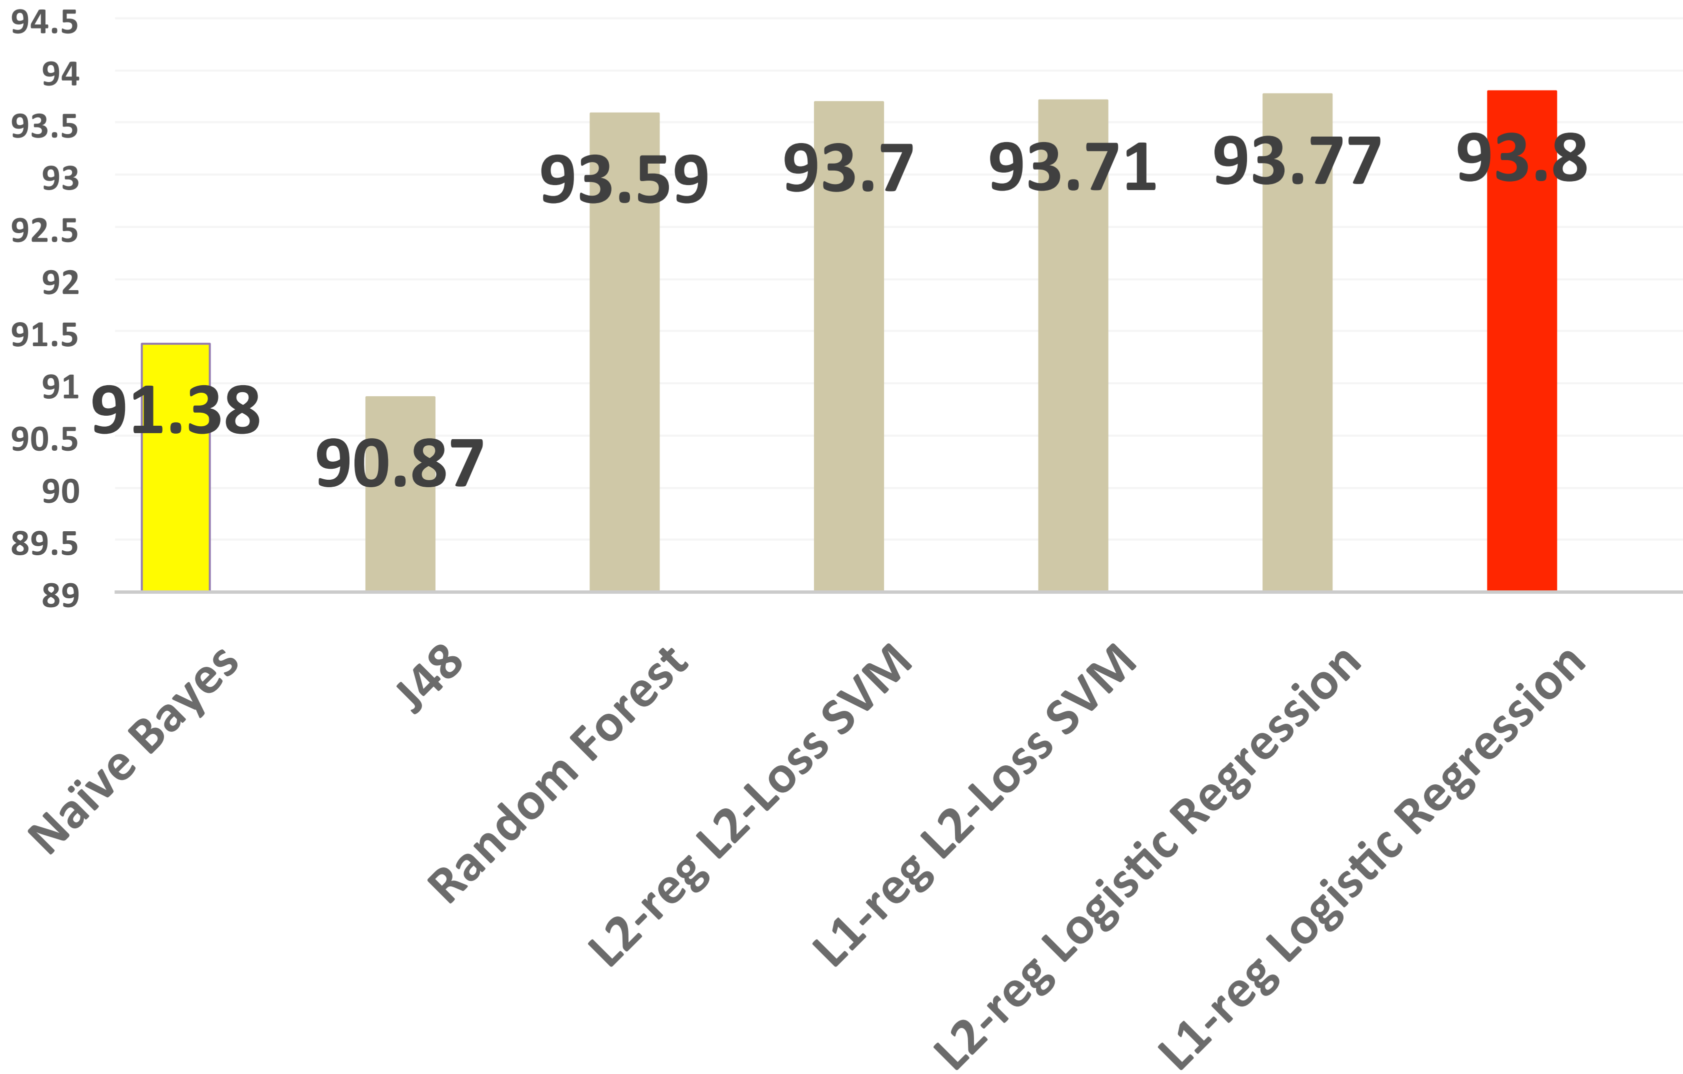

Summary

- Many a time, big data may not be sufficiently big for a data mining task
- Gathering more data is often necessary for effective data mining
- Social media data provides unique opportunities such as numerous sites and abundant user-generated content
- Traditionally available data can be equally tapped for making data “thicker”

Reza Zafarani and Huan Liu. ``Connecting Users across Social Media Sites: A Behavioral-Modeling Approach", SIGKDD, 2013.

Importance of Provenance Data in Social Media

- Social media shows its promise of producing positive and meaningful change in the world
  - Humanitarian assistance and disaster relief
  - Community outreach
  - User connection and information sharing
- As a neutral mechanism, its dark side should not be ignored
  - Promoting chaotic mass behavior
  - Escalation of rumors
  - Crowd manipulation and rumor propagation
  - Social media is a natural environment to social hacking
- Provenance can help uncover deceptive activities

Deception in Social Media

- Deception can cause social unrest
  - Incidents reported during Assam Exodus in which virulent messages and doctored photos ultimately lead to mass exodus of population from major cities in India
  - Deception was widely seen during the recent protests in Bangladesh
- Deception can also cause emotional distress to an individual or a group

Types of Deception

- Based on operations on Information with an intention of deception, we have
  - Manipulation: Modifying existing information
  - Change-of-context: Presenting correct information with incorrect context
  - Fabrication: Creating a new story from multiple pieces of information
  - Or some combinations of the above

Example (1)

- A gory video clip of riots in Indonesia was posted on YouTube, labeled as a video of the Assam riots.
- Many Twitter users, including some influential people, were retweeting, or spreading the misinformation.
- The influentials include Indian politicians, journalists, TV news reporters, and social activists
- Deception Type: Change-of-context

Example (2)

http://breakingbrown.com/2013/06/court-allows-shirley-sherrod-to-continue-with-defamation-case-against-right-wing-blogger/

- Resignation of Shirley Sherrod (July, 2010), where Blogger→ Out-of-context Video → Resignation
- It caused embarrassment to the government
- Deception Type: Change-of-context

Example (3)

- False link information

Fabricated news: Girl raped at generation square in Shahbag

No news available at the link provided

- Deception Type: Fabrication

Example (4)

- Image manipulation during recent Bangladesh protest

Fabricated news: Chhatro League activists raped a police woman

Actual news: A police woman fainted while she was on duty during a strike on July 7, 2011

- Deception Type: Manipulation

Example (5)

- Image manipulation during recent Bangladesh protest

Fabricated news: Young man died while acting of being hanged

Actual news: Young man acted symbolic hanging of war criminals

- Deception Type: Manipulation

Example (6)

- False news report claims

Fabricated news: Clerics of Kaaba formed human wall to protest atrocities committed against Islamic scholars

Actual news: Covering the Kaaba

- Deception Type: Change-of-context

Example (7)

- Fake images during Hurricane Sandy

- Power outage in North-eastern America.
- Originally from Aug 14,2003
- **Deception Type: Change-of-context**

- The Statue of Liberty
- Originally from a Hollywood movie
- **Deception Type: Fabrication**

- Sharks in the subway of New Jersey
- Originally from the collapse of a shark tank at The Scientific Center in Kuwait

  Deception Type: Manipulation

Example (8)

- Fake image to exploit religious sentiment
  - Face appearing on the moon surface

- Deception Type: Manipulation

Example (9)

- Samsung paid off a $1.05 billion judgment awarded to Apple in a patent infringement lawsuit entirely in nickels.

- Deception Type: Change-of-context

Categorization: Type vs. Impact

	Neutral	Mischievous	Disastrous
Manipulation	7c	8	4,5
Change-of-Context	7a, 9	1	2,6
Fabrication	7b		3

1	Assam Exodus – YouTube Video
2	Shirley Sherrod Resignation
3	Fake Rape Story at Shahbag (BD)
4	Fake Police Woman Story (BD)
5	Fake Hanging Story (BD)

6	Clerics of Kaba (BD)
7a	Power Outage in North-America (HS)
7b	The Statue of Liberty (HS)
7c	Sharks in the Subway Station (HS)
8	Moon Incident (BD)
9	Samsung Apple Humor Story

Abstraction with Real Examples

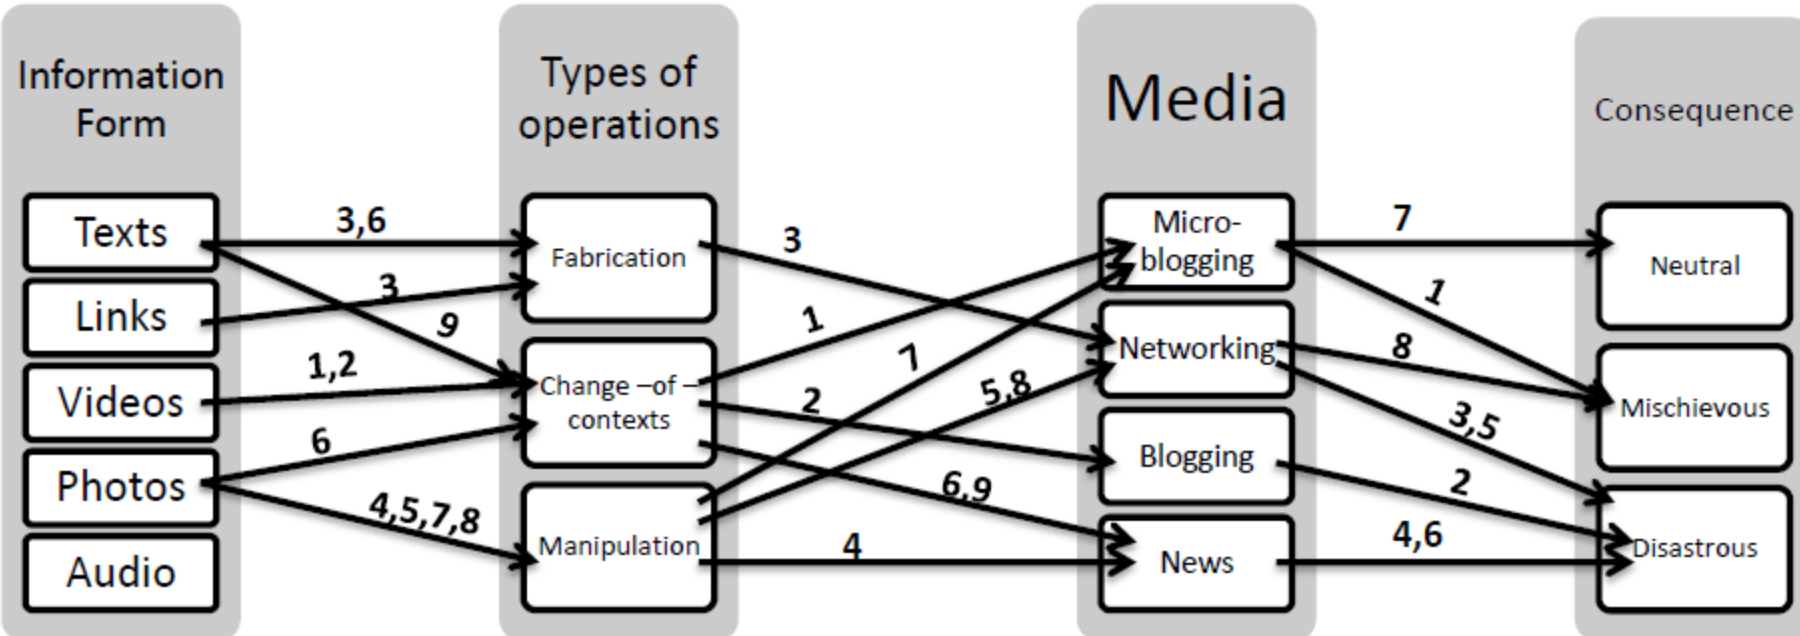

1	Assam Exodus – YouTube Video
2	Shirley Sherrod Resignation
3	Fake Rape Story at Shahbag (BD)
4	Fake Police Woman Story (BD)
5	Fake Hanging Story (BD)

6	Clerics of Kaba (BD)
7	Sharks in the Subway Station (HS)
8	Moon Incident (BD)
9	Samsung Apple Humor Story

Neutralizing Negative Aspects

- Recent events show that social media is used to mobilize crowds in both positive and negative ways
  - Arab Spring, Assam Exodus, and Bangladesh protests
- Users need to develop a healthy skepticism about information they receive, learn to check sources, and refine their skills of discernment
- A crowd able to carefully self-police itself is the absolute best defense – easier said than done …
- New technologies are needed to facilitate self-policing of social media and to discovering scams, hoaxes, and exploitations

Challenges

- Types of information – text, link, audio, photo, video, and mixture
- Types of operations – fabrication, change-of-context, and manipulation
- Number of sources
- Types of social media – micro-blogging, networking, blogging, and news
- Consequences – neutral, mischievous and disastrous
- Incomplete information – partial, noisy, …

Deception from a Social Science Perspective

- Relying on a physical observation
- Common cues for Deception Detection
  - Body language
  - Emotional gestures and contradiction
  - Interactions and reactions
  - Verbal contexts
  - Information (passively observable)
    - Fewer first-person pronouns
    - More negative emotion words (such as hate, worthless and sad)
    - Connection
- Characteristics of social media can be used to generate cues for deception detection in social media

Not Available in Social Media

Useful in Social Media

Computational Tasks

- Provenance Search in Social Media
  - Incomplete information, and
  - Number of sources
- Deception Detection with Provenance Data
  - Multimedia
  - Different types of operations
  - Various types of social media
- Impact Assessment of Suspicious Messages
  - Not all messages generate similarly critical consequences

Information Provenance

- It informs ownership, sources, or origin
- Social media often lacks provenance
- Users can't easily verify provenance
  - Reasons include big data, storage management, and unavailability of provenance records
- Provenance can inform additional value and trust assessment of the information
- It is challenging to search for provenance of information

Traditional Information Provenance

- The problem is well understood in databases, workflows, e-sciences and semantic webs
- Their focus is on organizing data storage to help finding provenance later
- In social media, users generate content and make connections without a centralized control
- Hence, identifying information provenance in social media is a new challenge

Challenges in Information Provenance

- Social media users rarely publish provenance data (why do we bother?)
- Information is spread in a dynamic network, resulting in a large search space
- Available information is incomplete and noisy
- Information propagates over multiple social media sites

Summary: Some Key Research Challenges

- Provenance Attributes
  - Are these attributes useful for identifying the provenance of information?
- Network Information
  - How can we seek the provenance of information using the network information alone?
- Using Both Information
  - How can content, attributes, and networks be helpful for identifying the provenance of information?

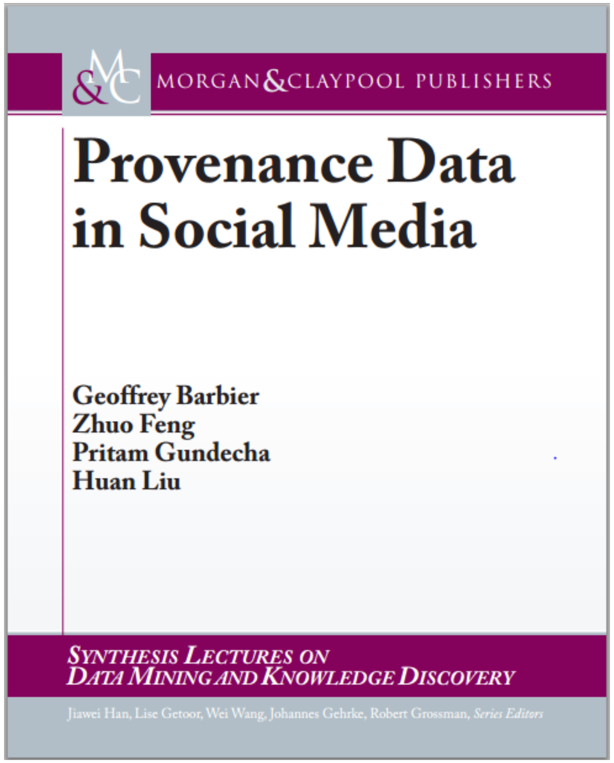

Pritam Gundecha, Feng Zhuo, and Huan Liu. ``Seeking Provenance of Information in Social Media", CIKM, 2013.

A New Type of Spammers in Social Media

- Social spammers send out unwanted spam content appearing on social networks and any website with user-generated content to targeted users, often corroborating to boost their legitimacy, credibility, and social influence.
- Spam content “can be manifested in many ways, including bulk messages, profanity, insults, hate speech, malicious links, fraudulent reviews, fake friends, and personally identifiable information” -- Wikipedia

Examples from Twitter

"Spam describes a variety of prohibited behaviors that violate the Twitter Rules."

According to Twitter, some common tactics are:

- **Posting harmful links** (including links to phishing or malware sites)
- **Aggressive following behavior** (mass following and mass un-following for attention)
- **Abusing the @reply or @mention function** to post unwanted messages to users
- **Creating multiple accounts** (either manually or using automated tools)
- **Posting repeatedly to trending topics** to try to grab attention
- **Repeatedly posting duplicate updates**
- **Posting links with unrelated tweets**

https://support.twitter.com/articles/64986-reporting-spam-on-twitter

Spamming on Twitter – An Example

Twitter spam bot replies to offer prizes related to events such as NFL or Miley Cyrus

http://www.symantec.com/connect/blogs/twitter-spam-bots-target-nfl-and-miley-cyrus-fans

Characteristics of Social Spammers

Content information:

- Short text
- Unconventional use of language
- Texts in a thread
- Adaptive to specific events

Characteristics of Social Spammers

Social network information:

- It is easier to establish an arbitrarily large number of social trust relations via Twitter follower markets

Rank	Provider	Highlights	Quick Review	Pricing	Links
#1	Devumi Social Media Marketing www.devumi.com	✓ Up to 300k Twitter Followers ✓ 100% Money-Back Guarantee ✓ Replacement Guarantee ✓ No Password Required ✓ Excellent Customer Support ✓ Saw Results in 24 Hours	Test Account: @DevumiReview Our Experience ✓ Completed within 24 Hours ✓ Over-Delivered by 30% ✓ Quick & Friendly Customer Support ✓ Offers Sponsored Mentions ✓ No Losses in Followers	From $12 **Accepts:** Credit Card, PayPal & 2Checkout	Visit Site Read Review
#2	FastFollowerz www.fastfollowerz.com	✓ Up to 1M Twitter Followers ✓ 100% Money-Back Guarantee ✓ Replacement Guarantee ✓ No Password Required ✓ Good Customer Support ✓ Saw Results in 48 Hours	Test Account: @FFzReview Our Experience ✓ Completed within 48 Hours ✓ Over-Delivered by 7% ✓ Quick & Friendly Customer Support ✖ Does **not** deliver Active followers. ✖ Unstable (Loss Followers Quickly)	From $16 **Accepts:** 2Checkout	Visit Site Read Review
#3	Twitter Boost Followers, Retweets & @Mentions www.twitterboost.co	✓ Up to 100k Twitter Followers ✓ 100% Money-Back Guarantee ✓ Replacement Guarantee ✓ No Password Required — Decent Customer Support ✓ Saw Results in 1 days	Test Account: @TwBoostReview Our Experience ✓ Completed within 1 Days ✓ Over-Delivered by 20% ✓ Daily Retweets Service ✓ No Drops in Followers yet.	From $7 **Accepts:** Credit Card, PayPal & Bitcoin	Visit Site Read Review

Stringhini, Gianluca, et al. "Follow the green: growth and dynamics in twitter follower markets." *IMC*, 2013.

Characteristics of Social Spammers

Social network information:

- Collaborative link farming [WWW 2012] widely exists on Twitter: spammers try to infiltrate the Twitter network by building social relationships with normal users and spammers themselves
- In social media, many users simply follow back when they are followed by someone for the sake of courtesy -- ***reflexive reciprocity*** [IJCAI 2013]

Ghosh, Saptarshi, et al. "Understanding and combating link farming in the twitter social network." *WWW*, 2012.
Xia Hu, Jiliang Tang, Yanchao Zhang, and Huan Liu. "Social spammer detection in microblogging." *IJCAI*, 2013.

Combating Social Spammers for Users

In a world without social spammers, from users' perspective,

- Information on social media services will be easier accessible and more interesting
- Social media will be less prone to cyber-attacks when acquiring useful information

Combating Social Spammers for Companies

Spam can inflict damages to companies:

- Spammers on social media can easily damage a brand and turn fans and followers into foes

- When advertisements of products from a company are mixed with spam information, it can have a profoundly negative impact on your social media marketing return on investment ROI

How to collectively make use of content and relations for social spammer detection?

Xia Hu, Jiliang Tang, Yanchao Zhang, and Huan Liu. "Social spammer detection in microblogging." *IJCAI*, 2013.

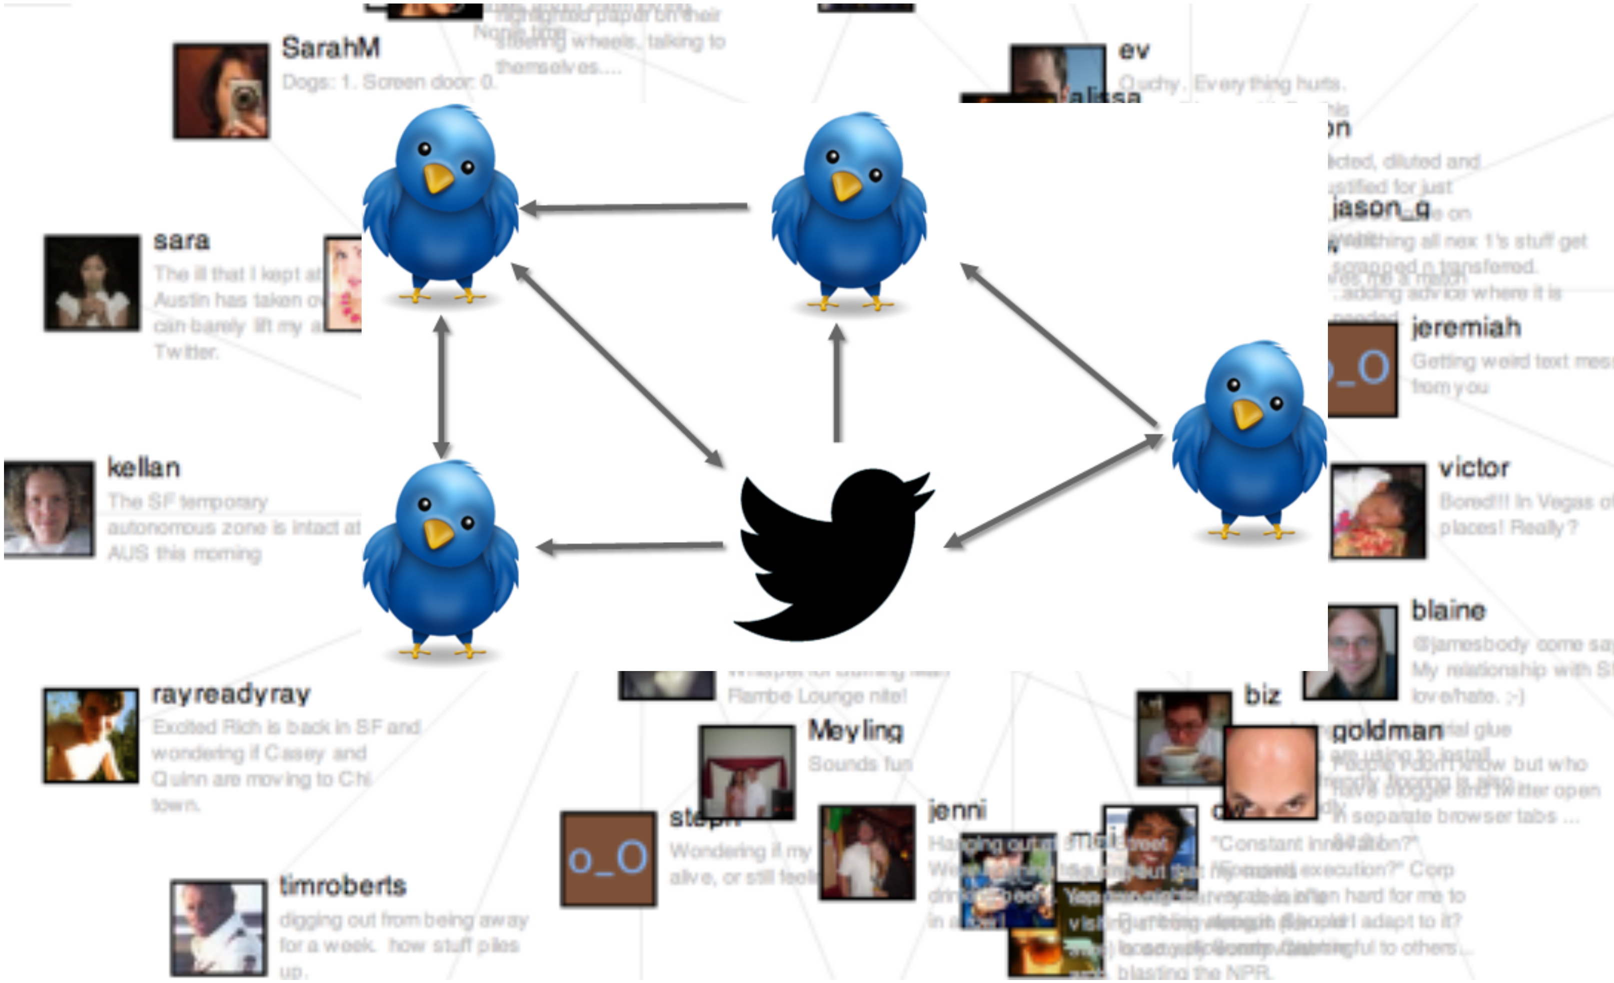

How to effectively collect labeling data?

Xia Hu, Jiliang Tang, Huiji Gao, and Huan Liu. "ActNeT: Active Learning for Networked Texts in Microblogging." *SDM*, 2013.

"Spam describes a variety of prohibited behaviors that violate the Twitter Rules. *Behaviors that constitute spamming will continue to evolve* as we respond to new tactics by spammers." -- Twitter

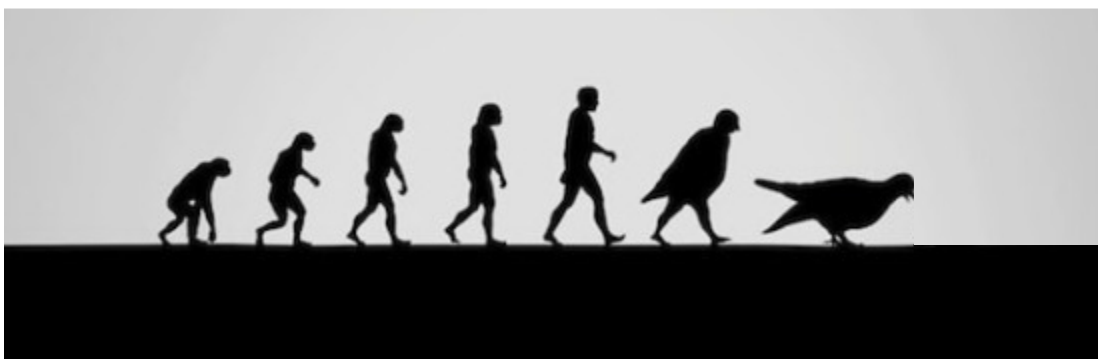

How to effectively handle evolving spammers/patterns in a built model?

Xia Hu, Jiliang Tang, and Huan Liu. "Online Social Spammer Detection." *AAAI*, 2014.

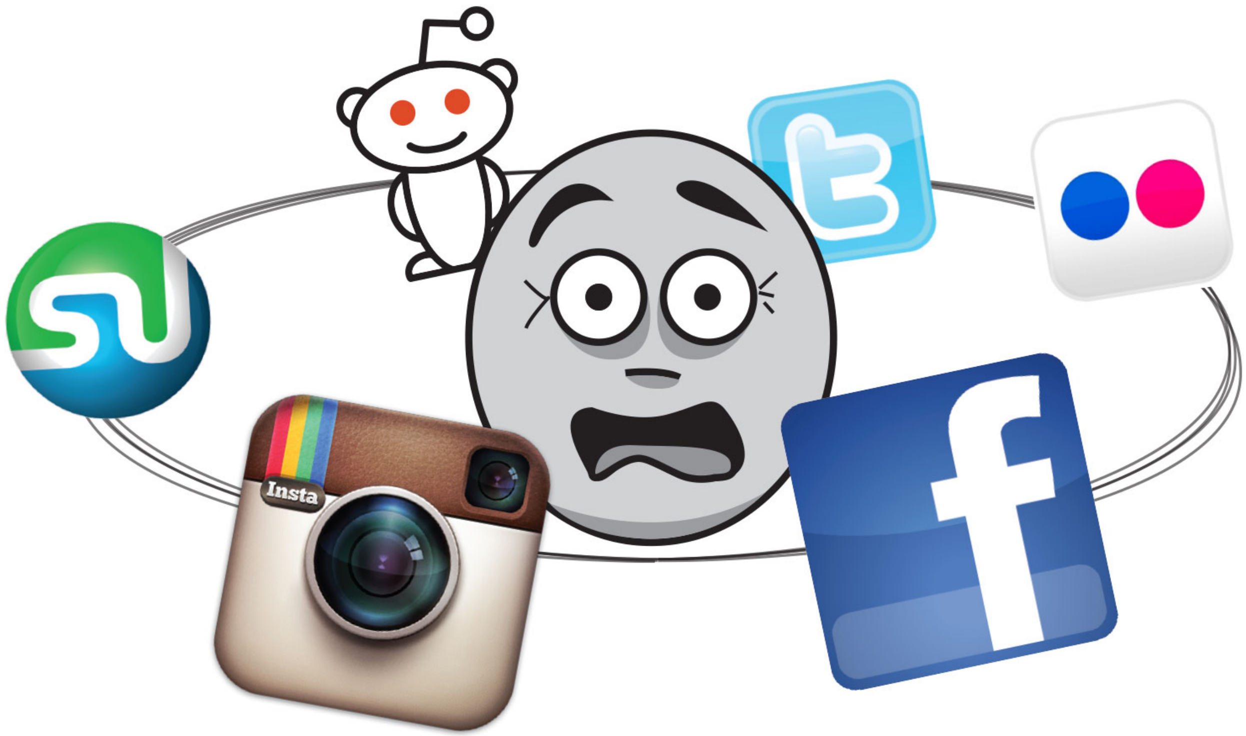

Is cross-media information helpful?

How to make use of cross-media resources?

Xia Hu, Jiliang Tang, and Huan Liu. "Leveraging Knowledge across Media for Spammer Detection in Microblogging." SIGIR, 2014

Summary

- Social spammers are a new type of spammers who take advantage of properties of social media
- They can wreak havoc on business and individuals
- It is challenging to detect social spammers without affecting normal users

A tutorial at PAKDD2014 in May on **Mining Social Spammers in Social Media**

http://www.public.asu.edu/~xiahu/tutorials/tut-pakdd14.htm

Empowering Humanitarian Assistance & Disaster Relief

- Twitter is one of major social media platforms
- Used in times of disaster.
- Recent events:
  - Arab Spring
  - Hurricane Sandy
  - Boston Bombing

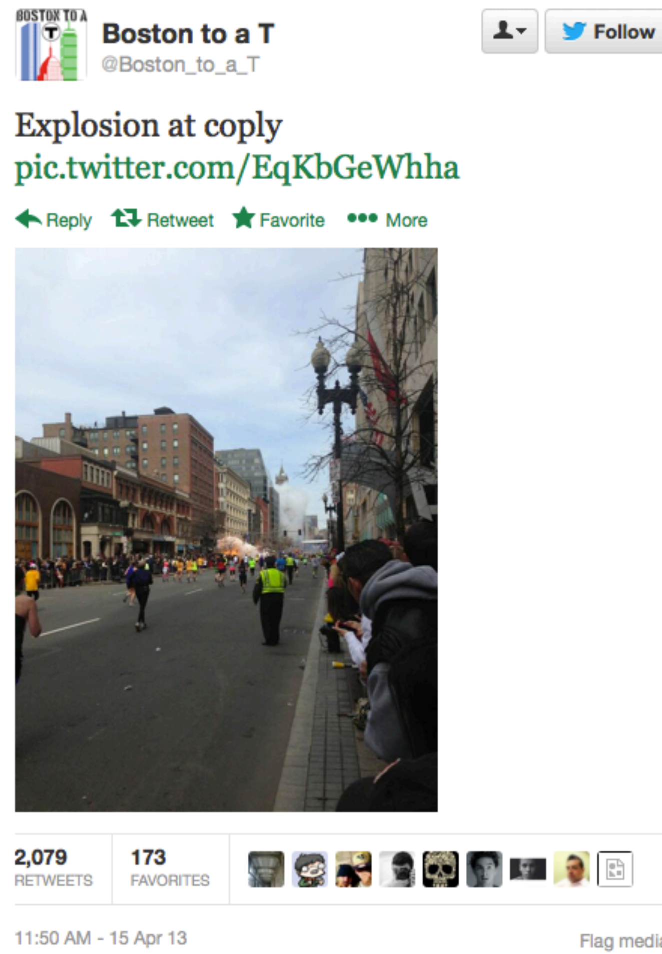

TweetTracker

- Collects crisis-related tweets
- Highlights central keywords, URLs, and users
- Aggregates and visualizes Twitter data

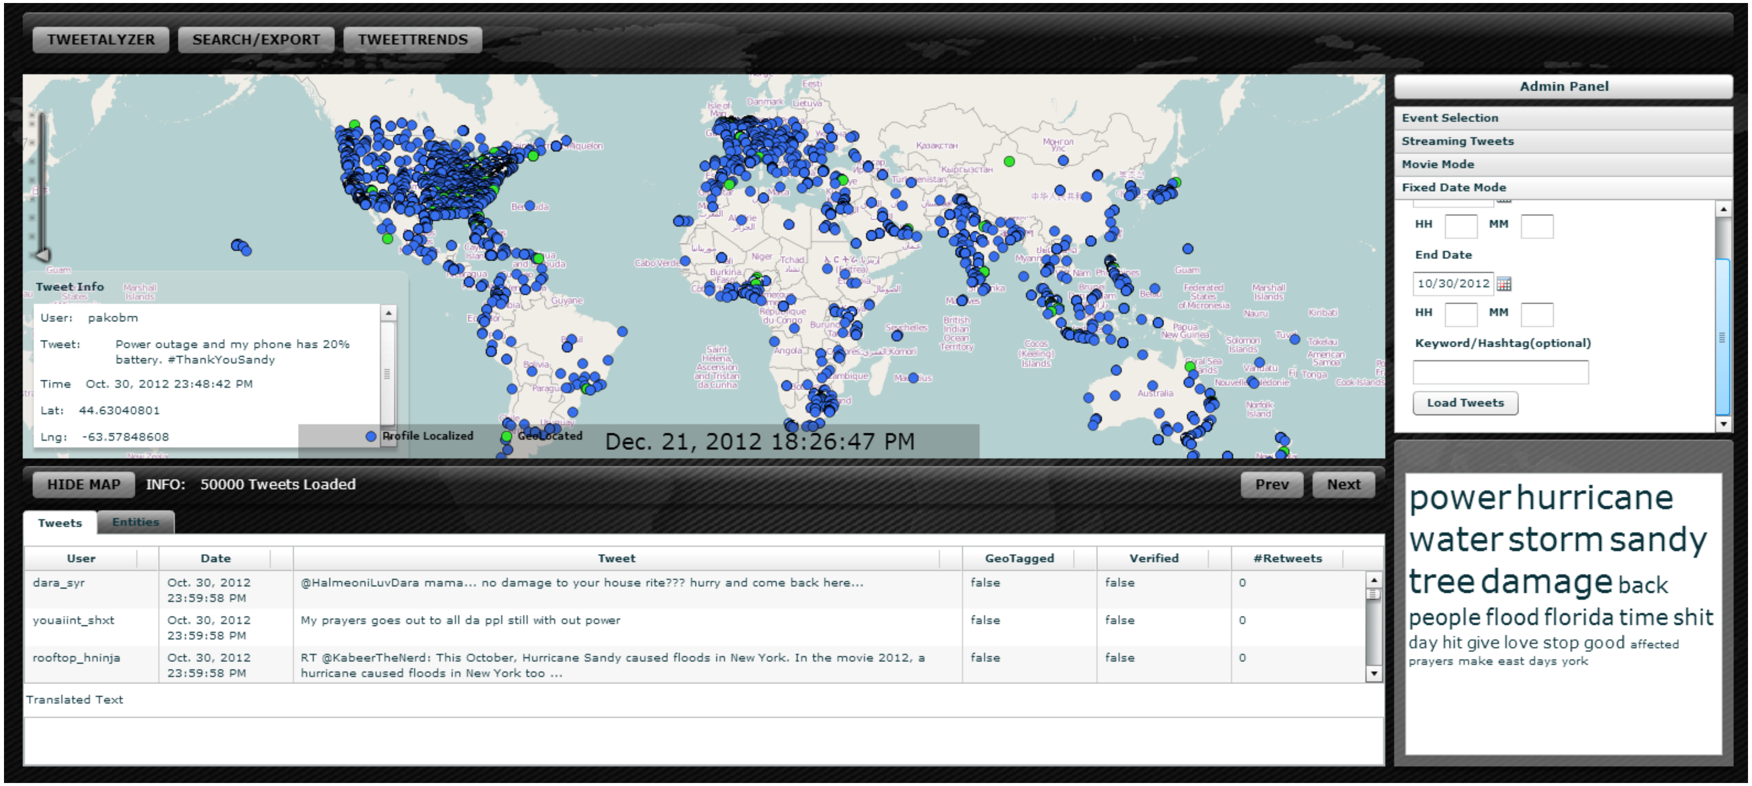

TweetXplorer

- Highlights important users and locations.
- Retweet network reveals key users.
- Map shows prominent locations.

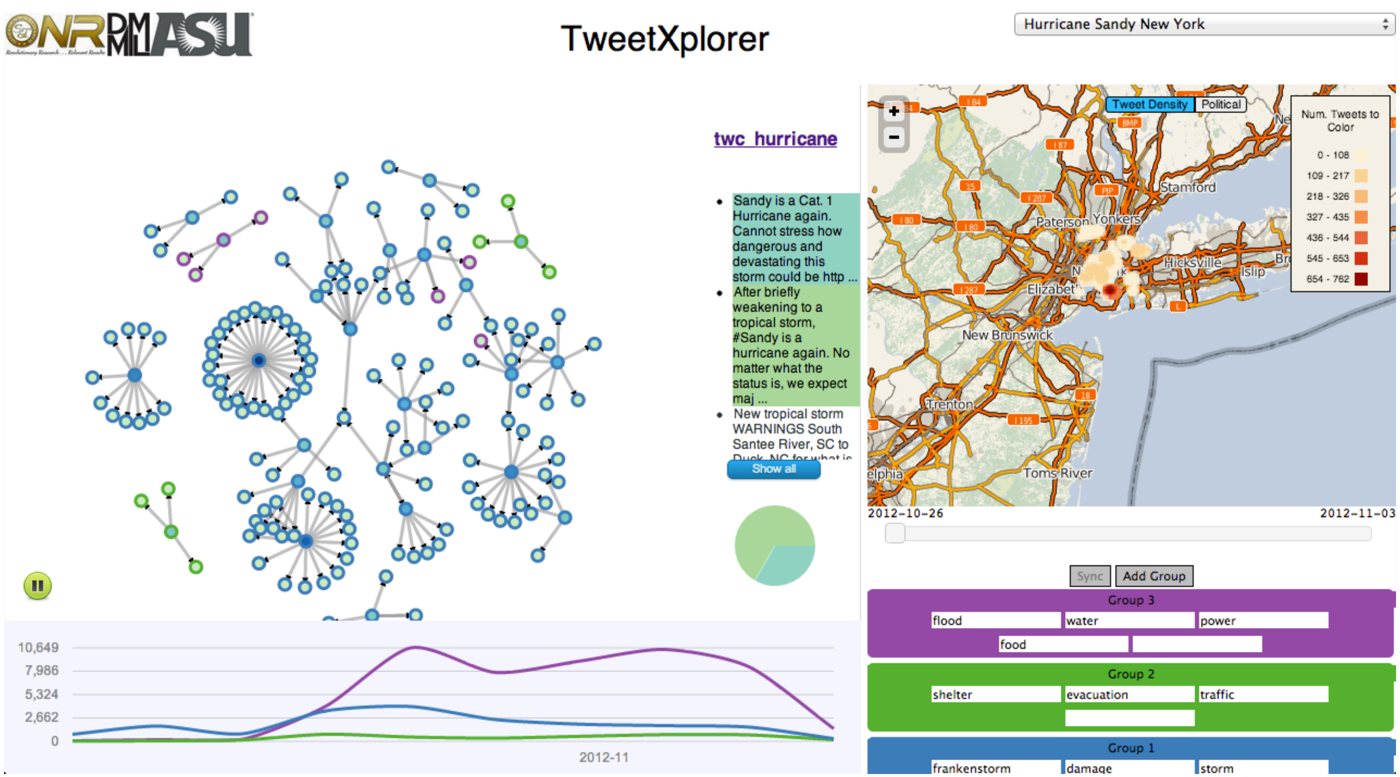

Challenges lead to New Opportunities

- Indexed over 2 billion crisis-related tweets
- Used by Humanity Road, and government organizations
- Licensed to 6 companies
- Awarded 2014 President's Team Award for Innovation at ASU

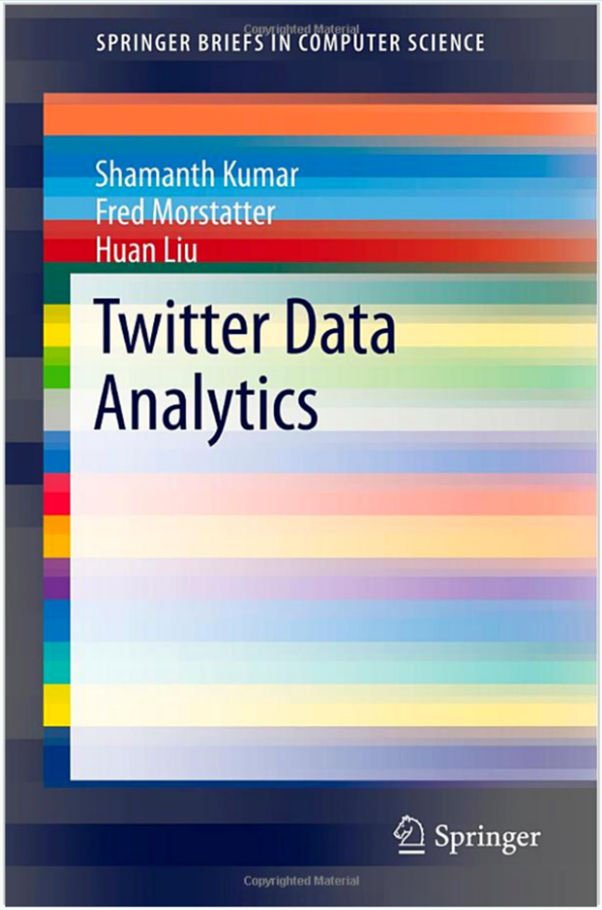

http://tweettracker.fulton.asu.edu/tda/

Twitter Data Analytics

- Leads reader through common tasks of mining Big Twitter Data
  - Collection
  - Storage
  - Analysis
  - Visualization

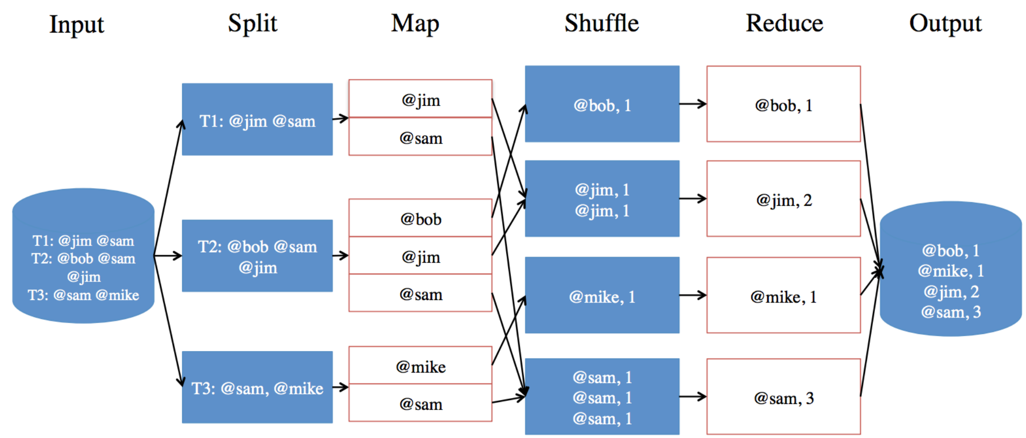

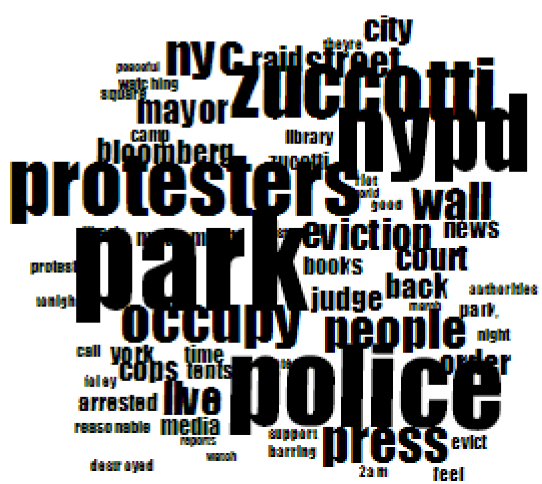

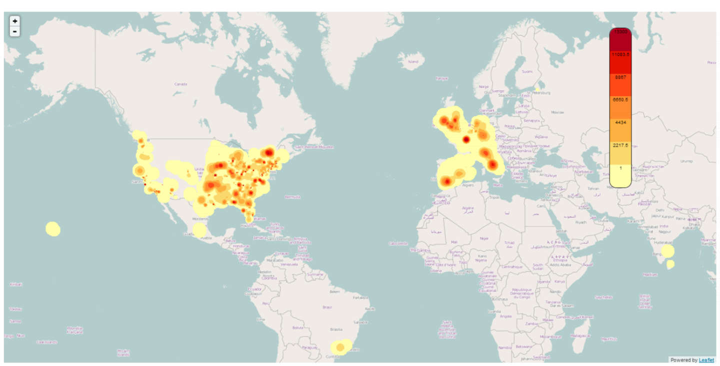

Concluding Remarks

- A Big-Data Paradox
- The Need for Provenance Data
- Detecting Social Spammers
- Empowering Humanitarian Assistance and Disaster Relief

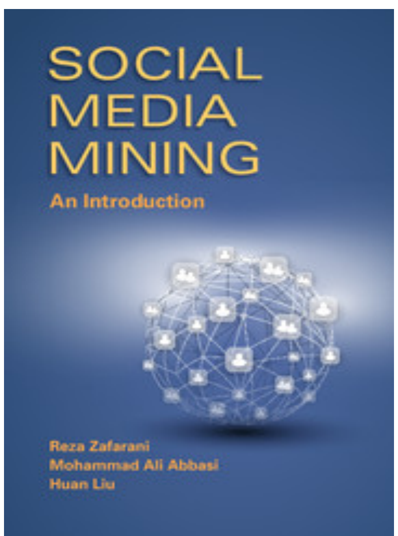

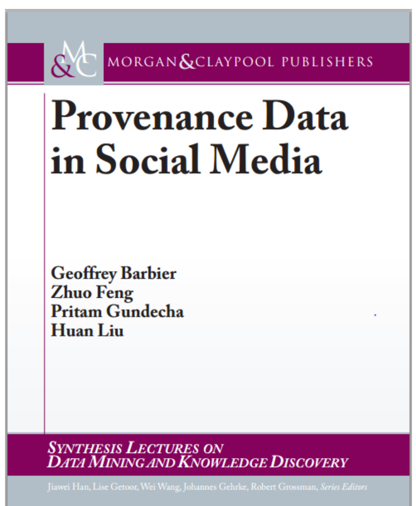

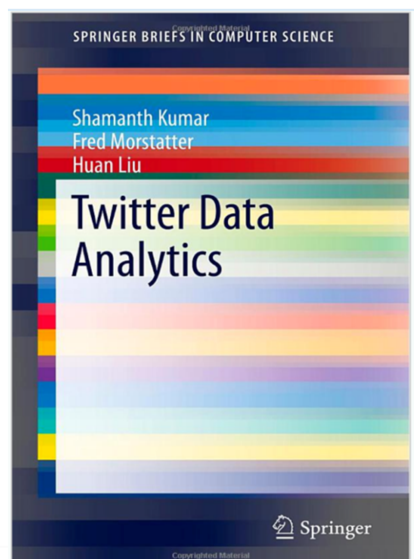

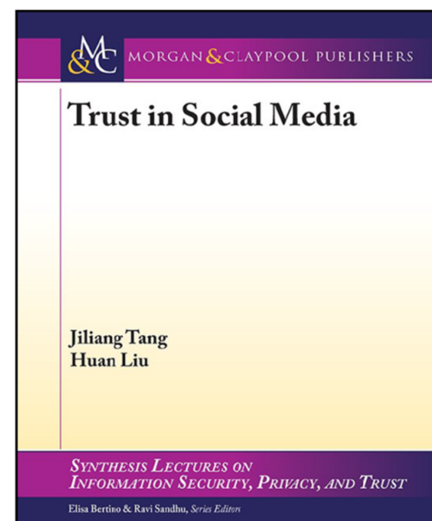

THANKS …

- Organizers for this wonderful opportunity to share our research work
- Acknowledgments
  - Grants from NSF, ONR, ARO
  - DMML members and project leaders
  - Collaborators


Reza Zafarani


Geoffrey Barbier

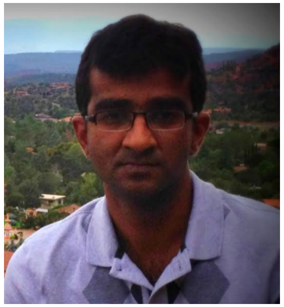
Pritam Gundecha


Xia Hu


Shamanth Kumar


Fred Morstatter