



Malicious Behavior on the Web: Characterization and Detection

Srijan Kumar (@srijankr)

Justin Cheng (@jcccf)

Jure Leskovec (@jure)

Slides are available at <http://snap.stanford.edu/www2017tutorial/>



Social media has brought us
closer together...



Vieweg, et al. (2010); Kittur, et al. (2013); Burke & Kraut (2016)

Social Media Allows Us..

- Social Media disseminates through Social interaction
- Web is no longer a static library that people passively browse
- Web is a place where people:
 - Consume and create content
 - Interact with other people:
 - Internet forums, Blogs, Social networks, Twitter, Wikis, Podcasts, Slide sharing, Bookmark sharing, Product reviews, Comments, ...

Social Media: Opportunities

- Any user can share and contribute content, express opinions, link to others
- This means: Can data-mine opinions and behaviors of millions of users to gain insights into:
 - Human behavior
 - Marketing analytics
 - Product sentiment

Real-time citizen journalism

- **Citizen journalism** provides more valuable information than newswire services
- **Challenge:**
 - Many redundant posts, users have to wade through hundreds of posts to locate useful information
- **Goal:**
 - Mine this data in real-time and produce well organized summaries



Real Time Citizen Journalism in Mumbai Terrorist Attacks

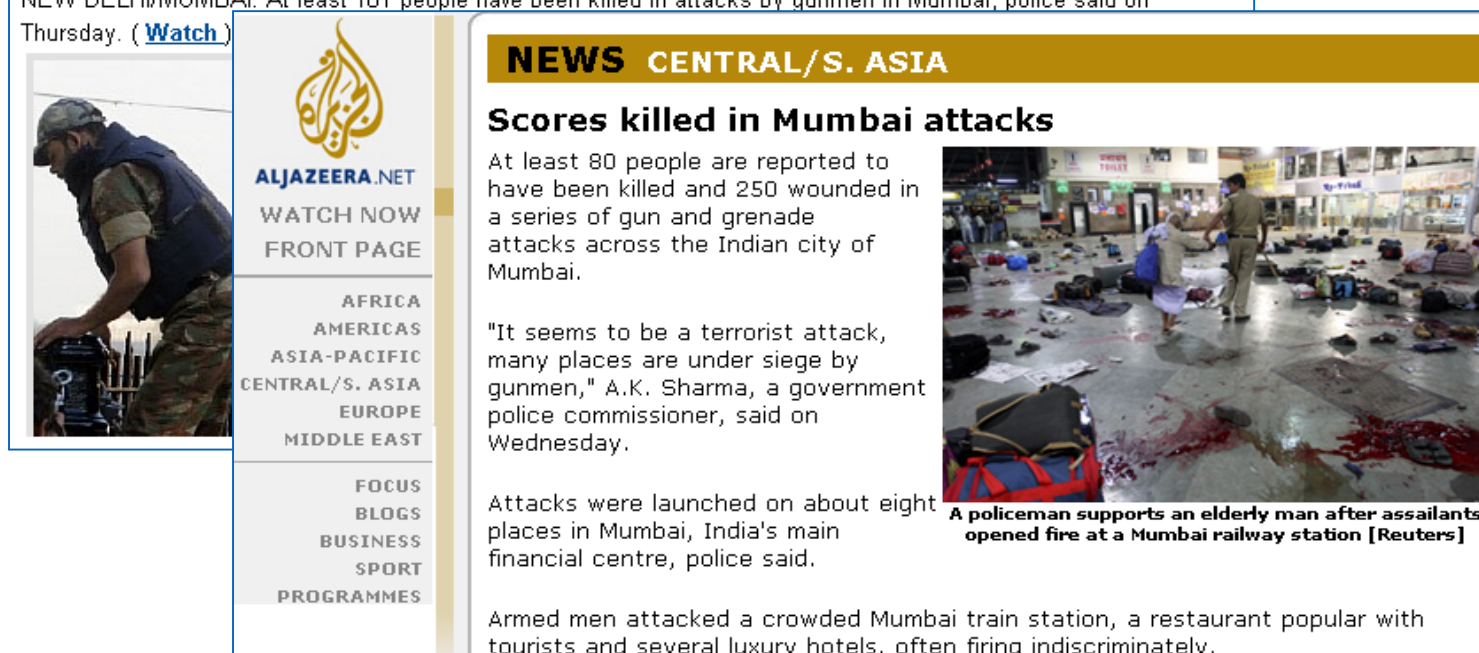
Welcome to Gauravonics Blog! Subscribe to my combined feed [in a feed reader](#) or [by e-mail](#) and you'll never miss a single post. Thanks for visiting!

(The Mumbai terror attack has finally ended after more than 60 hours.

Even as I continue to track instances of citizen journalism in the Mumbai terror attack on this post, I'm trying to make sense of what happened in a [work-in-progress case study](#) and a [Flickr set of screenshots](#) on the role of social media in the Mumbai terror attack. I'm also compiling [reactions on Indian news media's coverage of the terror attack](#).

For more, see my interviews on the role of citizen journalism in the terror attack with [Los Angeles Times](#), [CBS News](#), [BBC](#), [DNA](#), [LiveMint](#), [Associated Press](#), [Journalism.co.uk](#), [Tehelka](#), [NPR](#), [CNN](#), [CUNY TV](#) and [Star Telegram](#).

Finally, the role of the online community in India has not ended with the Mumbai terror



Applications: Reputation management

- Consumer Brand Analytics

- What are people saying about our brand?

- Marketing Communications

- Significant spending on marketing, advertising: Companies trying to position their products
 - Brand analytics helps to determine whether such campaigns are effective

- Product reviews

- Automatically mine product reviews for information on product features, new requests, ...
 - Easy to use, Light weight, Sturdy, Good price, ...



...but there's also a dark side
to social media





"On the Internet, nobody knows you're a dog."

Not everyone has good intentions...


How Trolls Are Ruining the Internet

When Will the Internet Be Safe for Women?

FAKE NEWS IS ABOUT TO GET EVEN SCARIER THAN YOU EVER DREAMED

Time (2016); The Atlantic (2016); Vanity Fair (2017)

Antisocial behavior is prevalent online



Pew Research (2014)

Antisocial behavior is widespread

39% of experts said they expect the online future will be “more shaped” by negative activities

Not everything is true...




Growing Number of Social Media Users are Misinformed About Current Events

- People who rely mainly on social media for news are often misinformed on current events or have formed an opinion on a topic without researching or fact checking.
- Journalists sometimes report actual news stories that are based on inaccurate social media updates.

More social media users rely on links from friends/family contacts

Facebook news driven by friends and family

Percent who get most of their news links from...



Graphic Source: Pew Research Center's Project for Excellence in Journalism / 2012 State of the News Media

While social media can help break real news...

- In 2009, Twitter user Janis Krums (@jkrums) tweeted photos of the US Airways “Miracle on the Hudson” plane that crashed in New York’s Hudson River. He was on one of the ferries that responded to the crash site before traditional news outlets knew about the incident.
- SOURCE: Social Media Today
<http://socialmediatoday.com/dlawrence/475239/social-media-now-major-source-breaking-news>



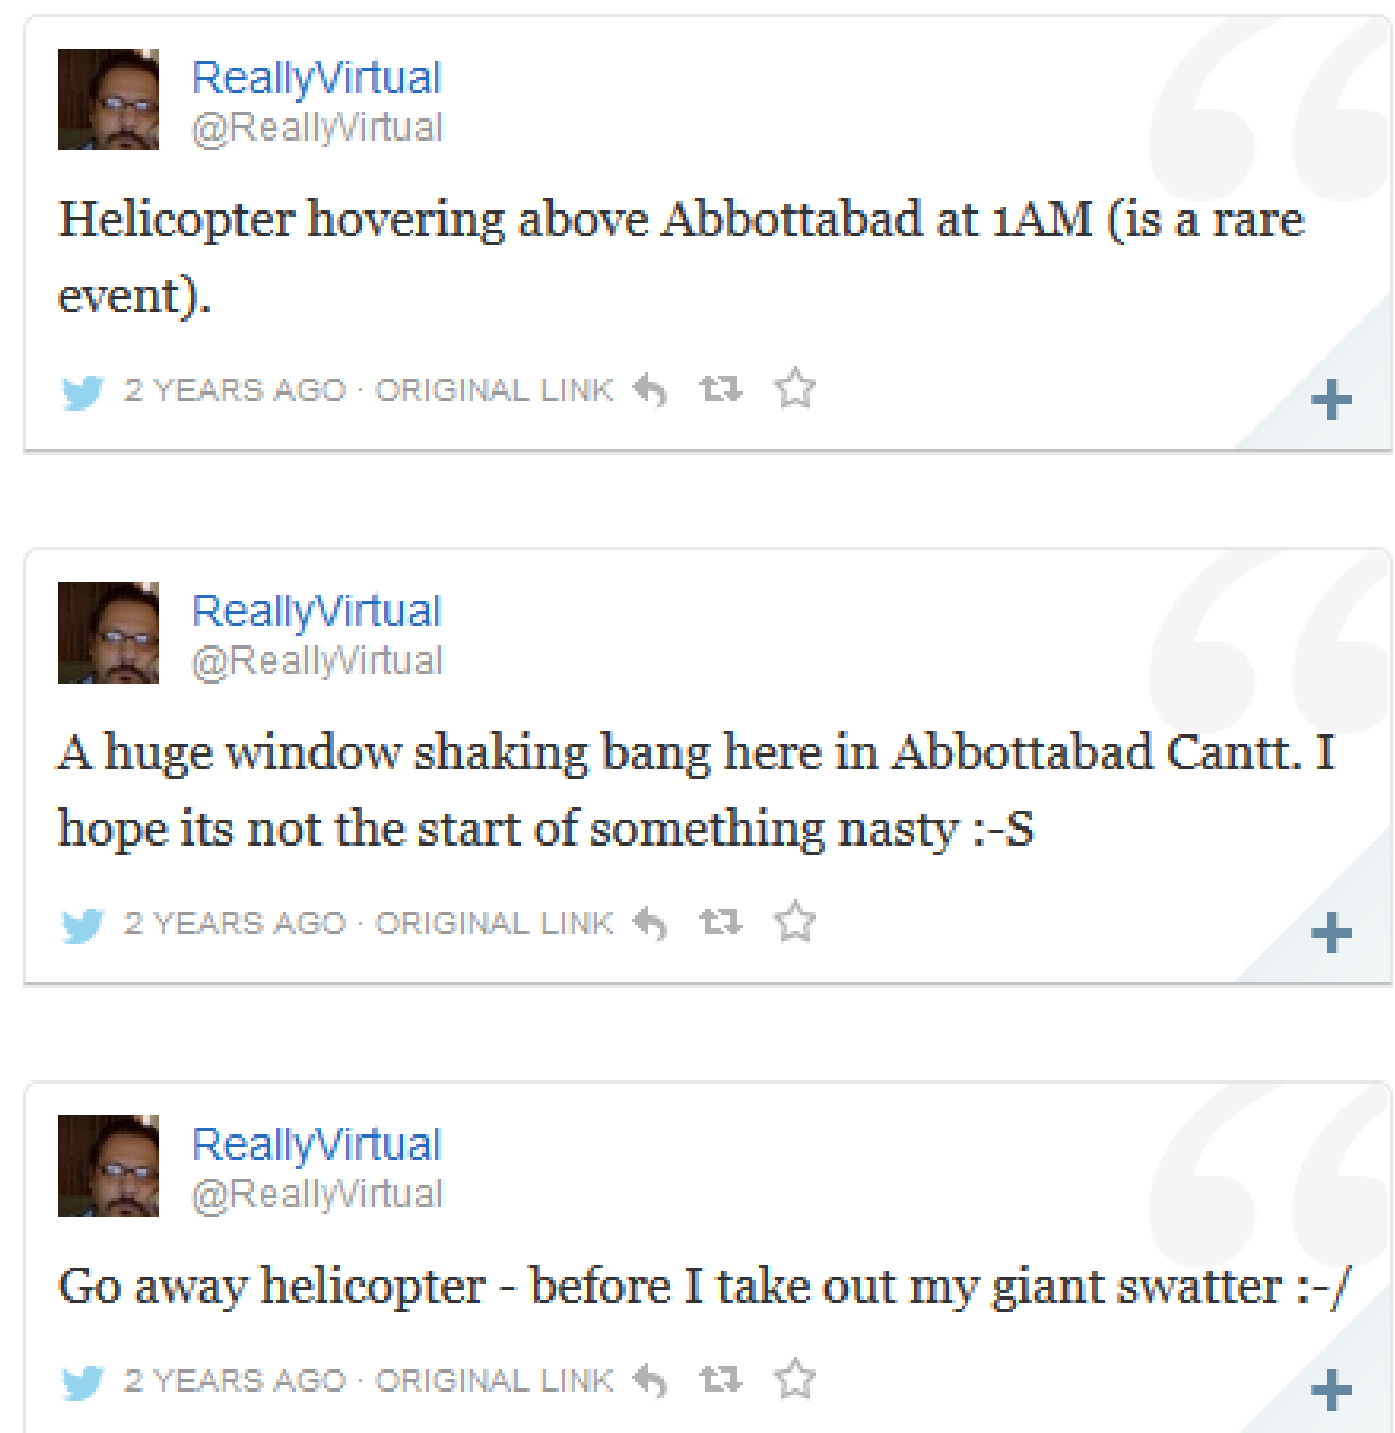
There's a plane in the Hudson. I'm on the ferry going to pick up the people. Crazy.

880,678 1486 days ago

While social media can help break real news...

- In 2011, Twitter user @reallyvirtual tweeted about a helicopter annoying him over his home in Abbottabad, Pakistan, at 1am local time. He was live-tweeting about the US raid that killed Osama Bin Laden.

- SOURCE: Washington Post.
http://www.washingtonpost.com/blogs/blogpost/post/sohaib-athar-tweeted-the-attack-on-osama-bin-laden%20--%20without-knowing-it/2011/05/02/AF4c9xXF_blog.html



...Not everything published online is necessarily true.

- Jackie Chan died in March 2011...according to numerous Facebook posts.

The actor then had to release a statement, via his Facebook page, to prove that the rumors were false and that he was alive.

Source: People.com

<http://www.people.com/people/article/0,,20477260,00.html>



...Not everything published online is necessarily true.

- Morgan Freeman died in December 2010 and the news was credited to CNN... according to numerous “tweets” on Twitter.

CNN then had to clarify that they never reported the news and that Morgan Freeman was, in fact, alive.


Source: Mediate

<http://www.mediaite.com/online/cnn-says-no-it-never-sent-out-tweet-reporting-death-of-actor-morgan-freeman-whos-alive/>



Fake News Goes “Viral” Online

- Many social media users “retweet”, “share”, and “like” these erroneous reports.
- These users include average citizens and journalists who don’t fact-check before spreading the news.



This tutorial:

Antisocial Behavior on the Web: Characterization and Detection

This tutorial: Goals

Towards design of better social systems

Can we use these insights to design better online interactions?

- How do we identify trolls/vandals/rumors and/or predict antisocial behavior?
- What kinds of interventions could minimize antisocial behavior?

Tutorial Outline

Malicious users

Trolling

Sockpuppets

Vandals

Misinformation

Fake reviews

Hoaxes

<http://snap.stanford.edu/www2017tutorial>

Goals of this Tutorial


- Discuss a set of methods for analyzing malicious behavior online
- Summarize insights obtained by various studies into different aspects of malicious behavior online:
 - Malicious users
 - Malicious content

Challenges in analyzing antisocial behavior




Population imbalance

Smaller proportion of behavior ($< 10\%$) is antisocial



Limited labels

Little known information about antisocial behavior



Deceptive behavior

Antisocial behavior tends to masquerade as benign