

# Algorithmic challenges in topic distillation and classification

Luca Becchetti      Andrea Vitaletti  
“Sapienza” University of Rome

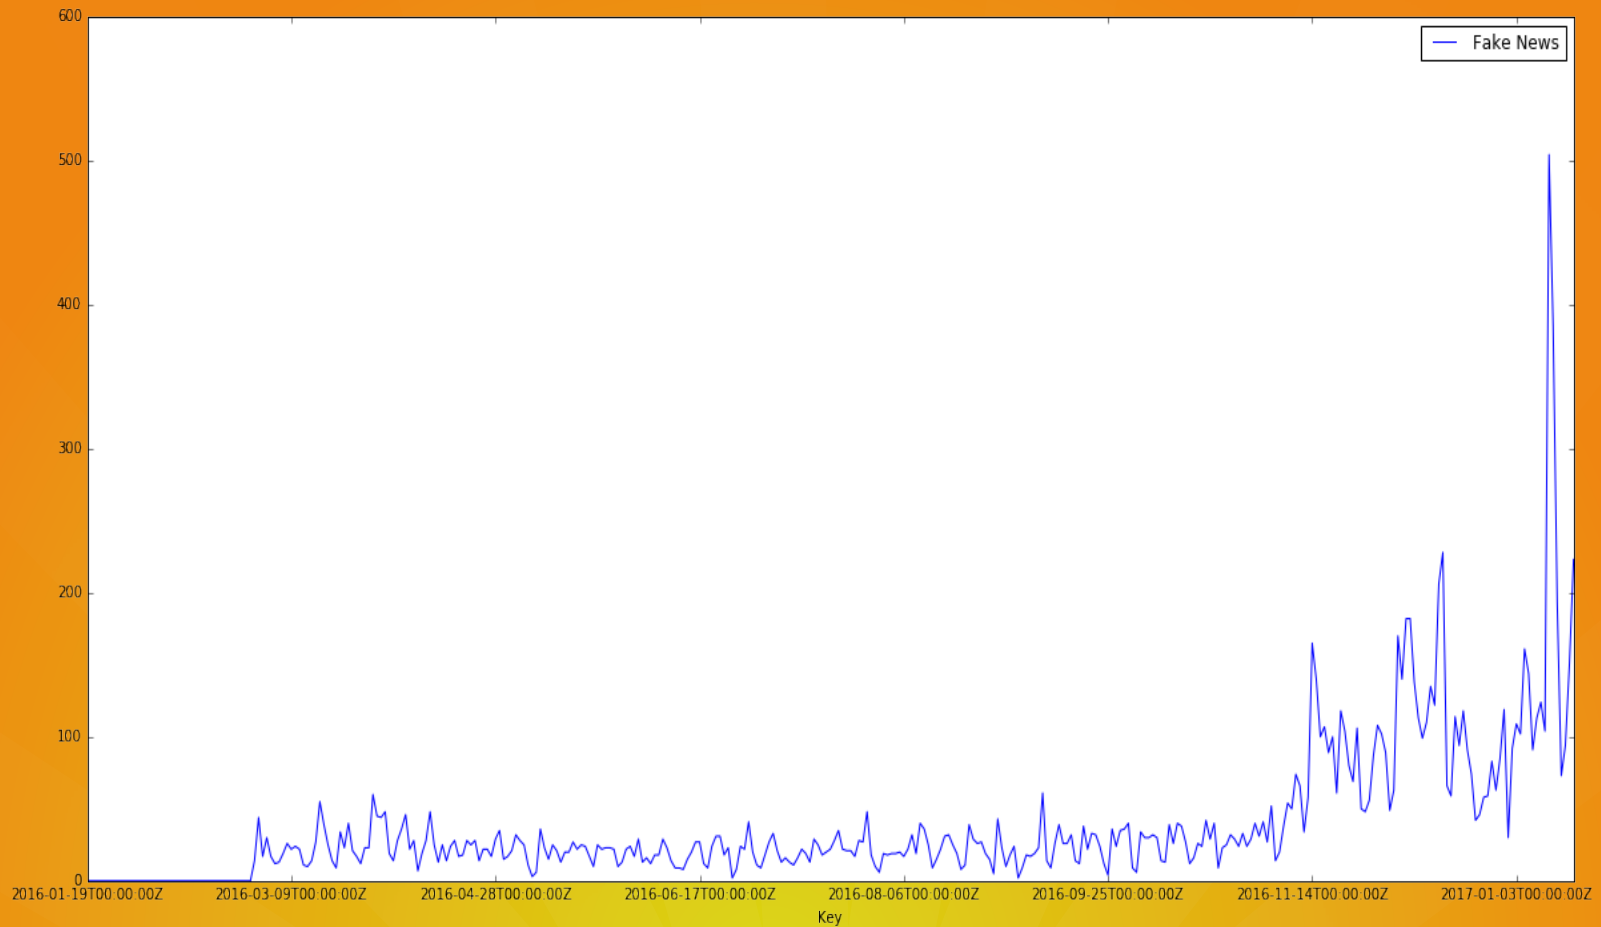
# Outline

- Extracting topics from documents (Web pages, Twitter streams etc.) can be challenging
- Detecting documents containing fake news is not easier ...

In this talk

- 1) Overview of related techniques for topic extraction and document classification
- 2) Challenges in document contents' quality assessment

# Fake news in the news



Courtesy Volocom.it: <http://www.volocom.it/>

# A number of related problems (very partial list)

- Spam detection
  - Castillo, Carlos, et al. "Know your neighbors: Web spam detection using the web topology." Proceedings of ACM SIGIR, 2007.
- Fake/fictitious product opinions/reviews
  - Jindal, Nitin, and Bing Liu. "Review spam detection." Proceedings of ACM WWW 2007
- Wikipedia hoaxes
  - Kumar, Srijan, Robert West, and Jure Leskovec. "Disinformation on the web: Impact, characteristics, and detection of wikipedia hoaxes." Proceedings of WWW 2016

# General approach

- Underlying problems
  - Topic identification
  - Classification
- Meta-algorithm
  - Use data mining/IR techniques to identify salient features
  - Use Machine Learning techniques to classify/profile objects (e.g., news articles, posts etc.) based on their features

# Dissecting Web contents

Source

The screenshot shows a web browser address bar with the URL: [thelastlineofdefense.org/breaking-obamas-being-sued-by-malia-and-sashas-biological-father/](http://thelastlineofdefense.org/breaking-obamas-being-sued-by-malia-and-sashas-biological-father/). Below the address bar is a breadcrumb trail: Home > Everything Bagels With Cream Cheese > BREAKING: Obamas Being Sued By Malia And Sasha's Biological Father. The main headline is "BREAKING: Obamas Being Sued By Malia And Sasha's Biological Father". Below the headline is a byline: "Stryker" followed by "Everything Bagels With Cream Cheese, Hate group, Loose Meat Hoagies" and a comment icon with "0". Below the byline is a call to action: "Join The Resistance And Share This Article Now!". At the bottom of the article snippet, there is a share count of "13.3k SHARES" and two buttons: "SHARE ON FACEBOOK" and "SHARE ON TWITTER".

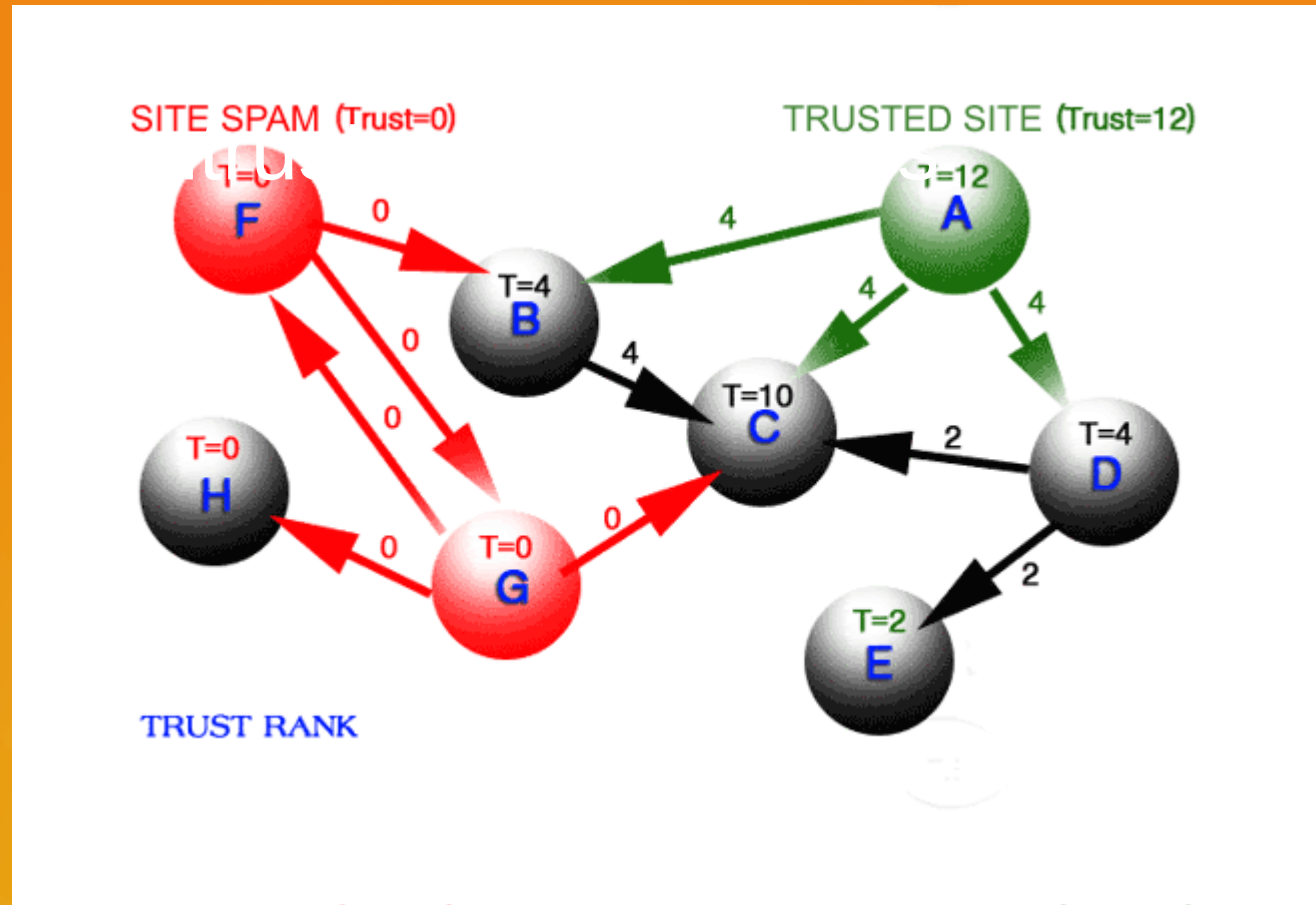
Text/contents

Network/  
propagation

Debunked by **Snopes**

Exploiting (static) network structure

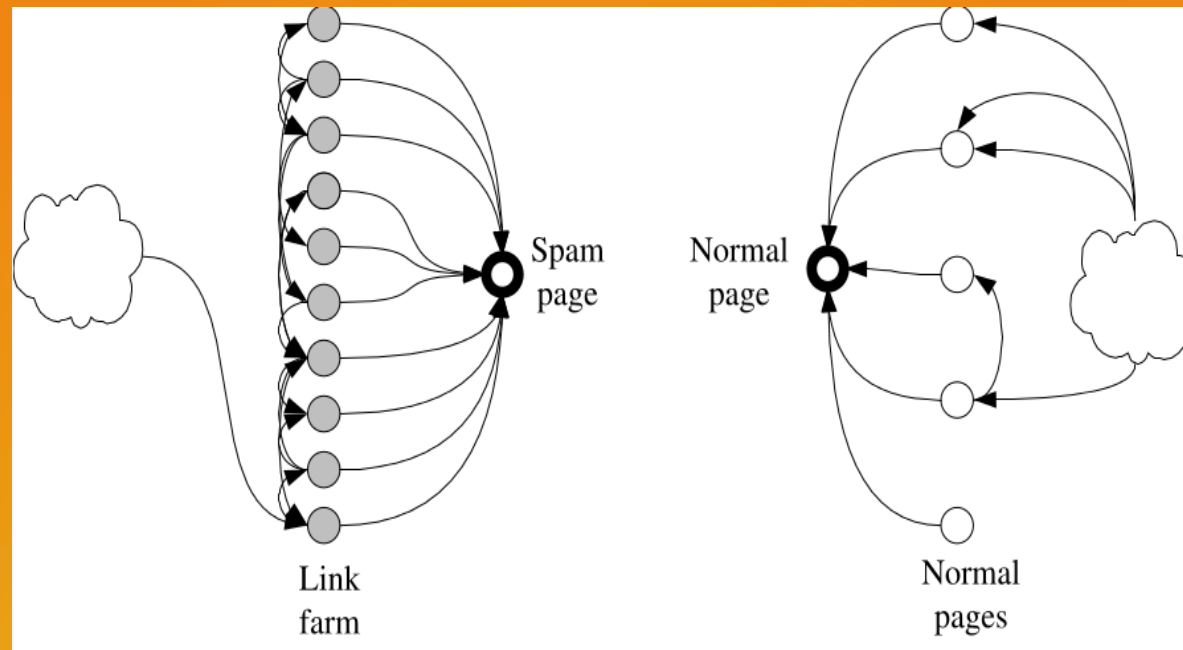
# Assigning authority scores



- Representative approach
  - Gyöngyi, Zoltán, Hector Garcia-Molina, and Jan Pedersen. "Combating web spam with trustrank." VLDB 2004.
- Need to know a representative set of "spam" sources



# Exploiting structure



E.g.:

- B., Castillo, Donato, Leonardi, Baeza-Yates. "Link analysis for web spam Detection." ACM Transactions on the Web, 2008.
- Ghosh, Saptarshi, et al. "Understanding and combating link farming in the twitter social network." WWW 2012.

# Network and propagation

# Network dynamics

- Different types of news/rumors exhibit different propagation patterns
  - E.g.: Del Vicario, Michela, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H. Eugene Stanley, and Walter Quattrociocchi. "The spreading of misinformation online." Proceedings of the National Academy of Sciences 113, no. 3 (2016): 554-559.
- Competing information campaigns
  - E.g.: Budak, Ceren, Divyakant Agrawal, and Amr El Abbadi. "Limiting the spread of misinformation in social networks." WWW 2011



Leveraging text/contents

# Analyzing text

- Topic distillation
  - A key ingredient in textual analysis
  - Algebraic techniques
    - E.g., Deerwester, Scott, et al. "Indexing by latent semantic analysis." Journal of the American society for information science 41.6 (1990): 391
    - Many variants/extensions since then
  - Ad hoc techniques
    - E.g., Pervin, Nargis, et al. "Fast, scalable, and context-sensitive detection of trending topics in microblog post streams." ACM Transactions on Management Information Systems, 2013
  - Statistical inference/machine learning (see previous talk)
- NLP/statistical inference
  - E.g., Qazvinian, Vahed, et al. "Rumor has it: Identifying misinformation in microblogs." Proceedings of the Conference on Empirical Methods in Natural Language Processing. Association for Computational Linguistics, 2011.

# Is Fake News a Machine Learning Problem?

(Zachary Lipton UCSD)

- Fake news challenge
  - Dean Pomerleau (CMU), Delip Rao (Twitter)
  - <http://www.fakenewschallenge.org/>
  - Idea: apply data mining and machine learning techniques to fake news detection
- Contest's goals had to be revised after experts and scholars (including Pomerleau and Rao themselves) identified non-trivial issues in automating this task
  - For a less formal discussion, see [Cade Wetz's article on Wired magazine](#)

# Is Fake News a Machine Learning Problem? (2)

## (Zachary Lipton UCSD)

- Original version
  - Input: a string representing a claim or headline, e.g. “Climate change is a Martian takeover operation”
  - Output
    - Boolean fakeness indicator ( $\{0,1\}$ )
    - Confidence score
    - Provenance URL to support
    - A confidence threshold for accepting/rejecting
- Revised version

### Input

A headline and a body text - either from the same news article or from two different articles.

### Output

Classify the stance of the body text relative to the claim made in the headline into one of four categories:

1. **Agrees:** The body text agrees with the headline.
2. **Disagrees:** The body text disagrees with the headline.
3. **Discusses:** The body text discuss the same topic as the headline, but does not take a position
4. **Unrelated:** The body text discusses a different topic than the headline