traders jlbollen #twitter mood predicts the #DJIA FTW!

Johan Bollen (IU) and Huina Mao (IU)

jbollen@indiana.edu, huinmao@indiana.edu
School of Informatics and Computing
Center for Complex Networks and Systems Research
Indiana University

May 24, 2011

So what's the story?

We upload paper to arxiv in October 2010

Johan Bollen, Huina Mao, and Xiao-Jun Zeng. Twitter mood predicts the stock market. Journal of Computational Science, 2(1), March 2011, Pages 1-8

3 days later:

So what's the story?

We upload paper to arxiv in October 2010

Johan Bollen, Huina Mao, and Xiao-Jun Zeng. Twitter mood predicts the stock market. Journal of Computational Science, 2(1), March 2011, Pages 1-8

3 days later:

msnbc...

So what's the story?

We upload paper to arxiv in October 2010

Johan Bollen, Huina Mao, and Xiao-Jun Zeng. Twitter mood predicts the stock market. Journal of Computational Science, 2(1), March 2011, Pages 1-8

3 days later:

msnbc...

bloomberg...

So what's the story?

We upload paper to arxiv in October 2010

Johan Bollen, Huina Mao, and Xiao-Jun Zeng. Twitter mood predicts the stock market. Journal of Computational Science, 2(1), March 2011, Pages 1-8

3 days later:

msnbc...

bloomberg...

cnn...

So what's the story?

We upload paper to arxiv in October 2010

Johan Bollen, Huina Mao, and Xiao-Jun Zeng. Twitter mood predicts the stock market. Journal of Computational Science, 2(1), March 2011, Pages 1-8

3 days later:

msnbc...

bloomberg...

cnn...

2 months later

How did we get here from there?

Computational social science at Indiana University:

Do societies experience varying mood states like individuals?
If so, can we assess such mood states from online materials and determine its socio-economic correlates?

How did we get here from there?

Computational social science at Indiana University:

Do societies experience varying mood states like individuals?
If so, can we assess such mood states from online materials and determine its socio-economic correlates?

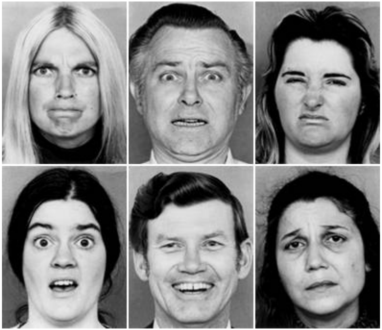

How did we get here from there?

Computational social science at Indiana University:

Do societies experience varying mood states like individuals?
If so, can we assess such mood states from online materials and determine its socio-economic correlates?

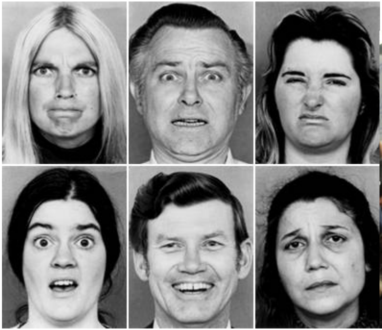

2009-2011: ongoing research on public sentiment in social media feeds

- Huina Mao, Alberto Pepe, and Johan Bollen. Structure and evolution of mood contagion in the Twitter social network. Proceedings of the International Sunbelt Social Network Conference XXX, Riva del Garda, Italy, July 2010.
- Johan Bollen, Huina Mao, and Alberto Pepe. Determining the public mood state by analysis of microblogging posts. Proceedings of the Proc. of the Alife XII Conference, Odense, Denmark, MIT Press, August 2010. (Reviewer-selected for Plenary Presentation)
- Johan Bollen, Alberto Pepe, and Huina Mao. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. ICWSM11, Barcelona, Spain, July 2011 (arXiv: 0911.1583) - poster.

Public mood analysis: lots of squiggly lines...

Artifact or reality?

Public mood analysis: lots of squiggly lines...

Artifact or reality?

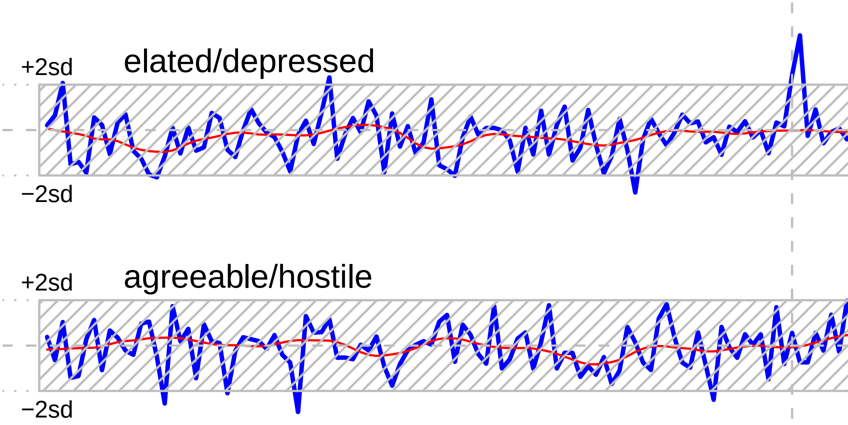

Cross-validating to broad socio-economic indices...

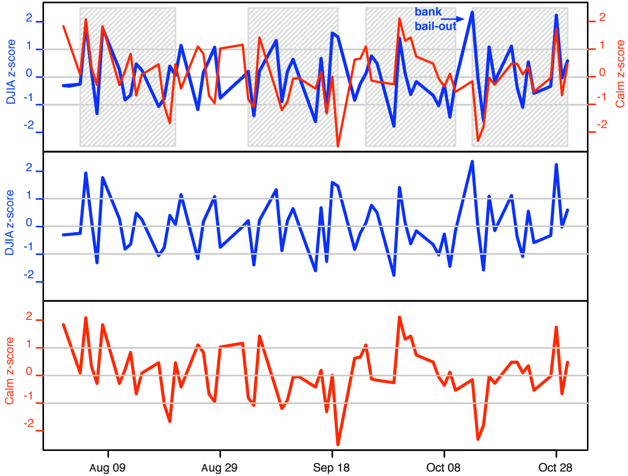

Prediction accuracy: 86.7%?!

Paul Hawtin of Derwent Capital Markets:

> *"For years, investors have widely accepted that financial markets are driven by fear and greed but we've never before had the technology or data to be able to quantify human emotion. This is the 4th dimension."*

Outline

Outline

1. Introduction
   - Microblogging: canary in a coal mine
   - Sentiment analysis
2. Methods
   - Data
   - Sentiment tracking instrument
3. Results
   - Case-studies
   - DJIA
4. Conclusions
   - Discussion
   - Literature

Microblogging: casu Twitter!

tweets and updates

users broadcast brief text updates to the public or to a limited group of contacts: 140 characters or less
Twitter, Facebook, Myspace

Examples

- "Our Rights from Creator (h/t @JLocke). Life, Liberty, PoH FTW! Your transgressions = FAIL. GTFO, @GeorgeIII. -HANCOCK et al."
- "at work feeling lousy"

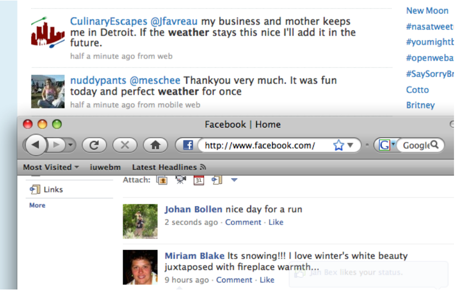

Analyzing the chatter

Twitter: +100M tweets per day, +150M users > many industrialized nations

Predicting the present

Mapping online traffic provides real-time information which translates to real-world outcomes

- Box office receipts from Twitter chatter: Asur (2010)
- Google trends: flu (verbal autopsies)
- Predicting consumer behavior from search query volume (Goel, 2010)
- Contagion of "Loneliness" and **happiness** in social networks (Cacioppo, 2010 - Bollen, 2011)

Extracting sentiment indicators from text

Happy tweets.
So...nothing quite feels like a good shower, shave and haircut...love it My beautiful friend. i love you sweet smile and your amazing soul i am very happy. People in Chicago loved my conference. Love you, my sweet friends @anonymous thanks for your follow I am following you back, great group amazing people
Unhappy tweets.
She doesn't deserve the tears but i cry them anyway I'm sick and my body decides to attack my face and make me break out!! WTF :(I think my headphones are electrocuting me. My mom almost killed me this morning. I don't know how much longer i can be here.

Different Approaches: Natural Language processing (n-grams) for reviews (Nasukawa, 2003), topics (Yi, 2003), Support Vector Machines: text classification (positive vs. negative) using pre-classified learning sets: Gamon (2004), Pang (2008), Blogs, web sites: mixed approaches. Mishne (2006), Balog (2006), Gruhl (2005),...

Sentiment and mood analysis is difficult for tweets

Individual tweets

- Length: 140 characters, lack of text content
- Diversity: no standardized training sets, dimensions of mood?
- Lack of topic specificity

Public mood from tweet collections and other microblog contents?

- We Feel Fine `http://www.wefeelfine.org/`
- Moodviews `http://moodviews.com`
- Myspace: Thelwall (2009), FB: United States Gross National Happiness `http://apps.facebook.com/usa_gnh/`, Michael Jackson (Kim, 2009)

Ratio of emotional tweets, over time.

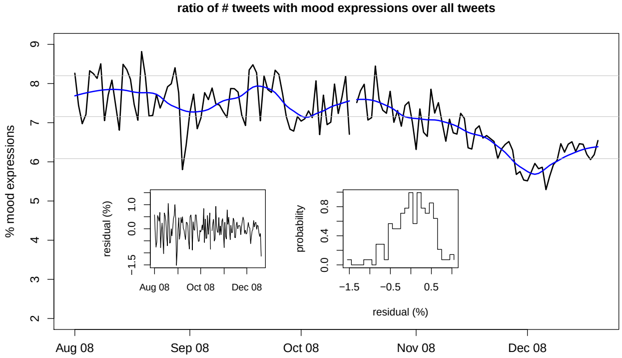

Ratio of tweets containing mood expressions vs. all tweets on a given day, including residuals from trendline.

What we did:

Trends in general public mood from a large-scale collection of tweets

- Each tweet= patient taking psychometric instrument for mood assessment

What we did:

Trends in general public mood from a large-scale collection of tweets

- Each tweet= patient taking psychometric instrument for mood assessment
- Large-scale collection of tweets: 10M, 2006-2008

What we did:

Trends in general public mood from a large-scale collection of tweets

- Each tweet= patient taking psychometric instrument for mood assessment
- Large-scale collection of tweets: 10M, 2006-2008
- Daily public mood assessment: Time series depicting fluctuations of public mood

What we did:

Trends in general public mood from a large-scale collection of tweets

- Each tweet= patient taking psychometric instrument for mood assessment
- Large-scale collection of tweets: 10M, 2006-2008
- Daily public mood assessment: Time series depicting fluctuations of public mood
- Correlations to socio-economic indicators?

Outline

Data sets

Collection of tweets:

- April 29, 2006 to December 20, 2008
- 2.7M users
- **Subset**: August 1, 2008 to December 2008 - 9,664,952 tweets

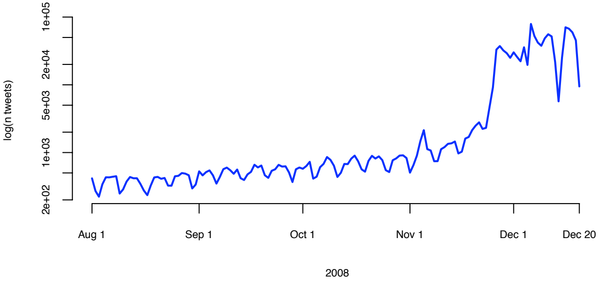

Each tweet:

ID	date-time	type	text
1	2008-11-28 02:35:48	web	Getting ready for Black Friday. Sleeping out at Circuit City or Walmart not sure which. So cold out.
2	2008-11-28 02:35:48	web	@anonymous I didn't know I had an uncle named Bob :-P I am going to be checking out the new Flip sometime soon
	...		

Mood assessment tool

Definition

Uses model derived from existing psychometric instrument (40 years of practice). Maps the content of Tweet to 6 dimensions of human mood. Uses "ancient magic" (just kidding).

calm
alert
sure
vital
kind
happy

Tool built "in-house", beyond mere term matching, learns from the web, lots of behind the scenes processing, continuous development.

Tweet:

I am so not bored. way too busy! I feel really great!

Tweet:

I am so not bored. way too busy! I feel really great!

$$\begin{array}{l} \text{composed/anxious} \\ \text{clearheaded/confused} \\ \text{confident/unsure} \\ \text{energetic/tired} \\ \text{agreeable/hostile} \\ \text{elated/depressed} \end{array} \quad \begin{pmatrix} 0.01725 \\ 0.05125 \\ 0.725625 \\ 0.666625 \\ 0.361 \\ 0.53175 \end{pmatrix}$$

Aggregating daily tweets into a mood time series

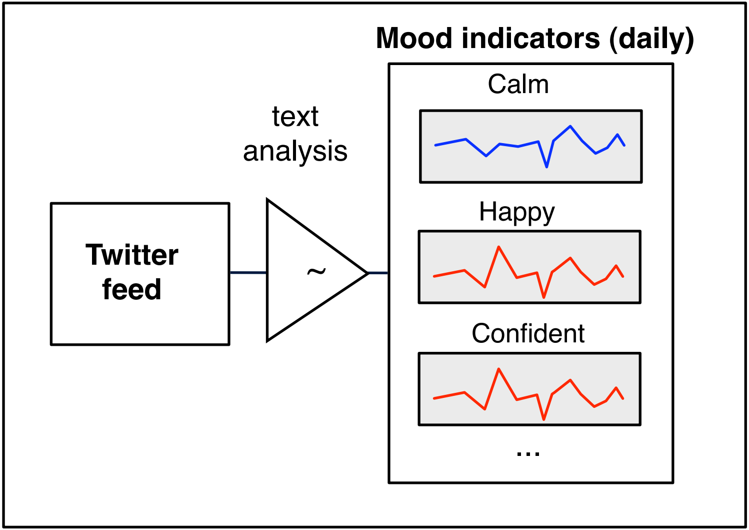

Outline

Public mood trends: overview

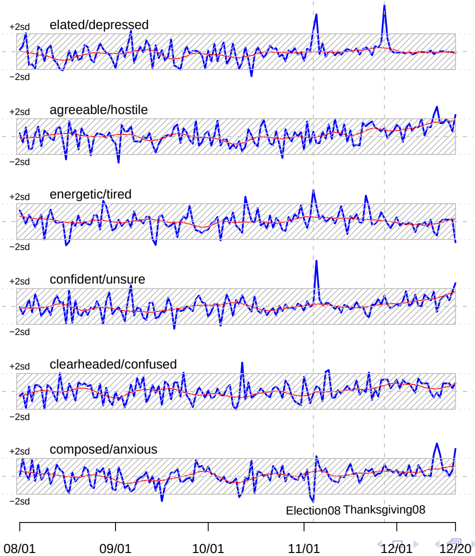

Case study 1: November 4th, 2008 - the presidential election

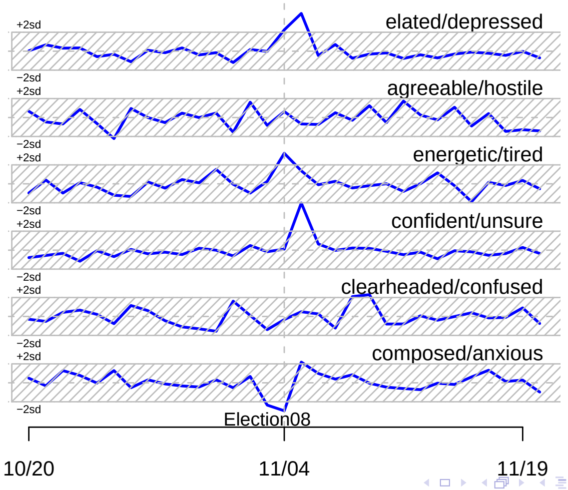

TFIDF scoring of tweet terms

2008 U.S. Presidential Election

Nov 03	Nov 04	Nov 05
robocal	poll	histori
business	plumber	won
voter	result	barack
cleanser	absente	prop
grandmoth	ballot	speech
russert	turnout	result
socialist	barack	president-elect
halloween	citizen	hologram
acknowledg	joe	victori
race	thoughtfulli	ecstat

Table: Top 10 TF-IDF ranking terms 1 day before, on and 1 day after election day.

Case study 2: November 27th, 2008 - Thanksgiving

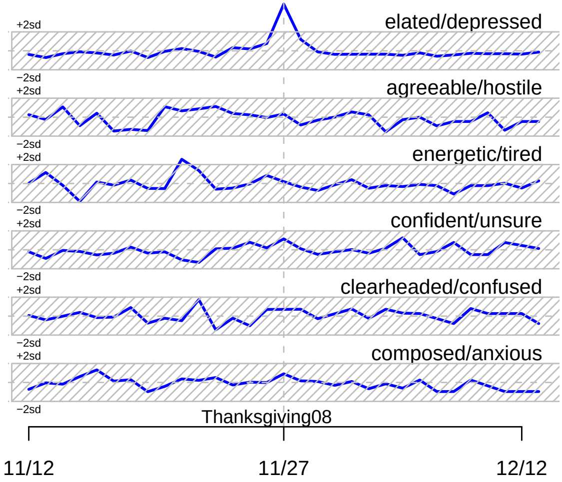

TerraMood:World Mood Analysis from Twitter

Apr 5 2011

Set Date

Select the mood:

calm alert sure vital kind happy

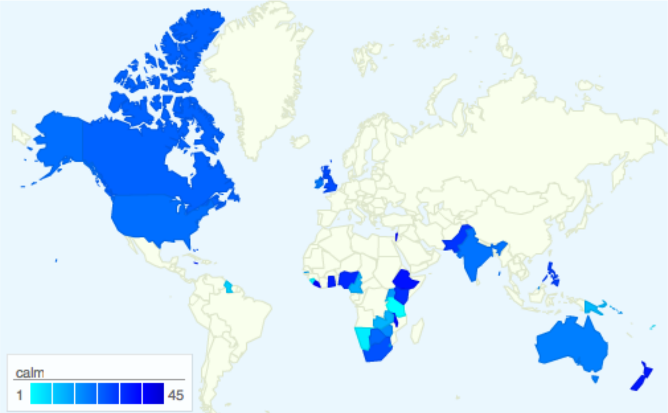

Outline

Comparison to DJIA

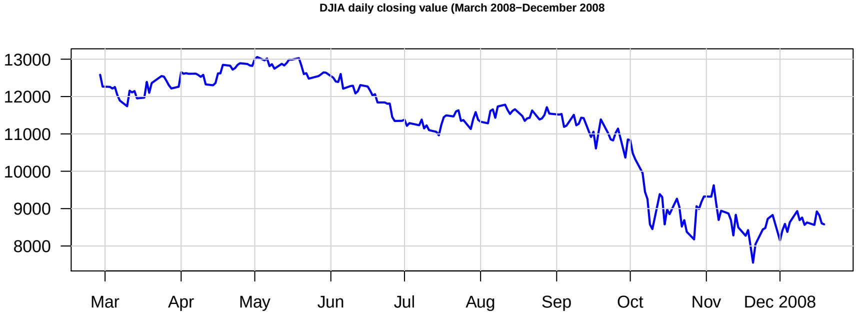

Comparison to DJIA

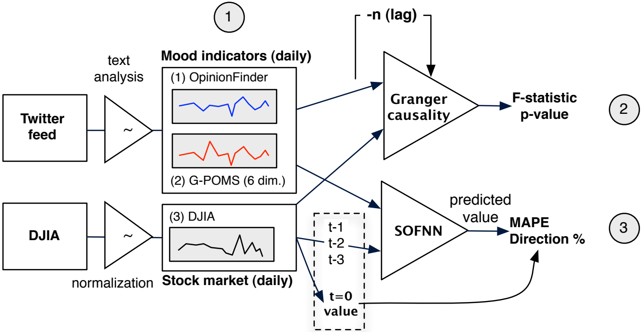

Figure: Methodological diagram outlining use of Granger causality analysis and Self-Organizing Fuzzy Neural Network to predict daily DJIA values from (1) past DJIA values at $t-1, t-2, t-3$, and various permutations of Twitter mood values (OpinionFinder and GPOMS).

bivariate-causal analysis: DJIA vs. public mood

Table: Calm (X_1), Alert (X_2),Sure (X_3), Vital (X_4), Kind (X_5), Happy (X_6)

lag	X_{OF}	X_1	X_2	X_3	X_4	X_5	X_6
1	0.703	**0.080***	0.521	0.422	0.679	0.712	0.300
2	0.633	**0.004****	0.777	0.828	0.996	0.935	0.697
3	0.928	**0.009****	0.920	0.563	0.897	0.995	0.652
4	0.657	**0.03****	0.54	0.61	0.87	0.78	0.68
5	0.235	**0.053***	0.753	0.703	0.246	0.837	**0.05***

Calm vs. DJIA

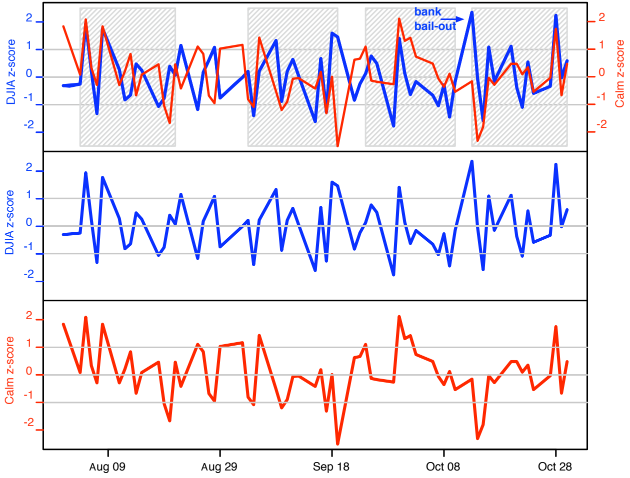
bank
bail-out
DJIA z-score
Calm z-score
2
1
0
-1
-2
Aug 09
Aug 29
Sep 18
Oct 08
Oct 28

Table: DJIA Daily Prediction Using SOFNN

Evaluation	I_{OF}	I_0	I_1	$I_{1,2}$	$I_{1,3}$	$I_{1,4}$	$I_{1,5}$	$I_{1,6}$
MAPE (%)	1.95	1.94	1.83	2.03	2.13	2.05	1.85	**1.79***
Direction (%)	73.3	73.3	**86.7***	60.0	46.7	60.0	73.3	80.0

Outline

Discussion

- "Fear vs. greed"
  - Social media offers window into public mood states
  - Potential to predict market behavior?
  - Results need to be proven in real trading, not just DJIA. Commodities, FOREX?

Discussion

- "Fear vs. greed"
  - Social media offers window into public mood states
  - Potential to predict market behavior?
  - Results need to be proven in real trading, not just DJIA. Commodities, FOREX?
- Some issues:
  - Time period chosen for analysis: Fall 2008 = worst case scenario
  - Short time period for test: subsequent analysis confirms prediction accuracy
  - Lack of historical data: prediction is at scale of 1-4 days, not decades
  - Algorithmic trading vs. hybrid (expert insight)? The crowd is not clairvoyant!
  - Gaming: 100M tweets/day, 150M users!

References

- Johan Bollen, Huina Mao, and Xiao-Jun Zeng. Twitter mood predicts the stock market. Journal of Computational Science, 2(1), March 2011, Pages 1-8, doi:10.1016/j.jocs.2010.12.007, arxiv: abs/1010.3003.
- Johan Bollen, Alberto Pepe, and Huina Mao. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. ICWSM11, Barcelona, Spain, July 2011 (arXiv: 0911.1583)
- Johan Bollen, Bruno Gonalves, Guangchen Ruan and Huina Mao. Happiness is assortative in online social networks. Artificial Life, In Press, Spring 2011 (arxiv:1103.0784)

THANK YOU!

Johan Bollen
jbollen@indiana.edu