

# CS47300: Web Information Search and Management

*Sentiment Analysis*

Prof. Chris Clifton

9 November 2020

Many slides courtesy Prof. [Bing Liu](#)



## Sentiment Analysis: What Is It?

- Goal: derive insight from the *corpus*
  - Too many documents to read/understand
  - Want understanding of the whole (as with pattern recognition)
- First insight into *meaning* of relevant document
  - Positive or Negative statement
- Query: iPhone X
  - Result 1: I love my iPhone X
  - Result 2: I hate my iPhone X

## Sentiment Analysis: Challenges

- Context, not just words
  - I hate my iPhone X
  - I hate being without my iPhone X
- Satire/Irony
  - I love having to charge it every 30 minutes
- Slang
  - Somebody give me an example, I wrote this last night and it is already out of date...
- Others?

3

## Introduction

- Two main types of textual information.
  - **Facts and Opinions**
    - Note: factual statements can imply opinions too.
- **Most current text information processing methods (e.g., web search, text mining) work with factual information.**
- **Sentiment analysis or opinion mining**
  - computational study of opinions, sentiments and emotions expressed in text.
- **Why opinion mining now?** Mainly because of the Web; huge volumes of opinionated text.

4

## Introduction – user-generated media

- **Importance of opinions:**
  - Opinions are important because whenever we need to make a decision, we want to hear others' opinions.
  - **In the past,**
    - **Individuals:** opinions from friends and family
    - **businesses:** surveys, focus groups, consultants ...
- **Word-of-mouth on the Web**
  - **User-generated media:** One can express opinions on anything in reviews, forums, discussion groups, blogs ...
  - **Opinions of global scale:** No longer limited to:
    - **Individuals:** one's circle of friends
    - **Businesses:** Small scale surveys, tiny focus groups, etc.

5

## A Fascinating Problem!

- **Intellectually challenging & major applications.**
  - A popular research topic in recent years in NLP and Web data mining.
  - 20-60 companies in USA alone
- It touches every aspect of NLP and yet is restricted and confined.
  - Little research in NLP/Linguistics in the past.
- Potentially a major technology from NLP.
  - **But** “not yet” and not easy!
  - Data sourcing and data integration are hard too!

6


## An Example Review

- “*I bought an iPhone a few days ago. It was such a nice phone. The touch screen was really cool. The voice quality was clear too. Although the battery life was not long, that is ok for me. However, my mother was mad with me as I did not tell her before I bought the phone. She also thought the phone was too expensive, and wanted me to return it to the shop. ...*”
- What do we see?
  - **Opinions**, **targets of opinions**, and **opinion holders**

7

## Target Object (Liu, Web Data Mining book, 2006)

- **Definition (object)**: An *object*  $o$  is a product, person, event, organization, or topic.  $o$  is represented as
  - a hierarchy of **components**, **sub-components**, and so on.
  - Each node represents a component and is associated with a set of **attributes** of the component



- An opinion can be expressed on any node or attribute of the node.
- To simplify our discussion, we use the term **features** to represent both components and attributes.

8

## What is an Opinion? (Liu, a Ch. in NLP handbook)

- An *opinion* is a quintuple

$(o_j, f_{jk}, so_{ijkl}, h_i, t_l)$ ,

where

- $o_j$  is a target object.
- $f_{jk}$  is a feature of the object  $o_j$ .
- $so_{ijkl}$  is the sentiment value of the opinion of the opinion holder  $h_i$  on feature  $f_{jk}$  of object  $o_j$  at time  $t_l$ .  $so_{ijkl}$  is +ve, -ve, or neu, or a more granular rating.
- $h_i$  is an opinion holder.
- $t_l$  is the time when the opinion is expressed.

9

## Objective – structure the unstructured

- **Objective:** Given an opinionated document,
  - Discover all quintuples  $(o_j, f_{jk}, so_{ijkl}, h_i, t_l)$ ,
    - i.e., mine the five corresponding pieces of information in each quintuple, and
  - Or, solve some simpler problems
- With the quintuples,
  - **Unstructured Text → Structured Data**
    - Traditional data and visualization tools can be used to slice, dice and visualize the results in all kinds of ways
    - Enable qualitative and quantitative analysis.

10

## Sentiment Classification: doc-level

(Pang and Lee, et al 2002 and Turney 2002)

- **Classify a document** (e.g., a review) based on the overall sentiment expressed by opinion holder
  - **Classes:** Positive, or negative (and neutral)
- In the model,  $(o_j, f_{jk}, so_{ijkl}, h_i, t)$ ,
- **Model assumes**
  - Each document focuses on a single object and contains opinions from a single opinion holder.
  - It considers opinion on the object,  $o_j$  (or  $o_j = f_{jk}$ )

11

## Subjectivity Analysis

(Wiebe et al 2004)

- Sentence-level sentiment analysis has two tasks:
  - **Subjectivity classification:** Subjective or objective.
    - **Objective:** e.g., *I bought an iPhone a few days ago.*
    - **Subjective:** e.g., *It is such a nice phone.*
  - **Sentiment classification:** For subjective sentences or clauses, classify positive or negative.
    - **Positive:** *It is such a nice phone.*
- **However.** (Liu, Chapter in NLP handbook)
  - subjective sentences  $\neq$  +ve or -ve opinions
    - E.g., *I think he came yesterday.*
  - Objective sentence  $\neq$  no opinion
    - **Imply -ve opinion:** *My phone broke in the second day.*

12

# Feature-Based Sentiment Analysis

- Sentiment classification at both document and sentence (or clause) levels are **not sufficient**,
  - they do not tell what people like and/or dislike
  - A positive opinion on an object does not mean that the opinion holder likes everything.
  - An negative opinion on an object does not mean .....
- **Objective:** Discovering all quintuples
 
$$(o_j, f_{jk}, so_{ijkl}, h_i, t_l)$$
- With all quintuples, all kinds of analyses become possible.

13

# Feature-Based Opinion Summary

(Hu & Liu, KDD-2004)

“I bought an *iPhone* a few days ago. It was such a nice *phone*. The *touch screen* was really cool. The *voice quality* was clear too. Although the *battery life* was not long, that is ok for me. However, my mother was mad with me as I did not tell her before I bought the phone. She also thought the phone was too *expensive*, and wanted me to return it to the shop. ...”

....

## Feature Based Summary:

### Feature1: **Touch screen**

Positive: 212

- The *touch screen* was really cool.
- The *touch screen* was so easy to use and can do amazing things.

...

Negative: 6

- The *screen* is easily scratched.
- I have a lot of difficulty in removing finger marks from the *touch screen*.

...


### Feature2: **battery life**

...

Note: We omit opinion holders

14

## Visual Comparison (Liu et al. WWW-2005)



## Demo: OpinionEQ

- It performs feature-based sentiment analysis.

**Demo 1:** Compare consumer opinions on three GPS systems, **Garmin, Magellan, Tomtom.**

- Based on a set of features, **price, map, software, quality, size,** etc.

**Demo 2:** Instant page analysis


- The user gives a URL, and the system identifies opinions on the page instantly.

- We also have a Twitter opinion monitoring system (not demoed)



## Demo 1: Compare 3 GPSs on different features

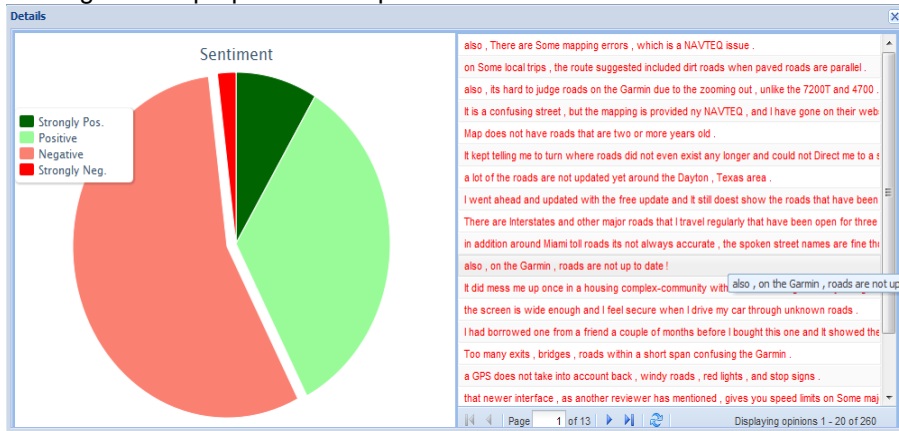
- Each bar shows the proportion of +ve opinion



18


## Demo 1: Detail opinion sentences

- You can click on any bar to see the opinion sentences. Here are negative opinion sentences on the maps feature of Garmin.
- The pie chart gives the proportions of opinions.




# Demo 1: # of feature mentions

- People talked a lot about prices than other features. They are quite positive about price, but not about maps and software.



# Demo 1: Aggregate opinion trend

- More complaints in July - Aug, and in Oct - Dec!

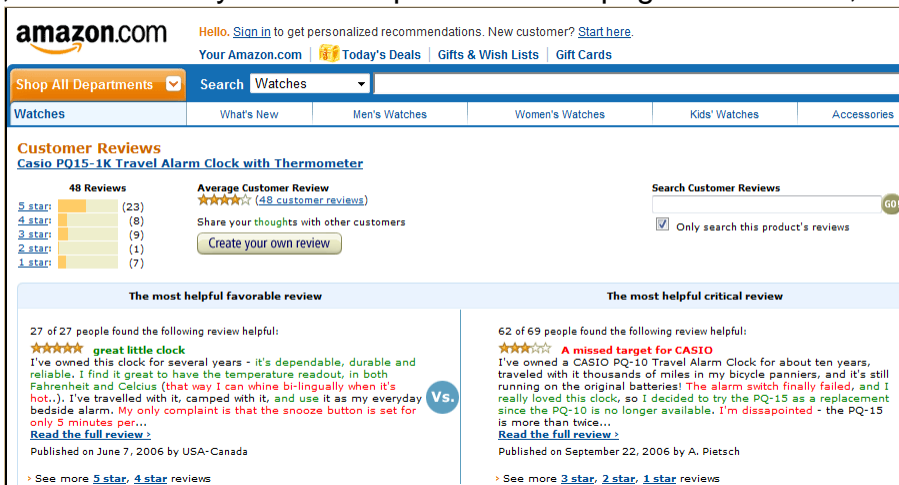


# Other goodies of OpinionEQ

- Allow the user to choose
  - Products/brands,
  - Features
  - Sites
  - Time periods
 for opinion comparison.
- Work on an individual feature for detailed analysis.
- Allow the user to see the full opinion text and also the actual page in the site from where the opinion text was extracted.

# Demo 2 – Instant page analysis

- Given a URL, it automatically identifies opinions on the page. Green: +ve, and red: -ve



The screenshot shows the Amazon.com product page for the Casio PQ15-1K Travel Alarm Clock with Thermometer. The page features a search bar, navigation tabs, and a section for customer reviews. Two reviews are highlighted with a 'Vs.' icon between them:

- The most helpful favorable review (Green):** A 5-star review titled "great little clock" published on June 7, 2006. The reviewer praises the clock's dependability, durability, and readability, and mentions a minor complaint about the snooze button.
- The most helpful critical review (Red):** A 3-star review titled "A missed target for CASIO" published on September 22, 2006. The reviewer expresses disappointment, stating that the clock is no longer available and that they are dissatisfied with the replacement.

## Demo 2 – Instant page analysis

Pros	Cons
<ol style="list-style-type: none"> <li>1. FREE 2-Day Shipping: See details. . <a href="#">(details)</a></li> <li>2. great little clock. . <a href="#">(details)</a></li> <li>3. I've owned this clock for several years - it's dependable, durable and reliable. <a href="#">(details)</a></li> <li>4. I find it great to have the temperature readout, in both Fahrenheit and Celcius (that way I can whine bi-lingually when it's hot. . <a href="#">(details)</a></li> <li>5. The alarm switch finally failed, and I really loved this clock, so I decided to try the PQ-15 as a replacement since the PQ-10 is no longer available. . <a href="#">(details)</a></li> <li>6. love it!, May 10, 2010. . <a href="#">(details)</a></li> <li>7. While the digits of the display are decent enough, the operations to set the clock for time, alarm or any other of the settings, are horrific. <a href="#">(details)</a></li> <li>8. LLBean has a much nicer, user friendly travel clock for less and I could kick myself for not having taken the time to buy another of those.. <a href="#">(details)</a></li> <li>9. I bought from Amazon due to the ease of billing, which was the ONLY reason I bought this instead of the LLBean clock.. <a href="#">(details)</a></li> <li>10. Very accurate timing but temperture is off, April 9, 2010. <a href="#">(details)</a></li> </ol>	<ol style="list-style-type: none"> <li>1. I find it great to have the temperature readout, in both Fahrenheit and Celcius (that way I can whine bi-lingually when it's hot. . <a href="#">(details)</a></li> <li>2. My only complaint is that the snooze button is set for only 5 minutes per. . <a href="#">(details)</a></li> <li>3. A missed target for CASIO. . <a href="#">(details)</a></li> <li>4. The alarm switch finally failed, and I really loved this clock, so I decided to try the PQ-15 as a replacement since the PQ-10 is no longer available. . <a href="#">(details)</a></li> <li>5. I'm dissappointed - the PQ-15 is more than twice. . <a href="#">(details)</a></li> <li>6. I always seem to loose track of time on the computer and phone and needed something.. <a href="#">(details)</a></li> <li>7. Not worth my time or money to return the clock, and I DO NOT recommend this clock at all.. <a href="#">(details)</a></li> <li>8. LLBean has a much nicer, user friendly travel clock for less and I could kick myself for not having taken the time to buy another of those.. <a href="#">(details)</a></li> <li>9. While you can't please everyone with a product, I would not recommend this to anyone. . <a href="#">(details)</a></li> <li>10. Bought it in December 1996 wholesale for \$13 (more than 13 years ago)and i'm still using it EVERYDAY.. Bottom line: Very Very Very accurate and durable BUT temperture show</li> </ol>

24

## Sentiment Analysis is Challenging!

- *“This past Saturday, I bought a **Nokia** phone and my girlfriend bought a **Motorola** phone with **Bluetooth**. We called each other when we got home. **The voice on my phone was not so clear, worse than my previous phone.** **The battery life was long.** **My girlfriend was quite happy with her phone.** **I wanted a phone with good sound quality.** **So my purchase was a real disappointment.** **I returned the phone yesterday.”***

25

## Senti. Analy. is not Just ONE Problem

- $(o_j, f_{jk}, so_{ijkl}, h_i, t_l)$ ,
  - $o_j$  - a target object: **Named Entity Extraction (more)**
  - $f_{jk}$  - a feature of  $o_j$ : **Information Extraction**
  - $so_{ijkl}$  is sentiment: **Sentiment determination**
  - $h_i$  is an opinion holder: **Information/Data Extraction**
  - $t_l$  is the time: **Data Extraction**
- **Co-reference resolution**
- **Relation extraction**
- **Synonym match (voice = sound quality) ...**
- **None of them is a solved problem!**

26

## Extraction of competing objects

- The user first gives a few objects/products as **seeds**, e.g., BMW and Ford.
- The system then identifies other competing objects from the opinion corpus.
- The problem can be tackled with **PU learning** (Learning from positive and unlabeled examples) (Liu et al 2002, 2003).
- See (Li et al. ACL-2010)

27

## Feature extraction

- We proposed a *double propagation* approach in (Qiu et al. IJCAI-2009).
- It exploits the dependency relations of opinions and features to extract features.
  - Opinions words modify object features, e.g.,
  - “This camera takes *great pictures*”
- The algorithm bootstraps using a set of seed opinion words (no feature input).
  - To extract features (and also opinion words)

28

## Rules from dependency grammar

	Relations and Constraints	Output	Examples
R1 <sub>1</sub>	$O \rightarrow O\text{-Dep} \rightarrow F$ s.t. $O \in \{O\}$ , $O\text{-Dep} \in \{MR\}$ , $POS(F) \in \{NN\}$	$f = F$	The phone has a <u>good</u> “screen”. <i>good</i> → <i>mod</i> → <i>screen</i>
R1 <sub>2</sub>	$O \rightarrow O\text{-Dep} \rightarrow H \leftarrow F\text{-Dep} \leftarrow F$ s.t. $O \in \{O\}$ , $O/F\text{-Dep} \in \{MR\}$ , $POS(F) \in \{NN\}$	$f = F$	“iPod” is the <u>best</u> mp3 player. <i>best</i> → <i>mod</i> → <i>player</i> ← <i>subj</i> ← <i>iPod</i>
R2 <sub>1</sub>	$O \rightarrow O\text{-Dep} \rightarrow F$ s.t. $F \in \{F\}$ , $O\text{-Dep} \in \{MR\}$ , $POS(O) \in \{JJ\}$	$o = O$	same as R1 <sub>1</sub> with <i>screen</i> as the known word and <i>good</i> as the extracted word
R2 <sub>2</sub>	$O \rightarrow O\text{-Dep} \rightarrow H \leftarrow F\text{-Dep} \leftarrow F$ s.t. $F \in \{F\}$ , $O/F\text{-Dep} \in \{MR\}$ , $POS(O) \in \{JJ\}$	$o = O$	same as R1 <sub>2</sub> with <i>iPod</i> is the known word and <i>best</i> as the extract word.
R3 <sub>1</sub>	$F_{i(j)} \rightarrow F_{i(j)}\text{-Dep} \rightarrow F_{j(i)}$ s.t. $F_{j(i)} \in \{F\}$ , $F_{i(j)}\text{-Dep} \in \{CONJ\}$ , $POS(F_{i(j)}) \in \{NN\}$	$f = F_{i(j)}$	Does the player play dvd with <u>audio</u> and “video”? <i>video</i> → <i>conj</i> → <i>audio</i>
R3 <sub>2</sub>	$F_i \rightarrow F_i\text{-Dep} \rightarrow H \leftarrow F_j\text{-Dep} \leftarrow F_j$ s.t. $F_i \in \{F\}$ , $F_i\text{-Dep} = F_j\text{-Dep}$ , $POS(F_j) \in \{NN\}$	$f = F_j$	Canon “G3” has a great <u>len</u> . <i>len</i> → <i>obj</i> → <i>has</i> ← <i>subj</i> ← <i>G3</i>
R4 <sub>1</sub>	$O_{i(j)} \rightarrow O_{i(j)}\text{-Dep} \rightarrow O_{j(i)}$ s.t. $O_{j(i)} \in \{O\}$ , $O_{i(j)}\text{-Dep} \in \{CONJ\}$ , $POS(O_{i(j)}) \in \{JJ\}$	$o = O_{i(j)}$	The camera is <u>amazing</u> and “easy” to use. <i>easy</i> → <i>conj</i> → <i>amazing</i>
R4 <sub>2</sub>	$O_i \rightarrow O_i\text{-Dep} \rightarrow H \leftarrow O_j\text{-Dep} \leftarrow O_j$ s.t. $O_i \in \{O\}$ , $O_i\text{-Dep} = O_j\text{-Dep}$ , $POS(O_j) \in \{JJ\}$	$o = O_j$	If you want to buy a <u>sexy</u> , “cool”, accessory- available mp3 player, you can choose iPod. <i>sexy</i> → <i>mod</i> → <i>player</i> ← <i>mod</i> ← <i>cool</i>

## Group synonym features (Zhai et al. 2010)

- Features that are domain synonyms should be grouped together.
- Many techniques can be used to deal with the problem, e.g.,
  - Topic modeling, distributional similarity, etc
- We proposed a semi-supervised learning method
  - Z. Zhai, B. Liu, H. Xu and P. Jia. Grouping Product Features Using Semi-Supervised Learning with Soft-Constraints. *COLING-2010*.

30

## Coreference resolution (Ding and Liu 2010)

- Different from traditional coreference resolution
  - Important to resolve objects and features
  - E.g., “I bought a *Canon S500* camera yesterday. *It* looked beautiful. I took a few *photos* last night. *They* were amazing”.
- Some specific characteristics of opinions can be exploited for better accuracy. See
  - X. Ding and B. Liu, Resolving Object and Attribute Coreference in Opinion Mining. *COLING-2010*.

31

## Identify opinion orientation

- For each feature, we identify the sentiment or opinion orientation expressed by a reviewer.
- Almost all approaches make use of opinion words and phrases. But notice again (a simplistic way):
  - Some opinion words have context independent orientations, e.g., “great”.
  - Some other opinion words have context dependent orientations, e.g., “small”
  - Many ways to use opinion words.
- Machine learning methods for sentiment classification at the sentence and clause levels are also applicable.

32

## Aggregation of opinion words

(Ding and Liu, 2008)

- **Input:** a pair  $(f, s)$ , where  $f$  is a product feature and  $s$  is a sentence that contains  $f$ .
- **Output:** whether the opinion on  $f$  in  $s$  is positive, negative, or neutral.
- Two steps:
  - Step 1: split the sentence if needed based on BUT words (but, except that, etc).
  - Step 2: work on the segment  $s_f$  containing  $f$ . Let the set of opinion words in  $s_f$  be  $w_1, \dots, w_n$ . Sum up their orientations  $(1, -1, 0)$ , and assign the orientation to  $(f, s)$  accordingly.
- In (Ding et al, WSDM-08), step 2 is changed to

$$\sum_{i=1}^n \frac{w_i \cdot o}{d(w_i, f)}$$
with better results.  $w_i \cdot o$  is the opinion orientation of  $w_i$ .  $d(w_i, f)$  is the distance from  $f$  to  $w_i$ .

33



## Basic Opinion Rules (Liu, Ch. in NLP handbook)

Opinions are governed by some rules, e.g.,

1. Neg → Negative
2. Pos → Positive
3. Negation Neg → Positive
4. Negation Pos → Negative
5. Desired value range → Positive
6. Below or above the desired value range → Negative

34

## Basic Opinion Rules (Liu, Ch. in NLP handbook)

7. Decreased Neg → Positive
8. Decreased Pos → Negative
9. Increased Neg → Negative
10. Increased Pos → Positive
11. Consume resource → Negative
12. Produce resource → Positive
13. Consume waste → Positive
14. Produce waste → Negative

36

## Two Main Types of Opinions

- **Direct Opinions:** direct sentiment expressions on some target objects, e.g., products, events, topics, persons.
  - E.g., “the picture quality of this camera is great.”
  - (many are much more complex).
- **Comparative Opinions:** Comparisons expressing similarities or differences of more than one object. Usually stating an ordering or preference.
  - E.g., “car x is cheaper than car y.”

37

## Comparative Opinions (Jindal and Liu, 2006)

- **Gradable**
  - **Non-Equal Gradable:** Relations of the type *greater* or *less than*
    - Ex: “optics of camera A is better than that of camera B”
  - **Equative:** Relations of the type *equal to*
    - Ex: “camera A and camera B both come in 7MP”
  - **Superlative:** Relations of the type *greater* or *less than all others*
    - Ex: “camera A is the cheapest camera available in market”

38

## Mining Comparative Opinions

- **Objective:** Given an opinionated document  $d$ ,. **Extract comparative opinions:**

$(O_1, O_2, F, po, h, t)$ ,

where  $O_1$  and  $O_2$  are the object sets being compared based on their shared features  $F$ ,  $po$  is the preferred object set of the opinion holder  $h$ , and  $t$  is the time when the comparative opinion is expressed.

- **Note:** not positive or negative opinions.

39

## Opinion Spam Detection

(Jindal and Liu, 2007, 2008)

- **Fake/untruthful reviews:**
  - **Write undeserving positive reviews** for some target objects in order to promote them.
  - **Write unfair or malicious negative reviews** for some target objects to damage their reputations.
- Increasing number of customers wary of fake reviews (biased reviews, paid reviews)

40

# An Example Practice of Review Spam

## Belkin International, Inc

- Top networking and peripherals manufacturer | Sales ~ \$500 million in 2008
- Posted an ad for writing fake reviews on amazon.com (65 cents per review)

Timer: 00:00:00 of 60 minutes Want to work on this HIT?  Want to see other HITs?

Write Product Reviews 25-50 Words  
Requester: Mike Bayard  
Qualifications Required: HIT approval rate (%) is not less than 95

**Write a Positive 5/5 Review for Product on Website**

Positive review writing.

- Use your best possible grammar and write in US English only
- Always give a 100% rating (as high as possible)
- Keep your entry between 25 and 50 words
- Write as if you own the product and are using it
- Tell a story of why you bought it and how you are using it
- Thank the website for making you such a great deal
- Mark any other negative reviews as "not helpful" once you post yours

Instructions:

The link below leads to a product on a website. Read-through the product's features and write a positive review for it using the guidelines above to the best of your ability. I have also provided the part number for this product and you can click on the links below to see it on several alternative websites. In order to post some reviews you will need to create an account on the site. You can use your own email address or open a new free webmail account (gmail, yahoo...) and use it to post with.

41



PURDUE  
UNIVERSITY

Department of Computer Science

# Experiments with Amazon Reviews

- June 2006
  - 5.8mil reviews, 1.2mil products and 2.1mil reviewers.
- A review has 8 parts
  - *<Product ID> <Reviewer ID> <Rating> <Date> <Review Title> <Review Body> <Number of Feedbacks> <Number of Helpful Feedbacks>*
- Industry manufactured products “mProducts”
  - e.g. electronics, computers, accessories, etc
  - 228K reviews, 36K products and 165K reviewers.

42

## Some Tentative Results

- Negative outlier reviews tend to be heavily spammed.
- Those reviews that are the only reviews of some products are likely to be spammed
- Top-ranked reviewers are more likely to be spammers.
- Spam reviews can get good helpful feedbacks and non-spam reviews can get bad feedbacks.

## Meeting Social Sciences

- **Extract and analyze political opinions.**
  - Candidates and issues
- **Compare opinions across cultures and lang.**
  - Comparing opinions of people from different countries on the same issue or topic, e.g., [Internet diplomacy](#)
- **Opinion spam (fake opinions)**
  - What are social, culture, economic aspects of it?
- **Opinion propagation** in social contexts
- **How opinions on the Web influence the real world**
  - Are they correlated?
- **Emotion analysis** in social context & virtual world

## Summary

- We briefly defined sentiment analysis problem.
  - **Direct opinions**: focused on feature level analysis
  - **Comparative opinions**: different types of comparisons
  - **Opinion spam detection**: fake reviews.
    - **Currently working with Google** (Google research award).
- A lot of applications.
- **Technical challenges are still huge.**
  - **But I am quite optimistic.**
- **Interested in collaboration with social scientists**
  - opinions and related issues are inherently social.

45

## More details can be found in

- B. Liu, “Sentiment Analysis and Subjectivity.” A Chapter in *Handbook of Natural Language Processing*, 2nd Edition, 2010.
  - (An earlier version) B. Liu, “Opinion Mining”, A Chapter in the book: *Web Data Mining*, Springer, 2006.
- **Download from**: <http://www.cs.uic.edu/~liub/FBS/sentiment-analysis.html>

46