

Sentiment Analysis and Opinion Mining

Slides modified from Bing Liu

Arjun Mukherjee[†]

Course webpage: <http://www.cs.uh.edu/~arjun/courses/nlp>

[†] Contains contents from [Liu, 2008] , and various other sources. Referenced in place.

Facts vs. Opinions

- Two main types of textual information on the Web.

Facts and Opinions

- Current search engines search for facts/statements (of assumed true facts)
Facts usually expressed with topic keywords.
- Search engines do not search for opinions

Facts vs. Opinions

- If we Google:
“How do people think of Motorola Cell phones?”
- We only get facts
- Need to delve deeper. **Q: Why?**

The image shows a Google search interface. The search bar contains the text "How do people think of Motorola Cell phones?". Below the search bar, there are tabs for "Web", "News", "Images", "Shopping", "Videos", "More", and "Search tools". The "Web" tab is selected. The search results show "About 68,600,000 results (0.48 seconds)".

The first result is "Mobile Phones - Motorola" from www.motorola.com/us/consumers/shop-all-mobile-phones/. It lists features like 4G LTE, Developer Editions/Unlocked Phones, and NFC. It also mentions "Droid Razr M. \$0.99. Shop ..." and "Moto X - Droid Maxx - Moto X Developer Edition - Droid Razr M".

The second result is "Inventor of cell phone: We knew someday everybody would ..." from www.cnn.com/2010/TECH/mobile/07/09/cooper.cell.phone.inventor/. It mentions "Jul 9, 2010 - Cooper and his team at Motorola, the communications company, created ... The company's first commercial cellular phone, the 'DynaTAC,' went on sale 10 ... Did people think it was unbelievable, impossible, unnecessary?". A red box highlights the text "Did people think it was unbelievable, impossible, unnecessary?". A red arrow points from the bullet point "We only get facts" to this box.

The third result is "[Weekend Poll] Has The Motorola-Lenovo Buyout Made You ..." from www.androidpolice.com/.../weekend-poll-has-the-motoro... It mentions "by David Ruddock - Feb 1, 2014 - And why do people think Lenovo would buy Motorola and then I was all set on making the Moto X line my phone of choice going forward, ...". A red box highlights the text "And why do people think Lenovo would buy Motorola and then I was all set on making the Moto X line my phone of choice going forward, ...". A red arrow points from the bullet point "We only get facts" to this box.

The fourth result is "Amazon.com: Verizon Motorola Droid X WiFi 3G Camera ..." from www.amazon.com. It mentions "What Other Items Do Customers Buy After Viewing This Item? Motorola ... Also, a lot of people think they simply must have their music on their iPods. ... I have a good DSLR and it doesn't bother me that the camera in my phone is not that great." A red box highlights the text "Also, a lot of people think they simply must have their music on their iPods. ... I have a good DSLR and it doesn't bother me that the camera in my phone is not that great." A red arrow points from the bullet point "We only get facts" to this box.

Facts vs. Opinions

- If we Google:
“How do people think of Motorola Cell phones?”
- We only get facts
- Need to delve deeper. **Q:Why?**
 - ❑ Opinions are hard to express with a few keywords
 - ❑ Need to mine/understand query and also retrieve opinions from web user contents/opinions
 - ❑ This is opinion retrieval/search – just the surface of opinion mining!

Google

How do people think of Motorola Cell phones?

Web News Images Shopping Videos More Search tools

About 68,600,000 results (0.48 seconds)

Mobile Phones - Motorola
www.motorola.com/us/consumers/shop-all-mobile-phones/ Motorola
Mobile Phones ... Features. 4G LTE; Developer Editions/Unlocked Phones; NFC ...
Mobile Phones (10) ... Droid Razr M. \$0.99. Shop ...
Moto X - Droid Maxx - Moto X Developer Edition - Droid Razr M

Inventor of cell phone: We knew someday everybody would ...
www.cnn.com/2010/TECH/mobile/07/09/cooper.cell.phone.inventor/ Jul 9, 2010 - Cooper and his team at Motorola, the communications company, created ... The company's first commercial cellular phone, the "DynaTAC," went on sale 10 ... Did people think it was unbelievable, impossible, unnecessary?

[Weekend Poll] Has The Motorola-Lenovo Buyout Made You ...
www.androidpolice.com/.../weekend-poll-has-the-motoro... Android Police
by David Ruddock - Feb 1, 2014 - And why do people think Lenovo would buy Motorola and then I was all set on making the Moto X line my phone of choice going forward, ...

Amazon.com: Verizon Motorola Droid X WiFi 3G Camera ...
www.amazon.com > ... > No-Contract Cell Phones > Phones > Amazon.com
What Other Items Do Customers Buy After Viewing This Item? Motorola ... Also, a lot of people think they simply must have their music on their iPods. ... I have a good DSLR and it doesn't bother me that the camera in my phone is not that great.

Opinions as User Generated Content

- Word-of-mouth on the Web
- One can express personal experiences and opinions on almost anything, at review sites, forums, discussion groups, blogs ... (called the user generated content.)
- They contain valuable information: Basis for deeper knowledge of human experience
- The web these days contain vast amount of user generated contents: Amazon.com, Facebook, Twitter, etc.
- **Our interest:** Mine opinions (sentiments) expressed in the user-generated content on various topics
- An intellectually very challenging problem.
- Practically very useful. **Q: Why?**

Opinion Mining Applications

- **Businesses and organizations:** product and service benchmarking. Market intelligence.
Business spends a huge amount of money to find consumer sentiments and opinions.
Can be used for product development/improvement in future releases
- **Individuals:** interested in other's opinions when
Purchasing a product or using a service,
Finding opinions on political topics,
- **Ads placements:** Placing ads in the user-generated content
Place an ad when one praises a product.
Place an ad from a competitor if one criticizes a product.
- **Opinion retrieval/search:**
Providing general search for opinions using deep web architectures.

Opinion Mining Defined

- Opinion mining (also known as sentiment analysis) refers to the use of natural language processing, and text analysis/mining techniques to identify and extract subjective information (emotions, appraisals, evaluations, attitudes, etc.) in source materials (text documents).
- **Intellectually challenging & major applications.**
 - A very popular research topic in recent years in NLP and Web data mining.
 - 100+ companies in USA alone
- It touches everything aspect of NLP and yet is restricted and confined.
 - Little research in NLP/Linguistics in the past.
- Potentially a major technology from NLP.
 - But it is not easy!**

Opinionated Evaluations: Two types

- **Regular Opinions:** sentiment expressions on some entities, e.g., products, events, topics, persons.
E.g., “the picture quality of this camera is **great**”
Subjective
- **Comparisons:** relations expressing similarities or differences of more than one entity. Usually expresses an ordering.
E.g., “car x is **cheaper** than car y.”
“I **prefer** tea to coffee”
Objective or subjective.
- **Red colored** terms may be crudely referred to as sentiment/opinion words

Opinion Search and Retrieval

- Is opinion search as convenient as general Web search?
- Opinion search critical for decision making - want others' opinions
- Wouldn't it be nice? If we could issue Google search queries such as
Opinions: “**Motorola cell phones**”
Comparisons: “**Motorola vs. Nokia**”
- Some work not still enough room for improvement! (trending research area...)

amazon Prime

ARJUN's Amazon.com Today's Deals Gift Cards Sell Help

Shop by Department Search All ipod Go Hello, ARJUN Your Account

MP3 Players & Accessories Best Sellers New Arrivals MP3 Players Accessories Portable Speakers Digital Music Store Trade-In All Electronics

Electronics > Portable Audio & Video > MP3 Players & Accessories > MP3 Players

Apple iPod classic 160 GB Black (7th Generation)
by Apple
★★★★★ 2,003 customer reviews | 215 answered questions

4.3 out of 5 stars

5 star 1,382
4 star 305
3 star 84
2 star 72
1 star 160

See all 2,003 reviews

“It is an amazing little device, very easy to use and great sound quality.”
BT | 178 reviewers made a similar statement

“I love my iPod, all i wanted was to be able to listen to music and watch videos and this does what I want.”
Dallas Winters | 261 reviewers made a similar statement

“I like having the 160 GB of storage space.”
K. Adams | 229 reviewers made a similar statement

- 160 GB capacity for 40,000 songs, 25,000 photos, or 200 hours of video
- Up to 36 hours of music playback or 6 hours of video playback when fully charged
- 2.5-inch color LCD with LED backlight and 320 x 240 pixel resolution
- Supports AAC, Protected AAC, MP3, MP3 VBR, Audible, Apple Lossless, AIFF, and WAV audio formats
- Highly recyclable aluminum and stainless steel enclosure
- 2.5-inch LED-backlit display
- 160 GB capacity for 40,000 songs, 25,000 photos, or 200 hours of video. Up to 36 hours of music playback or 6 hours of video playback when fully charged

Show more

27 new from \$249.99 29 used from \$144.90 4 refurbished from \$230.00

Typical Opinion Search Queries

- Find the opinion of a person or organization (opinion holder) on a particular entity or an aspect of the entity.
E.g., what is Bill Clinton's opinion on abortion?
- Find positive and/or negative opinions on a particular entity (or some aspects of the entity), e.g.,
customer opinions on a digital camera.
public opinions on a political topic.
- Find how opinions on an entity change over time.
- How entity A compares with entity B?
Chrome vs. Internet Explorer
- **All of these are open research problems!**

Opinions in Consumer Reviews

- Reviews can be regarded as traditional surveys.
 - ❑ In traditional survey, returned forms are treated as raw data.
 - ❑ Analysis is performed to summarize the survey results.
 - E.g., % against or for a particular issue, etc.
- In opinion mining and retrieval,
 - Can a summary be produced?
 - What should the summary be?

Uchi

904 Westheimer Road, Houston, TX 77006

(713) 522-4808

Website

Menu

Improve this listing

Add to trip



Ranked #1 of 7,206 restaurants in Houston

261 Reviews

Certificate of Excellence 2014

Cuisines: Japanese

Dining options: Reservations, Late Night



40 visitor photos

Twin Peaks Restaurant

www.twinpeaksrestaurant.com/ Chicken Wings, Ice Cold Beer, & All The Sports You Can Handle!

Sponsored links *

261 reviews from our community

Write a Review

93% Recommend

Excellent		206
Very good		39
Average		11
Poor		4
Terrible		1

Rating summary

Food	
Service	
Value	
Atmosphere	

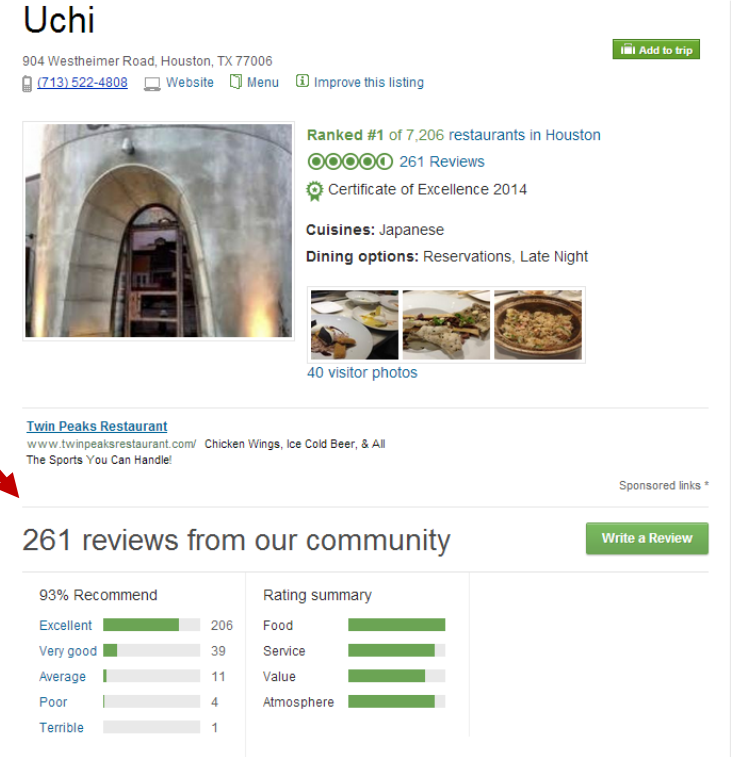
Structured Opinion Mining

- **Q: How to extract features and evaluations to construct a structured summary?**
- Lets look at a review and see what it contains

*“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, **entity/target of opinions**, and **opinion holders**



Uchi

904 Westheimer Road, Houston, TX 77006
(713) 522-4808 Website Menu Improve this listing

Ranked #1 of 7,206 restaurants in Houston
261 Reviews
Certificate of Excellence 2014

Cuisines: Japanese
Dining options: Reservations, Late Night

40 visitor photos

[Twin Peaks Restaurant](#)
www.twinpeaksrestaurant.com/ Chicken Wings, Ice Cold Beer, & All The Sports You Can Handle!

Sponsored links *

261 reviews from our community Write a Review

93% Recommend

Excellent	206
Very good	39
Average	11
Poor	4
Terrible	1

Rating summary

Food	
Service	
Value	
Atmosphere	

Opinion Mining: Abstractions [Hu and Liu, KDD-04]

- Basic components of an opinion

Opinion holder: The person or organization that holds a specific opinion on a particular entity.

Entity: on which an opinion is expressed

Opinion: a view, attitude, or appraisal on an entity from an opinion holder.

*“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. ...”*

- Opinion mining problem abstraction

Put existing research into a common framework


What do we see?

Opinions, **entity/target of opinions**, and **opinion holders**

- We use **consumer reviews of products** to develop the ideas. Other opinionated contexts are similar.

Target Entity

- **Definition (entity):** An *entity* e is a product, person, event, organization, or topic. e 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 *aspect* (*features*) to represent both components & attributes.

Model of an entity

- An entity e_i is represented with a finite set of aspects.
- Let the set of aspects, $A = \{a_1, a_2, \dots, a_n\}$.
- The entity can be expressed with any one of a final set of *entity expressions* $EE_i = \{oe_{i1}, oe_{i2}, \dots, oe_{ik}\}$.
- Each aspect $a_{ij} \in A_i$ of the entity can be expressed by any one of a finite set of *aspect expressions* $AE_{ij} = \{ae_{ij1}, ae_{ij2}, \dots, ae_{ijm}\}$.

"I bought an iPhone a few days ago. It was ~~suck~~ 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. ..."

Model of a review/opinionated document

- **Model of a review:** An **opinion holder** j comments on a subset of the **aspects** $S_j \subseteq A$ of entity e .

For each aspect $a_k \in S_j$ that j comments on, he/she chooses a word or phrase from EE_i to describe the entity, and chooses a word or phrase from AE_{ij} to describe the aspect, and expresses a positive, negative or neutral **opinion** on a_k .

*"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. ..."*

Model of an opinion [Liu, Ch. in NLP handbook]

- An *opinion* is a quintuple

$$(e_j, a_{jk}, so_{ijkl}, h_i, t_l),$$

where

e_j is a target entity.

a_k is an aspect of the entity e_j .

so_{ijkl} is the sentiment value of the opinion of the opinion holder h_i on aspect a_{jk} of entity e_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.

“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. ...”

Goal of Opinion Mining

- **Objective:** Given an opinionated document,
Discover all quintuples $(e_j, a_k, so_{ijkl}, h_i, t_l)$,
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.

*"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. ..."*



Aspect Based Summary[Hu and Liu, '04]:

aspect1: 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**.

...

aspect2: battery life

...

Goal of Opinion Mining

- **Objective:** Given an opinionated document,
Discover all quintuples $(e_j, a_k, so_{ijkl}, h_i, t_l)$,
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.

"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. ..."



Canon EOS 5D Mark III 22.3 MP Digital SLR Camera - Black - Body Only

\$2,495 online ★★★★★ 233 reviews

Reviews

233 reviews




What people are saying

value	■	"good seller nice to deal with you thanks again"
design	■	"D. Great feel and look well designed."
pictures	■	"Picture quality is superb."
zoom/lens	■	"The 24-105L is a good lens to shoot when travelling."
features	■	"The features of this dslr are great."
video	■	"Awesome, amazing indoor photos and video."
color	■	"The colors are vibrant."


Write a review

Visual Comparison [Liu et al., WWW '05]

Summary of
reviews of
Cell Phone 1



Comparison of
reviews of
Cell Phone 1
Cell Phone 2



Mining Opinions in fluid language is hard!

*“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.”*

Q: What is the overall opinion orientation/polarity? +/- ?

Q:How to mine/discover the polarity on aspects in a review?

Opinion Mining Sub-problems

- Recall an opinon is a quintuple: $(e_j, a_k, so_{ijkl}, h_i, t_l)$
- Where extracting each component is an NLP subproblem,
 - e_j - a target entity: Named Entity Extraction (more)
 - a_{jk} - a aspect of e_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
- Synonym match/grouping (voice = sound quality; front desk \approx reception) ...
- Employing topic models (trending area in Opinion Mining research)]
- **None of them is a solved problem!**

Opinion Mining Tasks

- At the document (or review) level:

Task: sentiment classification of reviews (overall polarity)

Classes: positive, negative, and neutral

Assumption: each document (or review) focuses on a single entity (not true in many discussion posts) and contains opinion from a single opinion holder.

- At the sentence level:

Task 1: identifying subjective/opinionated sentences

Classes: objective and subjective (opinionated)

Task 2: sentiment classification of sentences

Classes: positive, negative and neutral.

Assumption: a sentence contains only one opinion

not true in many cases.

Task 3: discover phrasal sentiment/issues (the wifi *signal dropped frequently* in the router)

Opinion Mining Tasks (contd.)

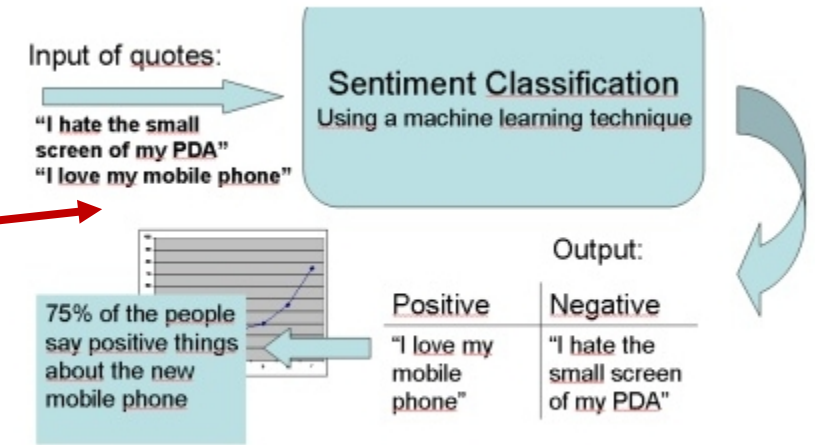
- **Some important NLP tasks at the aspect level:**

Task 1 (aspect extraction and grouping): Extract all aspect expressions of the entities, and group synonymous aspect expressions into clusters. Each aspect expression cluster of entity e_i indicates a unique aspect a_{ij} .

Task 2 (aspect sentiment classification): Determine whether each opinion on an aspect is positive, negative or neutral.

Task 3 (discover aspect specific sentiments): groupings of opinion words usually used to cast opinions/sentiments on a aspect (category)

Task 4 (opinion quintuple generation): Produce all opinion quintuples $(e_i, a_{ij}, o_{ijkl}, h_k, t_l)$ expressed in D .



- ❑ Many more specific tasks, but the above are usually key tasks.
- ❑ Also each of the task can be extended to include phrases (ngrams)

Opinion Mining Tasks (contd.)

- Some important NLP tasks at the aspect level:

Task 1 (aspect extraction and grouping): Extract all aspect expressions of the entities, and group synonymous aspect expressions into clusters. Each aspect expression cluster of entity e_i indicates a unique aspect a_{ij} .

Task 2 (aspect sentiment classification): Determine whether each opinion on an aspect is positive, negative or neutral.

Task 3 (discover aspect specific sentiments): groupings of opinion words usually used to cast opinions/sentiments on a aspect (category)

Task 4 (opinion quintuple generation): Produce all opinion quintuples $(e_i, a_{ij}, oo_{ijkl}, h_k, t_l)$ expressed in D .

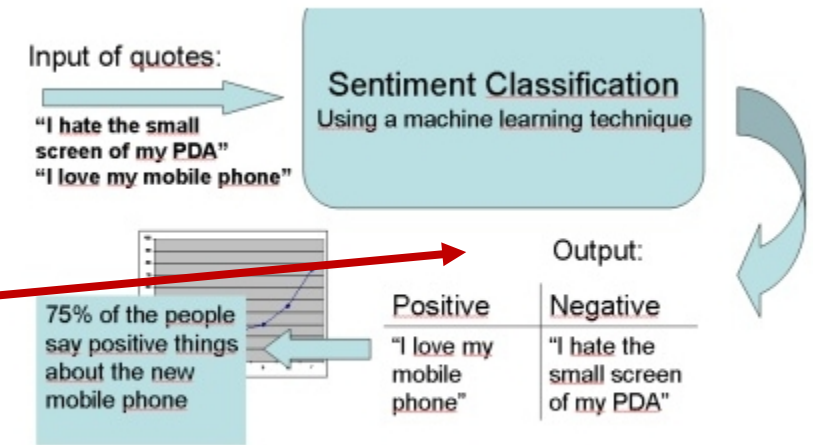
❑ Many more specific tasks, but the above are usually key tasks.

❑ Also each of the task can be extended to include phrases (ngrams)

Aspect (seeds)	ME-SAS	
	Aspect	Sentiment
Staff (staff service waiter hospitality upkeep)	attendant manager waitress maintenance bartender waiters housekeeping receptionist waitstaff janitor	friendly attentive polite nice clean pleasant slow courteous rude professional
Cleanliness (curtains restroom floor beds cleanliness)	carpets hall towels bathtub couch mattress linens wardrobe spa pillow	clean dirty comfortable fresh wet filthy extra stain front worn
Comfort (comfort mattress furniture couch pillows)	bedding bedcover sofa linens bedroom suites décor comforter blanket futon	comfortable clean soft nice uncomfortable spacious hard comfy dirty quiet

Opinion Mining: Document level classification

- Sentiment classification [Pang et al., 2002]
- Classify documents (e.g., reviews) based on the overall sentiments expressed by opinion holders (authors),
Positive, negative, and (possibly) neutral
Since in our model an entity e itself is also a aspect, then sentiment classification essentially determines the opinion expressed on e in each document (e.g., review).
- Similar but \neq topic-based text classification.
In topic-based text classification, topic words are important.
In sentiment classification, sentiment words are more important, e.g., great, excellent, horrible, bad, worst, etc.



Q: How would one implement this using a classifier?

Q: What are the features? Classes? E.g., using SVM/NB?

Opinion Mining: Document level classification

- Sentiment regression [\[Titov and McDonald, '08\]](#),
[\[Snyder, '07\]](#), [\[Blei and McAuliffe, \]](#)

- Regress documents (e.g., reviews) based on the overall sentiments expressed with the star rating,
Discover which terms indicate what kind of polarity.

Nikos' Fine Dining

Food	4/5	"Best fish in the city", "Excellent appetizers"
Decor	3/5	"Cozy with an old world feel", "Too dark"
Service	1/5	"Our waitress was rude", "Awful service"
Value	5/5	"Good Greek food for the \$", "Great price!"

Figure 1: An example aspect-based summary.

- Similar but \neq text regression.
Need to weigh sentiment/opinion words differently than other words.,
e.g., great, excellent -> 4, 5 star
VS.
terrible, horrible, bad, worst, etc. -> 1, 2 star
- **Q: How to regress star rating with text?**

★☆☆☆☆ Keyboard and mouse make this a terrible computer, March 9, 2014

By [Sasha](#) - [See all my reviews](#)

This review is from: **Lenovo X1 Carbon 14-Inch Touchscreen Ultrabook (20A7002QUS) Black (Personal Computers)**

Lenovo culled my negative review of the computer, so I'm opting to post on Amazon instead. The adaptive keyboard is a terrible idea. Every new layout is a major loss of productivity on initial purchase and then another moderate loss every time you move back and forth between various computers. It doesn't accomplish anything useful and feels like change for change's sake. A novelty at best.

On top of the terrible layout, the clickpad is miserable. The original X1 Carbon was OK - at least it had real buttons for the wiggle stick and so you could use the nub productively. The lack of physical buttons on the new X1 Carbon make the stick useless. You can't know which button you'll press, or if you're going to launch your cursor into space.

Looking and touching the machine is a joy. It's beautiful and feels good. But as soon as you boot it up and try to do real work you'll wish you didn't get this machine. The original X1 Carbon is a much better computer, despite having bad battery life and a lower resolution screen. I'm hoping Lenovo about-faces on the input-device choices they've been making recently.

Help other customers find the most helpful reviews

Was this review helpful to you?

[Report abuse](#) | [Permalink](#)

Opinion Mining: Unsupervised Methods

- Unsupervised review classification [Turney, 2002]
- Data: reviews from epinions.com on automobiles, banks, movies, and travel destinations.
- The approach: Three steps
- Step 1:
 - Part-of-speech tagging
 - Extracting two consecutive words (**two-word phrases**) from reviews if their tags conform to some given patterns, e.g., (1) JJ, (2) NN.

Opinion Mining: Unsupervised Methods

- Step 2: Estimate the semantic orientation (SO) of the extracted phrases

Use Pointwise mutual information

Semantic orientation (SO):

$$SO(phrase) = PMI(phrase, "excellent") \\ - PMI(phrase, "poor")$$

Using AltaVista near operator to do search to find the number of hits to compute PMI and SO.

- Step 3: Compute the average SO of all phrases
classify the review as **recommended** if average SO is positive, **not recommended** otherwise.
- Final classification accuracy:
 - automobiles - 84%
 - banks - 80%
 - movies - 65.83
 - travel destinations - 70.53%

Opinion Mining: Supervised Methods

- Sentiment classification using machine learning methods (Pang et al, EMNLP-02)
- Directly applied supervised classifiers (SVM, NB, etc.) to classify movie reviews into positive and negative using labeled data.
- Three classification techniques were tried:
 - Naïve Bayes
 - Maximum entropy [same as Logistic Regression]
 - Support vector machine
- Pre-processing settings: negation tag, unigram (single words), bigram, POS tag, position.
- SVM: the best accuracy 83% (unigram)

Opinion Mining: Sentence Level Classification

- Document-level sentiment classification is too coarse for most applications.
- Let us move to the sentence level.
- Much of the work on sentence level sentiment analysis focuses on identifying **subjective sentences** in news articles.

Classification: objective and subjective.

All techniques use some forms of machine learning.

E.g., using a naïve Bayesian classifier with a set of data features/attributes extracted from training sentences (Wiebe et al. ACL-99).

Opinion Mining: Beyond Sentence Level Classification

Sentiment classification at both document and sentence (or clause) levels are useful, but

They do not find what the opinion holder liked and disliked.

i.e., we need to discover aspect specific sentiments appearing in sentences

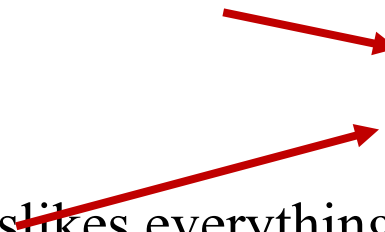
E.g., {nice->phone; clear->voice; long->battery}

An negative sentiment on an entity

does not mean that the opinion holder dislikes everything about the entity.

A positive sentiment on an entity

does not mean that the opinion holder likes everything about the entity.



*"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. ..."*

We need to go to the entity level.

But before we do that, having a opinion lexicon (annotated +/- words) is useful

Opinion Mining: Sentiment Lexicon

- **Opinion Words or Phrases** (also called polar words, opinion bearing words, etc). E.g.,
Positive: beautiful, wonderful, good, amazing,
Negative: bad, poor, terrible, cost someone an arm and a leg (idiom).
- They are instrumental for opinion mining (obviously)
- A comprehensive list of them is called a lexicon
- Three main ways to compile such a list:
Manual approach: not a bad idea, only an one-time effort
Corpus-based approaches [Hu and Liu, 2004]
Dictionary-based approaches [SentiWordNet, WordNet]
- **Important to note:**
Some opinion words are context independent (e.g., good).
Some are context dependent (e.g., long).
Some words have multiple senses (light has two senses weight or brightness)


Aspect Based Opinion Mining

- Again focus on reviews (easier to work in a concrete domain!)
- A major percentage of research in Opinion Mining/Sentiment Analysis focuses on user generated content (UGC) such as reviews, blogs, discussions. Why?
- UGC is actionable data (i.e., mined information useful for commercial applications)
- **Objective:** find what reviewers (opinion holders) liked and disliked
Product aspects and opinions on the aspects
- Since the number of reviews on an entity can be large, an opinion summary should be produced.
Desirable to be a structured summary.
Easy to visualize and to compare.
Analogous to but different from multi-document summarization.

Aspect Based Opinion Mining

- Task: Aspect Extraction. Q: What does it mean?
- Discover the green colored terms

"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. ..."



- Easier to discover frequent aspects (called features before):
those aspects that have been talked about by many reviewers [Hu and Liu, 2004].
- Use pattern mining [e.g., Apriori algorithm]
- Why the frequency based approach?
Different reviewers tell different stories (irrelevant)
When product aspects are discussed, the words that they use converge.
They are main aspects.
- Sequential pattern mining finds frequent phrases.
- Can further improve by finding frequent noun and noun phrases

Many companies implemented the approach (no POS restriction).

Aspect Based Opinion Mining

- Task: Aspect Extraction.
- Using part-of relationship and the Web
(Popescu and Etzioni, EMNLP-05)
- Improved (Hu and Liu, KDD-04) by removing those frequent noun phrases that may not be aspects: better precision (a small drop in recall).
- It identifies **part-of** relationship
 - ❑ Each noun phrase is given a pointwise mutual information score between the phrase and **part discriminators** associated with the product class, e.g., a scanner class.
 - ❑ The part discriminators for the scanner class are, “of scanner”, “scanner has”, “scanner comes with”, etc, which are used to find components or parts of scanners by searching on the Web: the KnowItAll approach, (Etzioni et al, WWW-04).

Aspect Based Opinion Mining

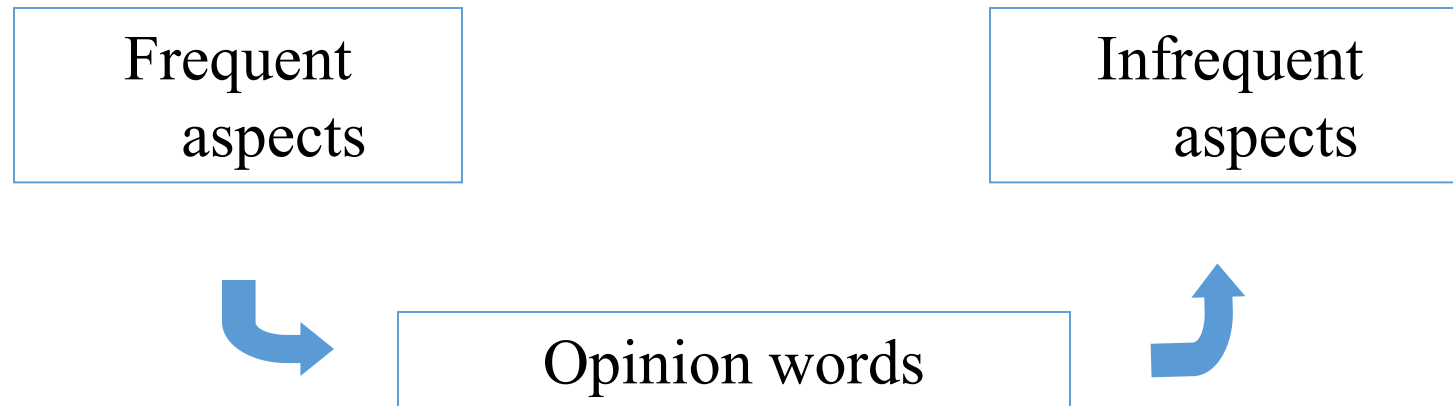
- Task: Aspect Extraction.

Q: How to discover infrequent but important product aspects?

- **Observation:** the same opinion word can be used to describe different aspects and entities.

“The pictures are absolutely **amazing**.”

“The software that comes with it is **amazing**.”



Aspect Based Opinion Mining

- Task: Aspect Extraction.
Q: Can we make use of topic modeling?
- **Observation:** A review usually contain evaluation of aspects which are semantically close/coherent. In the following hotel reviews:
“The **housekeeping staff** were rude. The **janitor** did not clean the bathroom well.”
“The **front desk receptionist** was great. **Lobby** was good too. The **lobby receptionist**...”
- Topic modeling employs higher-order-co-occurrence : “ w_1 co-occurring with w_2 which in turn co-occurs with w_3 denotes a second-order co-occurrence between w_1 and w_3 ” to discover synonymous words (aspects).
- Several topic modeling based approaches have been used for aspect extraction recently. [\[Lin and He, 2009\]](#), [\[Zhao et al., 2010\]](#), [\[Mukherjee and Liu, 2012\]](#)
- Very popular research topic.

Aspect Sentiment Classification


- For each aspect, we identify the sentiment or opinion orientation expressed by a reviewer.
- We work based on sentences, but also consider,
A sentence can contain multiple aspects.
Different aspects may have different opinions.
E.g., The **battery life** and **picture quality** are *great* (+), but the **view finder** is *small* (-).
- Almost all approaches make use of **opinion words and phrases**. But notice **again**:
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 them.

Public opinion in this country is everything.

—Abraham Lincoln

The Opinion Curse – Opinion Spam!

- Opinions \Rightarrow Virtual Currency



- Opinion Spam: Illegitimate activities (e.g., writing fake reviews/ratings) to deliberately mislead consumers
- In E-commerce, filtering opinion spam is vital



Which review is fake?

I want to make this review in order to comment on the excellent service that my mother and I received on the Serenade of the Seas, a cruise line for Royal Caribbean. There was a lot of things to do in the morning and afternoon portion for the 7 days that we were on the ship. We went to 6 different islands and saw some amazing sites! It was definitely worth the effort of planning beforehand. The dinner service was 5 star for sure. I recommend the Serenade to anyone who is looking for excellent service, excellent food, and a week full of amazing day-activities!

VS

Guacamole burger was quite tall; clam chowder was tasty. The appetizers weren't very good at all. And the service kind of lagged. A cross between Las Vegas and Disney world, but on the cheesy side. This Cafe is a place where you eat inside a plastic rain forest. The walls are lined with fake trees, plants, and wildlife, including animatronic animals. I could see it being fun for a child's birthday party (there were several that occurred during our meal), but not a place to go if you're looking for a good meal.

Which review is fake?

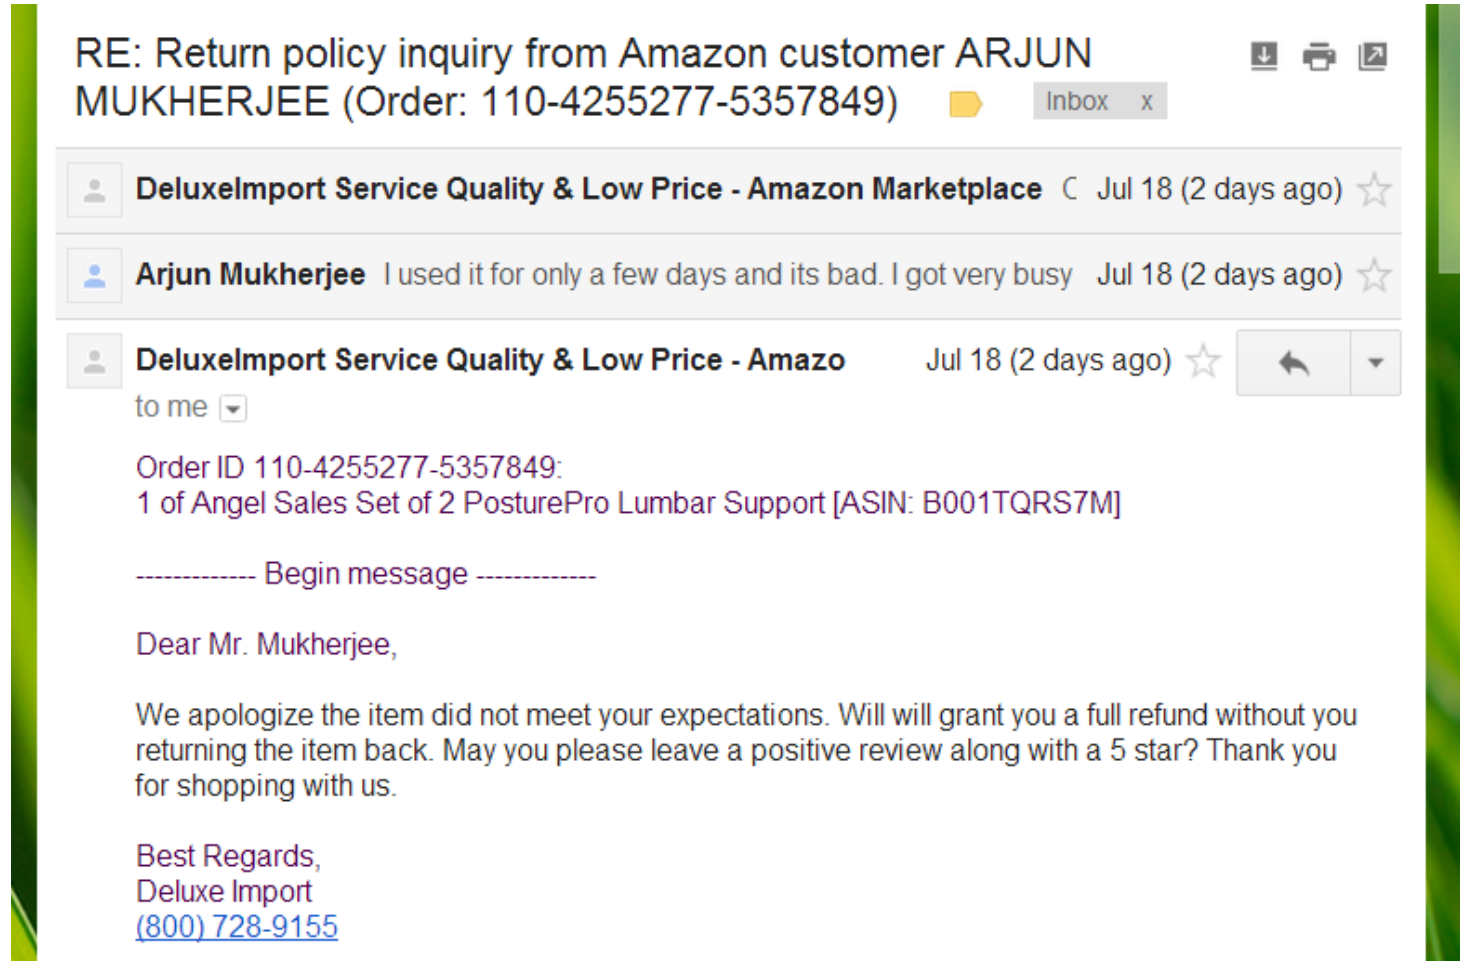
I want to make this review in order to comment on the excellent service that my mother and I received on the Serenade of the Seas, a cruise line for Royal Caribbean. There was a lot of things to do in the morning and afternoon portion for the 7 days that we were on the ship. We went to 6 different islands and saw some amazing sites! It was definitely worth the effort of planning beforehand. The dinner service was 5 star for sure. I recommend the Serenade to anyone who is looking for excellent service, excellent food, and a week full of amazing day-activities!

VS

Guacamole burger was quite tall; clam chowder was tasty. The appetizers weren't very good at all. And the service kind of lagged. A cross between Las Vegas and Disney world, but on the cheesy side. This Cafe is a place where you eat inside a plastic rain forest. The walls are lined with fake trees, plants, and wildlife, including animatronic animals. I could see it being fun for a child's birthday party (there were several that occurred during our meal), but not a place to go if you're looking for a good meal.

Opinion Spam Solicitations (Real Cases)

- I wanted to return an item I purchased from Amazon because it didn't work.
- Guess what did the seller say?



Opinion Spam Solicitations (Real Cases)

- The case of 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.

Opinion Spam Research

- Fake/untruthful reviews:
Write undeserving positive/negative reviews to promote/damage companies reputation.
- Fake reviews (biased reviews, paid reviews) are a big problem. Many cases in news. Filtering them is a pressing issue.
- First introduced in [\[Jindal and Liu, 2008\]](#) – used duplicate reviews as fake and others as non-fake.
- Using Amazon Mechanical Turk for generating ground truth fake reviews [\[Ott et al., 2011\]](#)
- Detecting group opinion spam [\[Mukherjee et al., 2012\]](#).
- Review spammer detection using burstiness modeling [\[Fei et al., 2013\]](#). Many more...
- A very popular topics of research.