Opinion Mining

Roberto Basili, Danilo Croce, Giuseppe Castellucci

WM&R
a.a. 2014/2015

main contribution from "Opinion Mining" by Bing Liu (Chpt. 11)
and "Opinion Mining and Sentiment Analysis" by B. Pang & L. Lee

Artificial Intelligence
at Roma Tor Vergata

SAG@RT
Semantic Analytics Group at
Artificial Intelligence Rome TorVergata

Outline

- Why Opinion Mining (OM)?
- What is OM?
- OM: Applications and Opportunities
- Technological challenges for OM
  - Levels of abstraction and Resources
  - Computational paradigms and Portability
  - Compositionality
- Active experiences
- Conclusions and Future Directions

A Web of people and opinions

- **31.7%** of the more than 200 million bloggers worldwide blog about opinions on products and brands (Universal McCann, July 2009)
- **71%** of all active Internet users read blogs.
- 2009 Survey of **25,000** Internet users in **50** countries: **70%** of consumers trust opinions posted online by other consumers (Nielsen Global Online Consumer, 2010).

Social Media Analytics

- Complex process for Social Media Analytics are necessary whereas …
- … Opinion Mining and Sentiment Analysis play a crucial role

Authority

- Does the opinion of one user (e.g. a blogger) actually matter?
- "*If a tree falls in a forest and no one is around to hear it, does it make a sound?*"
- Authority and reputation of users are key factors to understand and account for their opinions

What is OM?

- *Opinion Mining* or also *sentiment analysis* is **the computational study of opinions, sentiments and emotions expressed in text**
- How to model, code and compute the irrational aspects of our affects in an analytical way …
- It deals with rational models of emotions, rumors and trends within user communities
- … and with the word-of-mouth inside specific domains

What is OM? (2)

- Opinion Mining or Sentiment Analysis involve more than one linguistic task
- What is the *opinion* of a text
  - Who is author (or *opinion holder*, OH)
  - What is the *opinion target* (Object)
  - What are the *features* of the Object
  - What is the *subjective position* of the user wrt to the Object or the individual features
- What about the (dynamics of) opinions of large OH communities

Introduction – facts and opinions

- Two main types of information on the Web.
  - Facts and Opinions
- Current search engines search for facts (assume they are true)
  - Facts can be expressed with topic keywords.
- Search engines do not search for opinions
  - Opinions are hard to express with a few keywords
    - How do people think of Motorola Cell phones?
  - Current search ranking strategy is not appropriate for opinion retrieval/search.

Introduction – 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
  - Web/global scale
    - No longer limited to your circle of friends
- Our interest: to mine opinions expressed in the user-generated content
  - An intellectually very challenging problem.
  - Practically very useful.

Opinion search (Liu, Web Data Mining book, 2007)

- Can you search for opinions as conveniently as general Web search?
- Whenever you need to make a decision, you may want some opinions from others,
  - Wouldn't it be nice? you can find them on a search system instantly, by issuing queries such as
    - Opinions: "Motorola cell phones"
    - Comparisons: "Motorola vs. Nokia"
- Cannot be done yet!

Two types of evaluation

- Direct Opinions: sentiment expressions on some objects, 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 object. Usually expressing an ordering.
  - E.g., "car x is cheaper than car y."
  - Objective or subjective.
  - We will not cover in the class (read the textbook if you are interested)

Opinion Summarization through Visual Comparison (Liu et al. WWW-2005)

- Summary of reviews of Cell Phone 1

- Comparison of reviews of

  Cell Phone 1

  Cell Phone 2

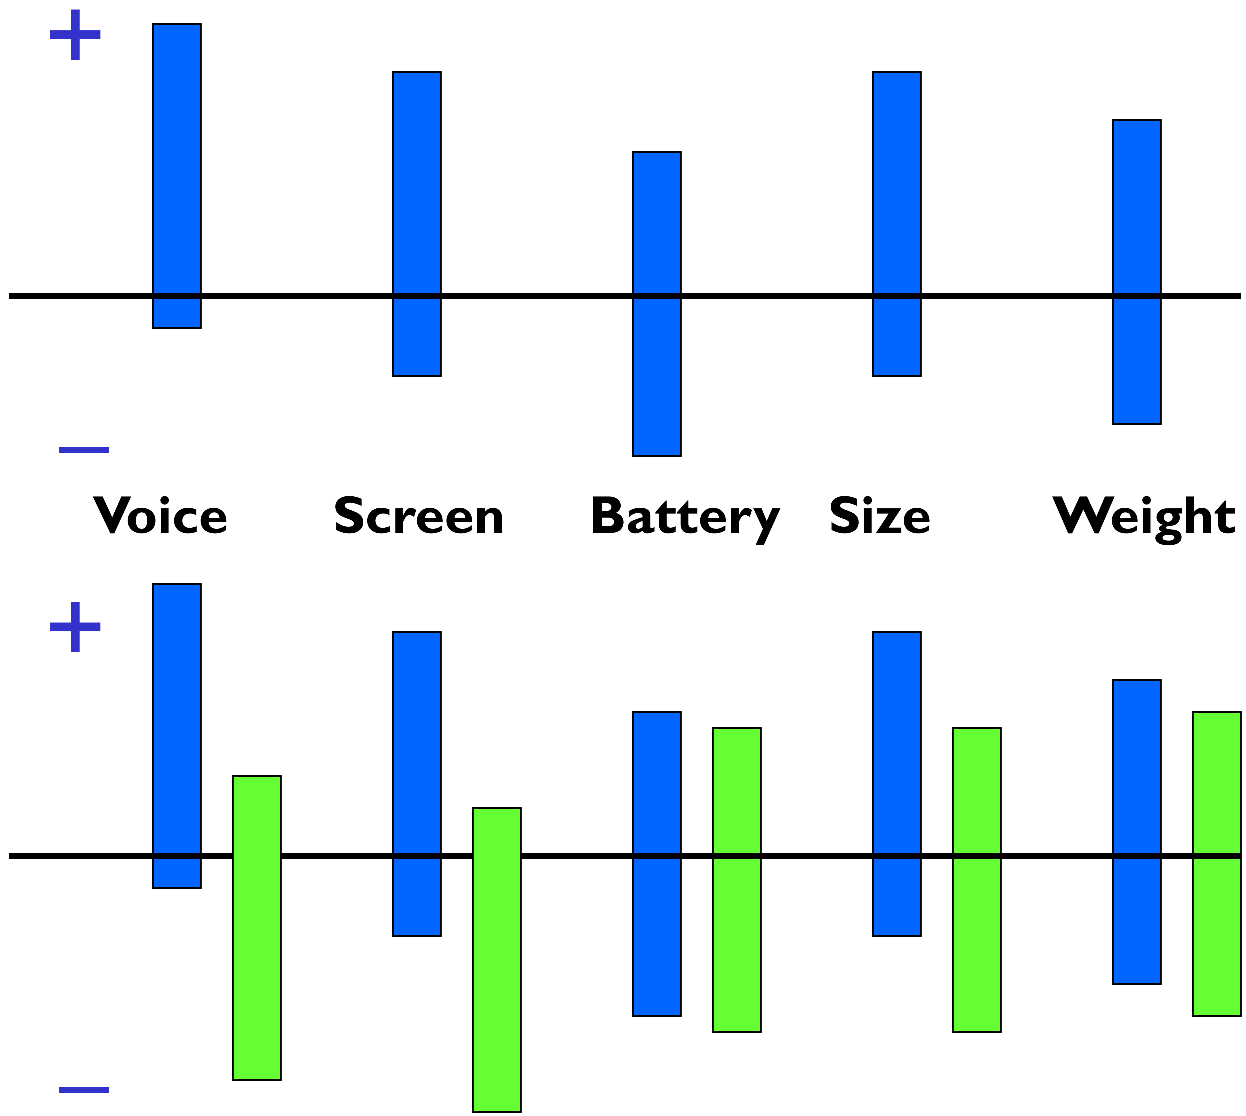

Find the opinion of a person on X

- In some cases, the general search engine can handle it, i.e., using suitable keywords.
  - Bill Clinton's opinion on abortion
- Reason:
  - One person or organization usually has only one opinion on a particular topic.
  - The opinion is likely contained in a single document.
  - Thus, a good keyword query may be sufficient.

Find opinions on an object X

We use product reviews as an example:

- Searching for opinions in product reviews is different from general Web search.
  - E.g., search for opinions on "Motorola RAZR V3"
- General Web search for a fact: rank pages according to some authority and relevance scores.
  - The user views the first page (if the search is perfect).
  - **One fact = Multiple facts**
- Opinion search: rank is desirable, however
  - reading only the review ranked at the top is dangerous because it is only the opinion of one person.
  - **One opinion ≠ Multiple opinions**

Search opinions (contd)

- Ranking:
  - produce two rankings
    - Positive opinions and negative opinions
    - Some kind of summary of both, e.g., # of each
  - Or, one ranking but
    - The top (say 30) reviews should reflect the natural distribution of all reviews (assume that there is no spam), i.e., with the right balance of positive and negative reviews.
- Questions:
  - Should the user reads all the top reviews? OR
  - Should the system prepare a summary of the reviews?

Reviews are similar to surveys

- Reviews can be regarded as traditional surveys.
  - In traditional survey, returned survey 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 search,
  - Can a summary be produced?
  - What should the summary be?

Features: opinions vs. mentions

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

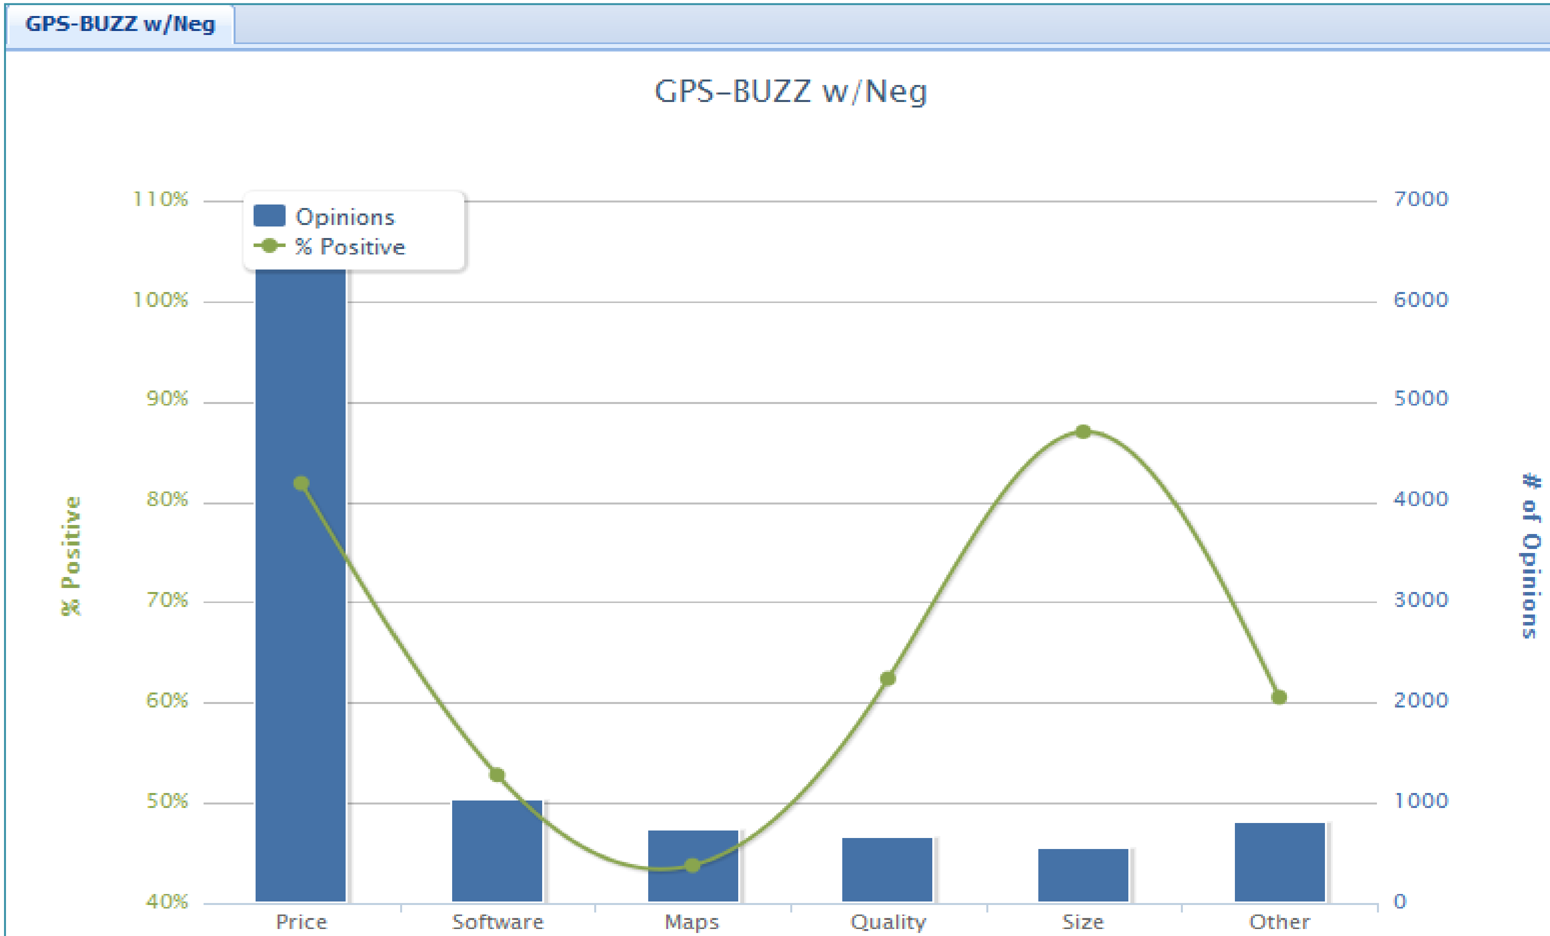

- It seems very appealing

but…

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."*

… and it could be a very complex task!!

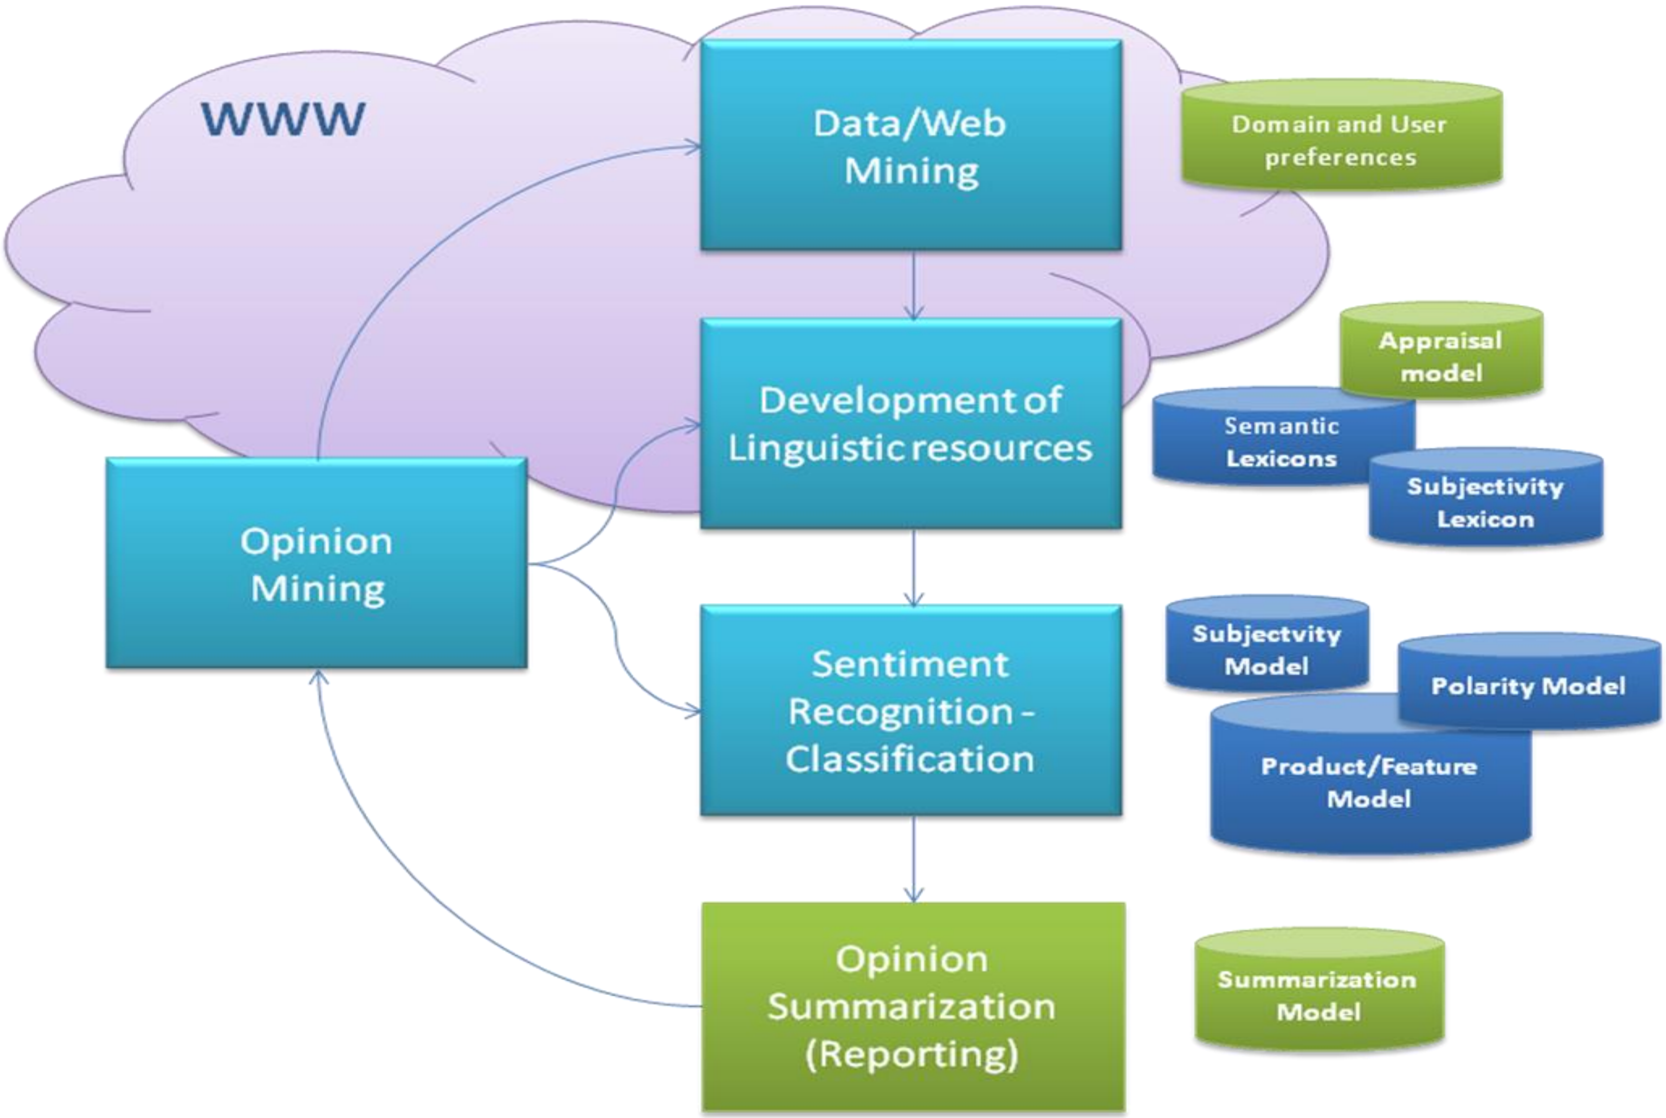

NL vs. Opinions

- Although subjectivity seems to preserve across domains and sublanguages, *subjectivity (or affective) lexicons are not fully portable*
  - Often polarity of some terms change across domains (e.g. *small* in laptops vs. TV screen)
- These problems triggers a number of inductive tasks
  - How to *model the uncertainty* of lexical information with respect to subjectivity
  - How to *validate (or adapt) existing lexicons* to newer domains
  - How to *acquire novel lexical information*
  - How to *support inference* according to the above lexical information

NL vs. Opinions

- Opinions can be treated as *uncertain events expressed* by a text such that …
- a modeling similar to Information Extraction tasks seems appropriate
- Machine Learning has been largely used in OM
- Sentiment Analysis has been mapped into a text classification task (see *genre* class.)
  - Subjectivity recognition
  - Polarity Assignment

Roadmap

- **Opinion mining – the abstraction**
- Domain level sentiment classification
- Sentence level sentiment analysis
- Feature-based sentiment analysis and summarization
- Summary

Opinion mining – the **abstraction**

(Hu and Liu, KDD-04)

- Basic components of an opinion
  - Opinion holder: A person or an organization that holds an specific opinion on a particular object.
  - Object: on which an opinion is expressed
  - Opinion: a view, attitude, or appraisal on an object from an opinion holder.
- Objectives of opinion mining: many ...
- We use consumer reviews of products to develop the ideas. Other opinionated contexts are similar.

Object/entity

- **Definition** (**object**): An object *O* is an entity which can be a product, person, event, organization, or topic. *O* is represented as a tree or taxonomy of components (or parts), sub-components, and so on.
  - Each node represents a component and is associated with a set of **attributes**.
  - *O* is the root node (which also has a set of attributes)
- An opinion can be expressed on any node or attribute of the node.
- To simplify our discussion, we use "**features**" to represent both components and attributes.
  - The term "feature" should be understood in a ***broad sense***,
    - Product feature, topic or sub-topic, event or sub-event, etc
- Note: the object *O* itself is also a feature.

A model of a review

- An object is represented with a finite set of features, $F = \{f_1, f_2, \ldots, f_n\}$.
  - Each feature f_i in F can be expressed with a finite set of words or phrases W_i, which are **synonyms**.

  That is to say: we have a set of corresponding synonym sets $W = \{W_1, W_2, \ldots, W_n\}$ for the features.
- **Model of a review**: An opinion holder j comments on a subset of the features $S_j \subseteq F$ of an object O.
  - For each feature $f_k \in S_j$ that j comments on, he/she
    - chooses a word or phrase from W_k to describe the feature, and
    - expresses a positive, negative or neutral opinion on f_k.

Opinion mining tasks

- At the document (or review) level:

  Task: sentiment classification of reviews

  - Classes: positive, negative, and neutral
  - Assumption: each document (or review) focuses on a single object *O* (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.
  - Then we can also consider clauses.

Opinion mining tasks (contd)

- At the feature level:

  Task 1: Identifying and extracting object features that have been commented on in each review.

  Task 2: Determining whether the opinions on the features are positive, negative or neutral in the review.

  Task 3: Grouping feature synonyms.

  - Produce a feature-based opinion summary of multiple reviews (more on this later).
- Opinion holders: identify holders is also useful, e.g., in news articles, etc, but they are usually known in user generated content, i.e., the authors of the posts.

More at the feature level

F: the set of features

W: synonyms of each feature

- **Problem 1:** Both *F* and *W* are unknown.
  - We need to perform all three tasks:
- **Problem 2:** *F* is known but *W* is unknown.
  - All three tasks are needed. Task 3 is easier. It becomes the problem of matching discovered features with the set of given features *F*.
- **Problem 3:** *W* is known (*F* is known too).
  - Only task 2 is needed.

Roadmap

- Opinion mining – the abstraction
- **Document level sentiment classification**
- Sentence level sentiment analysis
- Feature-based sentiment analysis and summarization
- Summary

Sentiment classification

- Classify documents (e.g., reviews) based on the overall sentiments expressed by authors,
  - Positive, negative, and (possibly) neutral
  - Since in our model an object *O* itself is also a feature, then sentiment classification essentially determines the opinion expressed on *O* in each document (e.g., review).
- Similar but not identical to *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.

Unsupervised review classification
(Turney, ACL-02)

- 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.

- Step 2: Estimate the semantic orientation of the extracted phrases
  - Use Pointwise mutual information

$$PMI(word_1, word_2) = \log_2\left(\frac{P(word_1 \wedge word_2)}{P(word_1)P(word_2)}\right)$$

  - 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 2: Estimate the semantic orientation of the extracted phrases

- Use Pointwise mutual information Semantic orientation (SO):
-

$$SO(phrase) = PMI(phrase, \text{“excellent”}) - PMI(phrase, \text{“poor”})$$

$$SO(phrase) = \log_2\left(\frac{hits(phrase \text{ NEAR “excellent”})\, hits(\text{“poor”})}{hits(phrase \text{ NEAR “poor”})\, hits(\text{“excellent”})}\right)$$

- 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%

Sentiment classification using machine learning methods (Pang et al, EMNLP-02)

- The paper applied several machine learning techniques to classify movie reviews into positive and negative.
- Three classification techniques were tried:
  - Naïve Bayes
  - *Maximum entropy (mixture model + Par Est)*
  - Support vector machine
- Pre-processing settings: negation tag, unigram (single words), bigram, POS tag, position.
- SVM: the best accuracy 83% (unigram)

Roadmap

- Opinion mining – the abstraction
- Document level sentiment classification
- **Sentence level sentiment analysis**
- Feature-based sentiment analysis and summarization
- Summary

Sentence-level sentiment analysis

- 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 focus 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).

Using learnt patterns (Rilloff and Wiebe, EMNLP-03)

- A bootstrapping approach.
  - A high precision classifier is used to automatically identify some subjective and objective sentences.
    - Two high precision (low recall) classifiers were used,
      - a high precision subjective classifier
      - A high precision objective classifier
      - Based on manually collected lexical items, single words and n-grams, which are good subjective clues.
  - A set of patterns are then learned from these identified subjective and objective sentences.
    - Syntactic templates are provided to restrict the kinds of patterns to be discovered, e.g., <subj> passive-verb.
  - The learned patterns are then used to extract more subject and objective sentences (the process can be repeated).

Subjectivity and polarity (orientation)
(Yu and Hazivassiloglou, EMNLP-03)

- For subjective or opinion sentence identification, three methods was tried:
  - Sentence similarity.
  - Naïve Bayesian classification.
  - Multiple naïve Bayesian (NB) classifiers.
- For opinion orientation (positive, negative or neutral) (also called polarity) classification, it uses a similar method to (Turney, ACL-02), but
  - with more seed words (rather than two) and based on log-likelihood ratio (LLR).
  - For classification of each word, it takes average of LLR scores of words in the sentence and use cutoffs to decide positive, negative or neutral.

Let us go further?

- Sentiment classifications at both document and sentence (or clause) level are useful, but
  - They do not find what the opinion holder liked and disliked.
- A negative sentiment on an object
  - does not mean that the opinion holder dislikes everything about the object.
- A positive sentiment on an object
  - does not mean that the opinion holder likes everything about the object.
- **We need to go to the feature level.**

But before we go further

- Many approaches to opinion, sentiment, and subjectivity analysis rely on **lexicons** of words that may be used to express subjectivity.

(1) He is a **disease** to every team he has gone to.
Converting to SMF is a **headache**.
The concert left me **cold**.
That guy is such a **pain**.

(2) Early symptoms of the **disease** include severe **headaches**, red eyes, fevers and **cold** chills, body **pain**, and vomiting.

But before we go further

- Let us discuss **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)
- Three main ways to compile such a list:
  - Manual approach: not a bad idea, only an one- time effort
  - Corpus-based approaches
  - Dictionary-based approaches
- **Important to note:**
  - **Some opinion words are context independent.**
  - **Some are context dependent.**

Corpus-based approaches

- Rely on syntactic or co-occurrence patterns in large corpuses. (Hazivassiloglou and McKeown, ACL-97; Turney, ACL-02; Yu and Hazivassiloglou, EMNLP-03; Kanayama and Nasukawa, EMNLP-06; Ding and Liu, 2007)
  - Can find domain (not context) dependent orientations (positive, negative, or neutral).
- (Turney, ACL-02) and (Yu and Hazivassiloglou, EMNLP-03) are similar.
  - Assign opinion orientations (polarities) to words/phrases.
  - (Yu and Hazivassiloglou, EMNLP-03) is different from (Turney, ACL-02) in that
    - using more seed words (rather than two) and using log-likelihood ratio (rather than PMI).

Corpus-based approaches (contd)

- Use constraints (or conventions) on connectives to identify opinion words (Hazivassiloglou and McKeown, ACL-97; Kanayama and Nasukawa, EMNLP-06; Ding and Liu, SIGIR-07). E.g.,
  - Conjunction: conjoined adjectives usually have the same orientation (Hazivassiloglou and McKeown, ACL-97).
    - E.g., "This car is *beautiful* **and** *spacious*." (conjunction)
  - AND, OR, BUT, EITHER-OR, and NEITHER-NOR have similar constraints
- Learning using
  - log-linear model: determine if two conjoined adjectives are of the same or different orientations.
  - Clustering: produce two sets of words: positive and negative
- Corpus: 21 million word 1987 Wall Street Journal corpus.

Dictionary-based approaches

- Typically use WordNet's synsets and hierarchies to acquire opinion words
  - Start with a small seed set of opinion words
  - Use the set to search for synonyms and antonyms in WordNet (Hu and Liu, KDD-04; Kim and Hovy, COLING-04).
  - Manual inspection may be used afterward.
- Use additional information (e.g., glosses) from WordNet (Andreevskaia and Bergler, EACL-06) and learning (Esuti and Sebastiani, CIKM-05).
- Weakness of the approach: Do not find domain and/or context dependent opinion words, e.g., small, long, fast.

OM resources: SentiWordnet

- SentiWN (Sebastiani & Esuli, 2008)

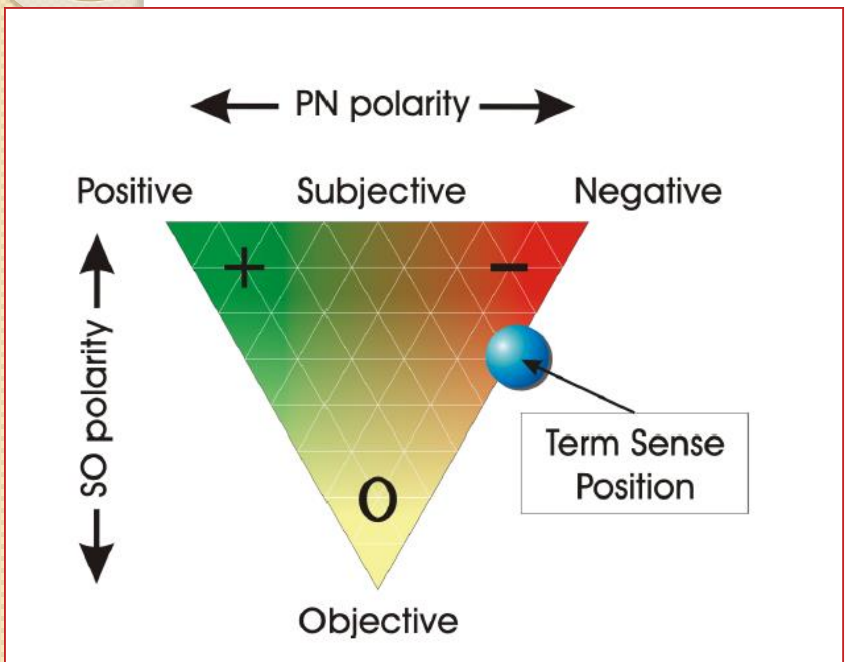

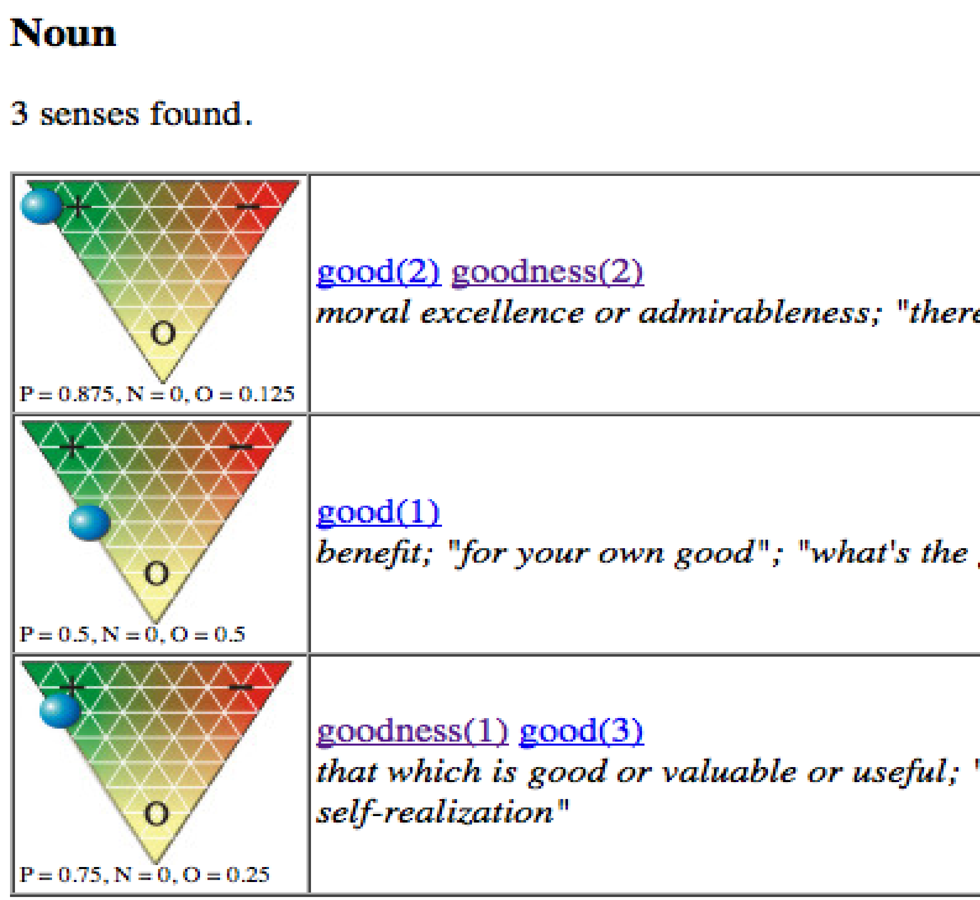

NCSR Lexicon (Mohammad & Turney, 2013)

- Saif Mohammad and Peter D. Turney. 2013. *Crowd-sourcing a word-emotion association lexicon*. Computational Intelligence, 29(3):436–465.

Term	positive	negative	anger	anticipation	disgust	fear	joy	sadness	surprise	trust
agitated	0	1	1	0	0	0	0	0	0	0
agitation	0	1	1	0	0	0	0	0	0	0
agnostic	0	0	0	0	0	0	0	0	0	0
ago	0	0	0	0	0	0	0	0	0	0
agonizing	0	1	0	0	0	1	0	0	0	0
agony	0	1	1	0	0	1	0	1	0	0
agree	1	0	0	0	0	0	0	0	0	0
agreeable	1	0	0	0	0	0	0	0	0	1
agreed	1	0	0	0	0	0	0	0	0	1
agreeing	1	0	0	0	0	0	0	0	0	1
agreement	1	0	0	0	0	0	0	0	0	1
agricultural	0	0	0	0	0	0	0	0	0	0
agriculture	1	0	0	0	0	0	0	0	0	0
aground	0	1	0	0	0	0	0	0	0	0
agua	0	0	0	0	0	0	0	0	0	0
ahead	1	0	0	0	0	0	0	0	0	0
aid	1	0	0	0	0	0	0	0	0	0
aiding	1	0	0	0	0	0	0	0	0	0

Roadmap

- Opinion mining – the abstraction
- Document level sentiment classification
- Sentence level sentiment analysis
- **Feature-based sentiment analysis and summarization**
- Summary

Feature-based opinion mining and summarization (Hu and Liu, KDD-04)

- Again focus on reviews (easier to work in a concrete domain!)
- Objective: find what reviewers (opinion holders) liked and disliked
  - Product features and opinions on the features
- Since the number of reviews on an object can be large, an opinion summary should be produced.
  - Desirable to be a structured summary.
  - Easy to visualize and to compare.
  - Analogous to multi-document summarization.

The tasks

- Recall the three tasks in our model.

  Task 1: Extracting object features that have been commented on in each review.

  Task 2: Determining whether the opinions on the features are positive, negative or neutral.

  Task 3: Grouping feature synonyms.

  - Summary
- Task 2 may not be needed depending on the format of reviews.

Different review format

- Format 1 - Pros, Cons and detailed review: The reviewer is asked to describe Pros and Cons separately and also write a detailed review. Epinions.com uses this format.
- Format 2 - Pros and Cons: The reviewer is asked to describe Pros and Cons separately. C|net.com used to use this format.
- Format 3 - free format: The reviewer can write freely, i.e., no separation of Pros and Cons. Amazon.com uses this format.

Format 1

My SLR is on the shelf

by camerafun4. Aug 09 '04

Pros: Great photos, easy to use, very small
Cons: Battery usage; included memory is stingy.

I had never used a digital camera prior to purchasing th
have always used a SLR ... **Read the full review**

Format 2

User rating
Perfect
10
out of 10

"It is a great digitbal still camera for this century"
September 1, 2004

Pros:
It's small in size, and the rotatable lens is great. It's very easy to use, and has fast response from the shutter. The LCD has increased from 1.5 in to 1.8, which gives bigger view. It has lots of modes to choose from in order to take better pictures.

Cons:
It almost has no cons, it would be better if the LCD is bigger and it's going to be best if the model is designed to a smaller size.

Format 3

GREAT Camera., Jun 3, 2004
Reviewer: **jprice174** from Atlanta, Ga.

I did a lot of research last year before I bought this camera... It kinda hurt to leave behind my beloved nikon 35mm SLR, but I was going to Italy, and I needed something smaller, and digital. The pictures coming out of this camera are amazing. The 'auto' feature takes great pictures most of the time. And with digital, you're not wasting film if the picture doesn't come out.

Feature-based Summary (Hu and Liu, KDD-04)

GREAT Camera., Jun 3, 2004

Reviewer: **jprice174** from Atlanta, Ga.

I did a lot of research last year before I bought this camera... It kinda hurt to leave behind my beloved nikon 35mm SLR, but I was going to Italy, and I needed something smaller, and digital.

The pictures coming out of this camera are amazing. The 'auto' feature takes great pictures most of the time. And with digital, you're not wasting film if the picture doesn't come out. …

….

Feature Based Summary:

Feature1: **picture**

Positive: 12

- The pictures coming out of this camera are amazing.
- Overall this is a good camera with a really good picture clarity.

…

Negative: 2

- The pictures come out hazy if your hands shake even for a moment during the entire process of taking a picture.
- Focusing on a display rack about 20 feet away in a brightly lit room during day time, pictures produced by this camera were blurry and in a shade of orange.

Feature2: **battery life**

…

Feature extraction from Pros and Cons of Format 1 (Liu et al WWW-03; Hu and Liu, AAAI-CAAW-05)

- **Observation**: Each sentence segment in Pros or Cons contains only one feature. Sentence segments can be separated by commas, periods, semi-colons, hyphens, '&''s, 'and''s, 'but''s, etc.

- Pros in Example 1 can be separated into 3 segments:

great photos <photo>
easy to use <use>
very small <small> $\Rightarrow$ <size>

- Cons can be separated into 2 segments:

battery usage <battery>
included memory is stingy <memory>

Extraction using label sequential rules

- Label sequential rules (LSR) are a special kind of sequential patterns, discovered from sequences.
- LSR Mining is supervised (Liu's Web mining book 2006).
- The training data set is a set of sequences, e.g.,

  "*Included memory is stingy*"

  is turned into a sequence with POS tags.

  ⟨{included, VB}{memory, NN}{is, VB}{stingy, JJ}⟩

  then turned into

  ⟨{included, VB}{$feature, NN}{is, VB}{stingy, JJ}⟩

Using LSRs for extraction

- Based on a set of training sequences, we can mine label sequential rules, e.g.,

$$\langle\{easy, JJ\}\{to\}\{*, VB\}\rangle \rightarrow \langle\{easy, JJ\}\{to\}\{\$feature, VB\}\rangle$$

[confidence = 95%]

Feature Extraction

- Only the right hand side of each rule is needed.
- The word in the sentence segment of a new review that matches $feature is extracted.

Extraction of features of formats 2 and 3

- Reviews of these formats are usually complete sentences

  e.g., "the pictures are very clear."
  - Explicit feature: picture
- "It is small enough to fit easily in a coat pocket or purse."
  - Implicit feature: size
- Extraction: Frequency based approach
  - Frequent features
  - Infrequent features

Frequency based approach
(Hu and Liu, KDD-04)

- Frequent features: those features that have been talked about by many reviewers.
- Use sequential pattern mining
- Why the frequency based approach?
  - Different reviewers tell different stories (irrelevant)
  - When product features are discussed, the words that they use converge.
  - They are main features.
- Sequential pattern mining finds frequent phrases.
- Froogle has an implementation of the approach (no POS restriction).

Infrequent features extraction

- How to find the infrequent features?
- Observation: the same opinion word can be used to describe different features and objects.
  - "The pictures are absolutely amazing."
  - "The software that comes with it is amazing."

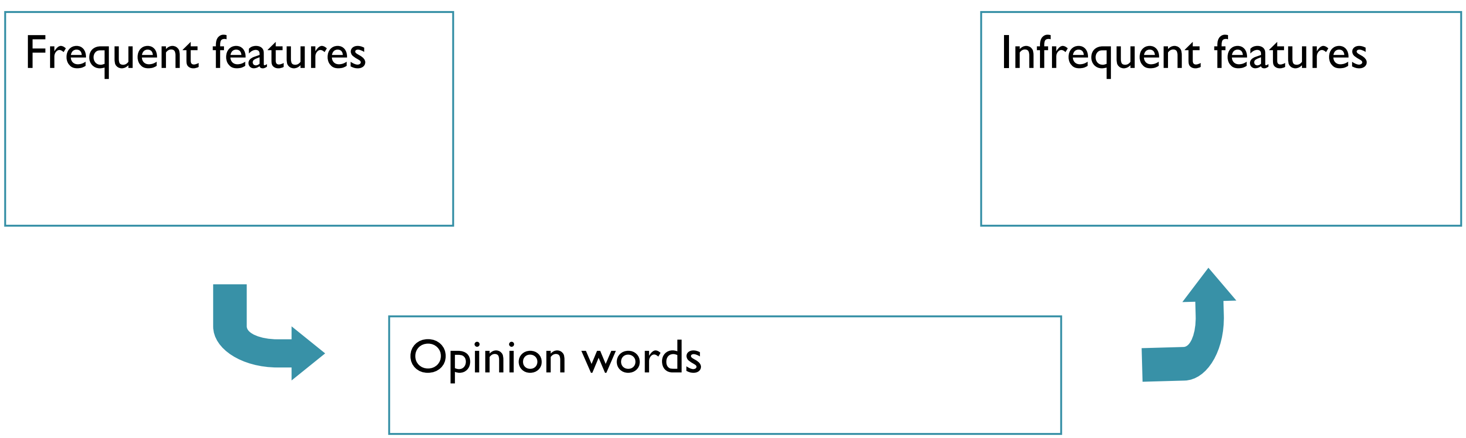

Identify feature synonyms

- Liu et al (WWW-05) made an attempt using only WordNet.
- Carenini et al (K-CAP-05) proposed a more sophisticated method based on several similarity metrics, but it requires a taxonomy of features to be given.
    - The system merges each discovered feature to a feature node in the taxonomy.
    - The similarity metrics are defined based on string similarity, synonyms and other distances measured using WordNet.
    - Experimental results based on digital camera and DVD reviews show promising results.
- Many ideas in information integration are applicable.

Identify opinion orientation on feature

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

Context dependent opinions

- Popescu and Etzioni (2005) used
  - constraints of connectives in (Hazivassiloglou and McKeown, ACL-97), and some additional constraints, e.g., morphological relationships, synonymy and antonymy, and
  - relaxation labeling to propagate opinion orientations to words and features.
- Ding and Liu (2007) used
  - constraints of connectives both at intra-sentence and inter-sentence levels, and
  - additional constraints of, e.g., TOO, BUT, NEGATION, to directly assign opinions to (*f*, *s*) with good results (> 0.85 of F-score).

Roadmap

- Opinion mining – the abstraction
- Document level sentiment classification
- Sentence level sentiment analysis
- Feature-based sentiment analysis and summarization
- Summary

OM: What's next

- Relatively stable workflows have been defined for most of the OM tasks
- However problematic issues still exist:
  - The *inner structure of a subjective statement*
  - *Affective lexicons and domains*
  - Recognizing opinions in *heterogeneous collections*
    - Texts (such as documents or reviews)
    - Sentences or UCGs (such as in blogs/tweets)

OM: Technological directions

- Open Issues:
  - **Adaptivity**: semi-supervised models
    - For the affective lexicon (e.g. Li et al., ACL 2009)
    - For the representation of target texts
    - For generalizing resource across langauges
  - **Fine-grained OM** through Structured learning (e.g. (Johansson & Moschitti, CoNLL 2010))
  - **Compositional subjectivity** models (e.g. (Neviarouskaya et al., COLING 2010))

Directions

- Exploit the geometry of semantic/word spaces (more on this later)
- Compositionality
  - Representation: profile or instance based
  - Operators
  - Source evidence
- Semi-supervised distribuitional models
- Support to interpretation
  - Cognitive aspects (e.g. which features)
  - Semiotics of expectations
- Computational Aspects
  - Scalability
    - Ways of computing SVD or embeddings
  - Distributed Wrokflows

Further References

- Social Media Analytics R. Lawrence, P. Melville, C. Perlich, V.Sindhwani, E.Meliksetian, P.Hsueh, Y. Liu Operations Research/Management Science Today, Feburary 2010

- Bing Liu, Sentiment Analysis and Subjectivity, Handbook of Natural Language Processing, Second Edition, (editors: N. Indurkhya and F. J. Damerau), 2011