A short Introduction to Sentiment Analysis

WM&R a.a. 2021/22

R. BASILI
DIPARTIMENTO DI INGEGNERIA DELL'IMPRESA
UNIVERSITÀ DI ROMA "TOR VERGATA"
EMAIL: BASILI@INFO.UNIROMA2.IT

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

Summary

Introduction to the overall notion of Sentiment Analysis

- The defintion of sentiment and subjectvity
- The model fo the tasks
- Types of OM tasks

Major Approaches to the different tasks

Resources for OM

Architectural and Technological Issues

Evaluation and Benchmarking Champaigns

Neural Approaches to SA

SA in Twitter

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 & Digital culture

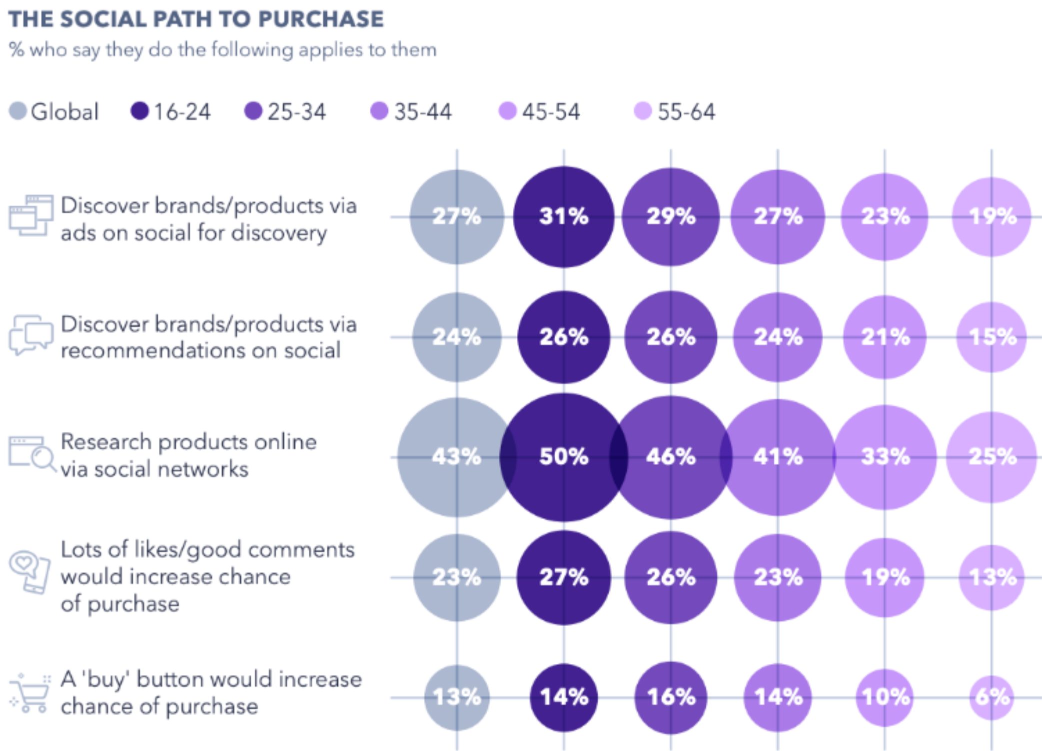

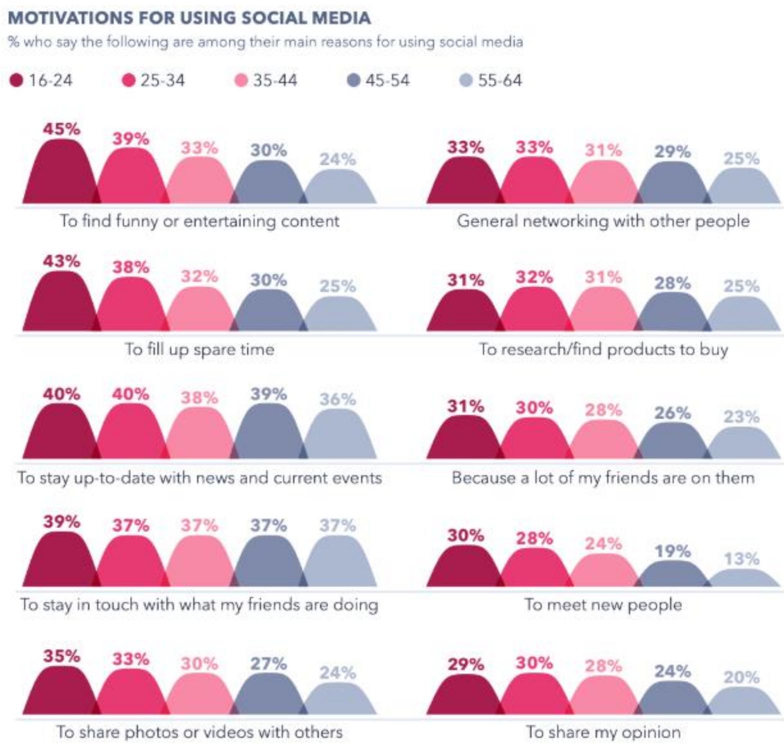

Source: https://blog.hootsuite.com/twitter-demographics/

"First, they do an on-line search."

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.

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

- Summary of reviews of Cell Phone 1

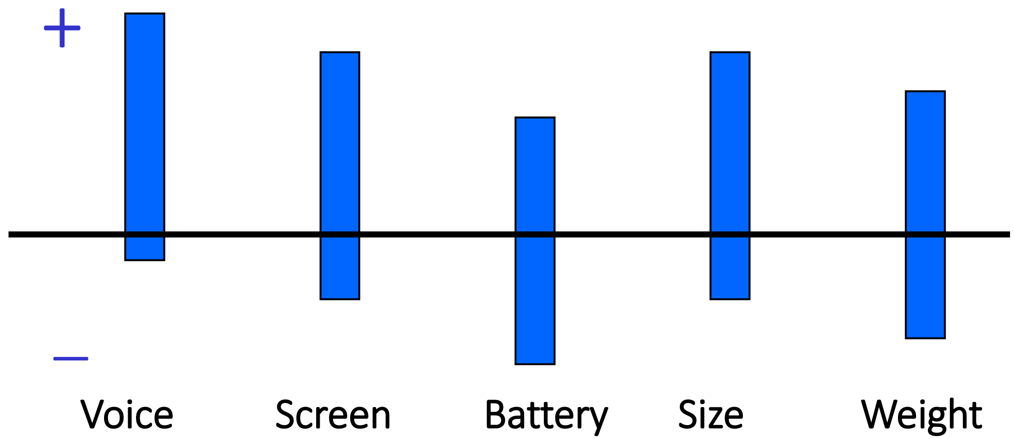

- Comparison of reviews of

Cell Phone 1

Cell Phone 2

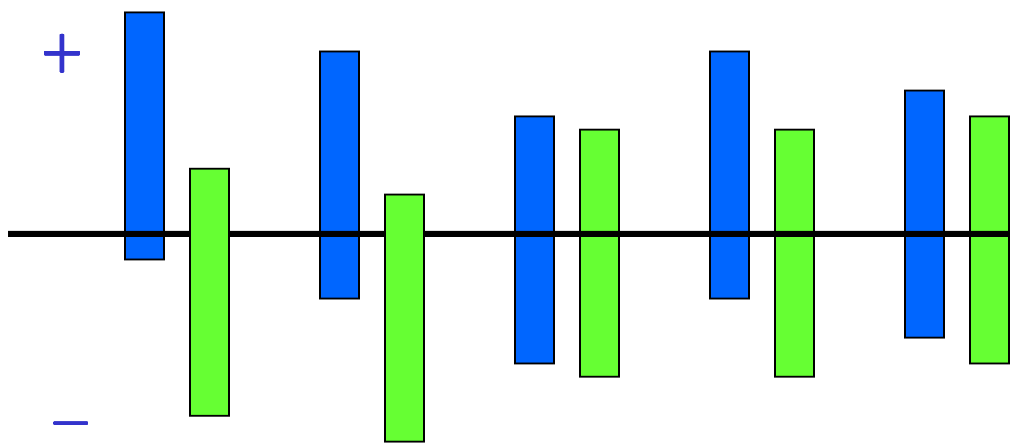

Find the opinion of a person on X

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.

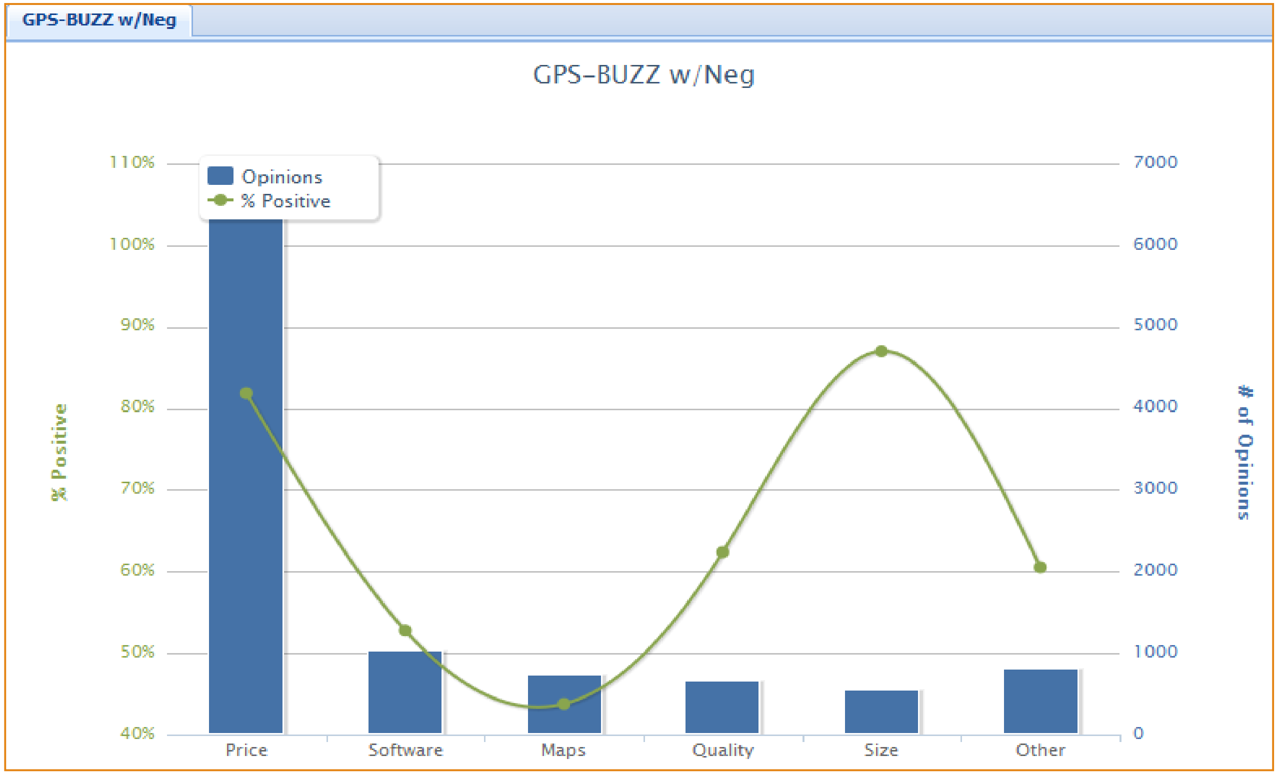

It seems very appealing

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 corresponds to a very complex process!!

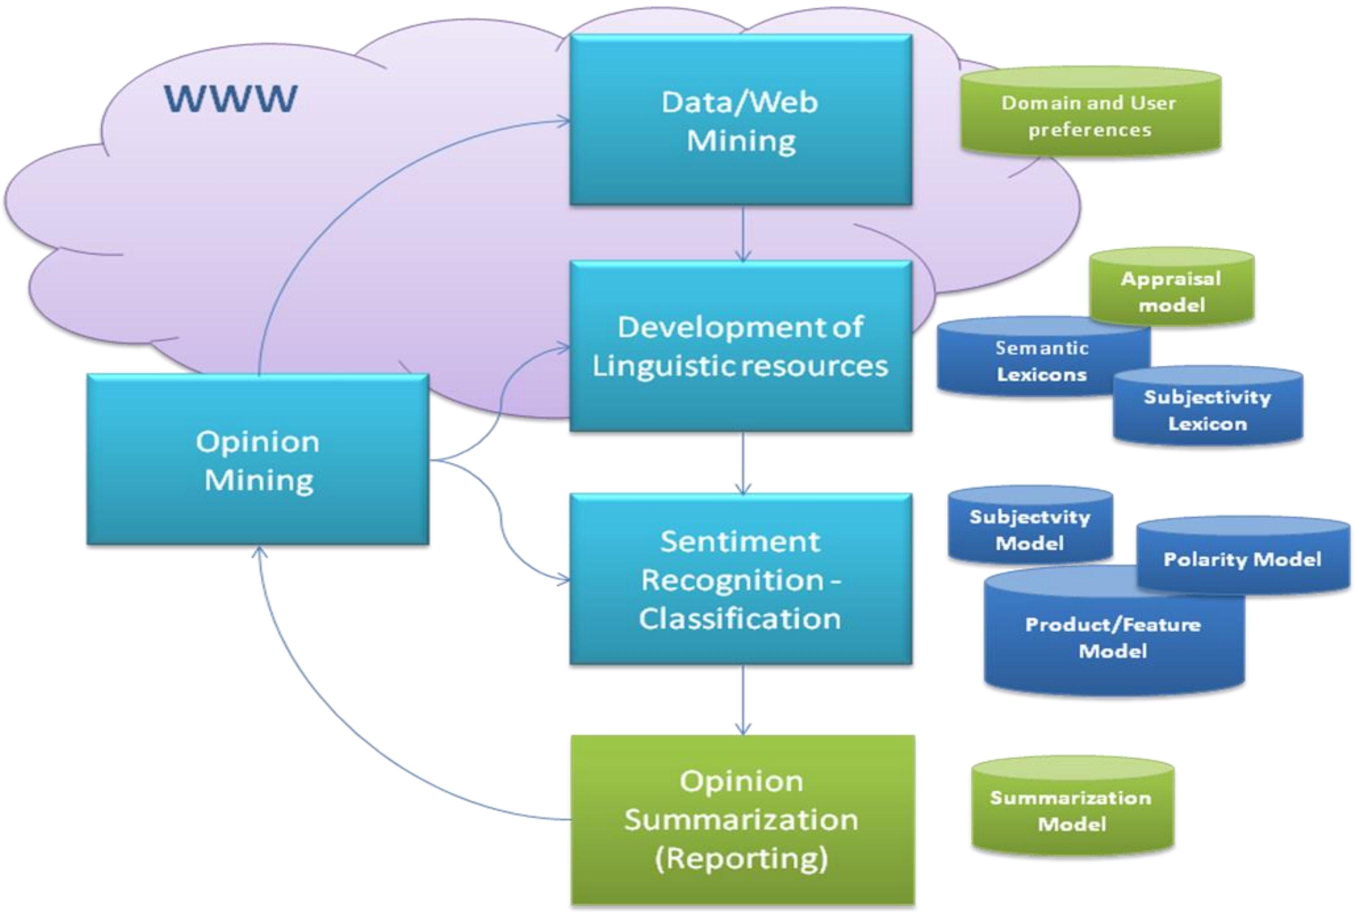

Tasks

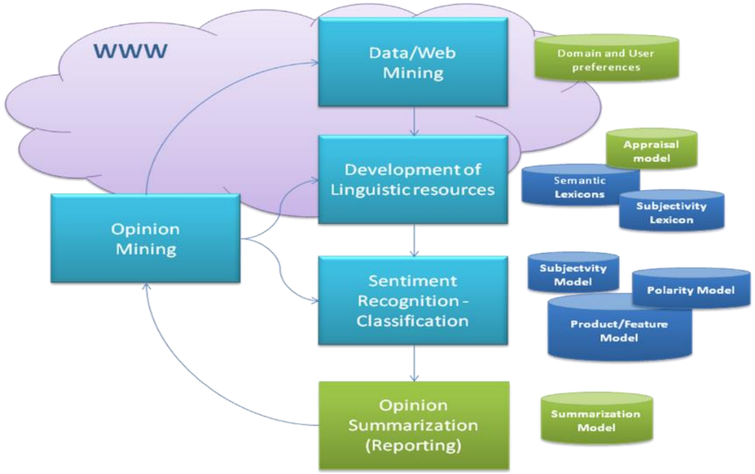

Data Gathering

- Objective: to access information relevant to unders
- Resources: Individual Profiles, Community sites, blogs

Linguistic Resources Development:

- Objective: to develop linguistic models (as ontologies, dictionaries, embeddings, …)
- Resources: general-purpose corpora, domain corpora, opinion datasets
- Outcome: Semantic Lexicons, Subjectivity Lexicons

Sentiment Recognition:

- SubTasks: Subjectivity, Aspect and Polarity Recognition, Opinion Summarization
- Resources: Subjectivity models, Appraisal models, Polarity Models

Opinion Summarization:

- Objective: Summarize opinions across large user communities

"I'd like your honest, unbiased and possibly career-ending opinion on something."

NL vs. Opinions

Although subjectivity seems to preserve across domains and sublanguages, knowledge about *subjectivity (e.g. affective lexicons) is not fully portable*

- For example, the polarity of some terms change across domains (e.g. *small* about laptops vs. TV screen)

These issues trigger 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

Two (closely related) notions

- Subjectivity and emotion.
- **Sentence subjectivity**: An *objective sentence* presents some factual information, while a *subjective sentence* expresses some personal feelings, views, emotions, or beliefs.
- **Emotion**: Emotions are people's subjective

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.

Opinion Ontologies

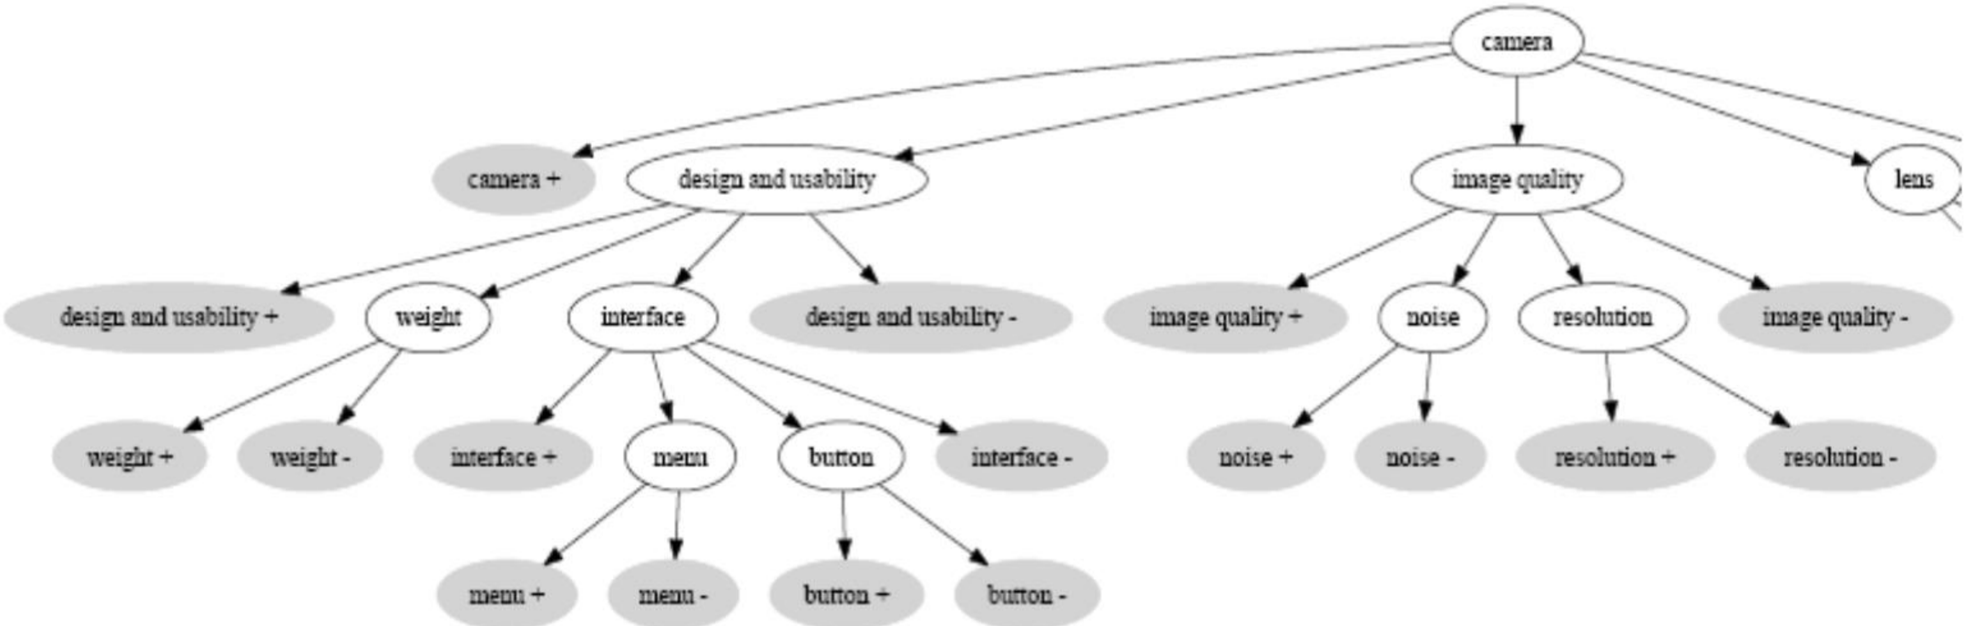

Roadmap

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,
- travel destinations.

The approach: Three steps

Step 1: Feature Extaction

- 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

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 for **estimation**
  - Search to find the **number of hits** in the indexed Web pages to compute PMI and SO
  - The "near" operator is applied to constraint the search

Step 2: Estimate the semantic orientation of the extracted phrases

Estimate the Pointwise Mutual Information
for Semantic orientation (SO):

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

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

Step 3: Estimate the SO of the entire text by averaging

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 supervised 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)

More recent approaches apply Convolutional Neural networks and LSTMs, improvement is significant (+5-10%)

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).

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**.

Sentiment (or opinion) lexicons

- Sentiment lexicon: lists of words and expressions used to express people's subjective feelings and sentiments/opinions.
  - Not just individual words, but also phrases and idioms, e.g., "cost an arm and a leg"
- There seems to be endless variety of sentiment bearing expressions.
  - We have compiled more than 6,700 individual words.
  - There are also a large number of phrases.

Affective Lexicons

They have been extensively used in the field either for lexicon-based approaches or in machine-learning solutions

- Additional features
- Bootstrapping: unsupervised solutions (see previous)

Can be created manually, automatically or semi-automatically

Can be domain-dependent or independent

A lot of them are already available:

- Manual
  - LIWC: Linguistic Inquiry and Word Count [10]
  - ANEW: Affective norms for English words [11]
- Automatic:
  - WordNet-Affect [9]
  - SentiWordNet [31] …

LIWC: Linguistic Inquiry and Word Count (https://liwc.wpengine.com/)

II. PSYCHOLOGICAL PROCESSES

Social Processes	talk, us, friend	455
Friends		
Family		
Humans		
Affective Processes		
Positive Emotions		
Negative Emotions		
Anxiety		
Anger		
Sadness		
Cognitive Processes		
Insight		
Causation		
Discrepancy		
Tentative		
Certainty		
Inhibition		
Inclusive		
Exclusive		
Perceptual Processes		
Seeing		
Hearing		
Feeling	touch, hold, felt	75
Biological Processes	eat, blood, pain	567
Body	ache, heart, cough	180

125 Affect						126 Posemo			127 Negemo			
abandon*	damn*	fume*	kindn*	privileg*	supporting	accept	freed*	partie*	abandon*	enrag*	maddening	snob*
abuse*	danger*	fuming	kiss*	prize*	supportive*	accepta*	freeing	party*	abuse*	envie*	madder	sob
abusi*	daring	fun	laidback	problem*	supports	accepted	freely	passion*	abusi*	envious	maddest	sobbed
accept	darlin*	funn*	lame*	profit*	suprem*	accepting	freeness	peace*	ache*	envy*	maniac*	sobbing
accepta*	daze*	furious*	laugh*	promis*	sure*	accepts	freer	perfect*	aching	evil*	masochis*	sobs
accepted	dear*	fury	lazie*	protest	surpris*	active*	frees*	play	advers*	excruciat*	melanchol*	solemn*
accepting	decay*	geek*	lazy	protested	suspicio*	admir*	friend*	played	afraid	exhaust*	mess	sorrow*
accepts	defeat*	genero*	liabilit*	protesting	sweet	ador*	fun	playful*	aggravat*	fail*	messy	sorry
ache*	defect*	gentle	liar*	proud*	sweetheart*	advantag*	funn*	playing	aggress*	fake	miser*	spite*
aching	defenc*	gentler	libert*	puk*	sweetie*	adventur*	genero*	plays	agitat*	fatal*	miss	stammer*
active*	defens*	gentlest	lied	punish*	sweetly	affection*	gentle	pleasant*	agoniz*	fatigu*	missed	stank
admir*	definite	gently	lies	radian*	sweetness*	agree	gentler	please*	agony	fault*	misses	startl*
ador*	definitely	giggl*	like	rage*	sweets	agreeab*	gentlest	pleasing	alarm*	fear	missing	steal*
advantag*	degrad*	giver*	likeab*	raging	talent*	agreed	gently	pleasur*	alone	feared	mistak*	stench*
adventur*	delectabl*	giving	liked	rancid*	tantrum*	agreeing	giggl*	popular*	anger*	fearful*	mock	stink*
advers*	delicate*	glad	likes	rape*	tears	agreement*	giver*	positiv*	angr*	fearing	mocked	strain*
affection*	delicious*	gladly	liking	raping	teas*	agrees	giving	prais*	anguish*	fears	mocker*	strange
afraid	deligh*	glamor*	livel*	rapist*	tehe	alright*	glad	precious*	annoy*	feroc*	mocking	stress*
aggravat*	depress*	glamour*	LMAO	readiness	temper	amaz*	gladly	prettie*	antagoni*	feud*	mocks	struggl*
aggress*	depriv*	gloom*	LOL	ready	tempers	amor*	glamor*	pretty	anxi*	fiery	molest*	stubborn*
agitat*	despair*	glori*	lone*	reassur*	tender*	amus*	glamour*	pride	apath*	fight*	mooch*	stunk
agoniz*	desperat*	glory	longing*	rebel*	tense*	aok	glori*	privileg*	appall*	fired	moodi*	stunned
agony	despis*	goddam*	lose	reek*	tensing	appreciat*	glory	prize*	apprehens*	flunk*	moody	stuns
agree	destroy*	good	loser*	regret*	tension*	assur*	good	profit*	argh*	foe*	moron*	stupid*
agreeab*	destruct*	goodness	loses	reject*	terribl*	attachment*	goodness	promis*	argu*	fool*	mourn*	stutter*
agreed	determina*	gorgeous*	losing	relax*	terrific*	attract*	gorgeous*	proud*	arrogan*	forbid*	murder*	submissive*
agreeing	determined	gossip*	loss*	relief	terrified	award*	grace	radian*	asham*	fought	nag*	suck
agreement*	devastat*	grace	lost	reliev*	terrifies	awesome	graced	readiness	assault*	frantic*	nast*	sucked
agrees	devil*	graced	lous*	reluctan*	terrify	beaut*	graceful*	ready	asshole*	freak*	needy	sucker*
alarm*	devot*	graceful*	love	remorse*	terrifying	beloved	graces	reassur*	attack*	fright*	neglect*	sucks
alone	difficult*	graces	loved	repress*	terror*	benefic*	graci*	relax*	aversi*	frustrat*	nerd*	sucky
alright*	digni*	graci*	lovely	resent*	thank	benefit	grand	relief	avoid*	fuck	nervous*	suffer

The VAD model

V: Pleasantry

A: Intensity

D: Control

Core Dimensions of Connotative Meaning

Influential factor analysis studies (Osgood et al., 1957; Russell, 1980, 2003) have shown that the three most important, largely independent, dimensions of word meaning:

- valence (V): positive/pleasure – negative/displeasure
- arousal (A): active/stimulated – sluggish/bored
- dominance (D): powerful/strong – powerless/weak

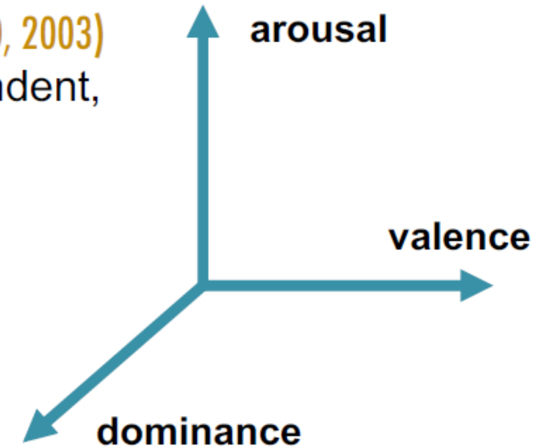

Thus, when comparing the meanings of two words, we can compare their V, A, D scores. For example:

- *banquet* indicates more positiveness than *funeral*
- *nervous* indicates more arousal than *lazy*
- *queen* indicates more dominance than *delicate*

VAD lexicons: examples of entries

Dimension	Word	Score↑	Word	Score↓
valence	*love*	1.000	*toxic*	0.008
	happy	1.000	*nightmare*	0.005
	happily	1.000	*shit*	0.000
arousal	*abduction*	0.990	*mellow*	0.069
	exorcism	0.980	*siesta*	0.046
	homicide	0.973	*napping*	0.046
dominance	*powerful*	0.991	*empty*	0.081
	leadership	0.983	*frail*	0.069
	success	0.981	*weak*	0.045

ANEW: Affective norms for English words

Description	Word No.	Valence Mean(SD)	Arousal Mean(SD)	Dominance Mean (SD)	Word Frequency
abduction	621	2.76 (2.06)	5.53 (2.43)	3.49 (2.38)	1
abortion	622	3.50 (2.30)	5.39 (2.80)	4.59 (2.54)	6
absurd	623	4.26 (1.82)	4.36 (2.20)	4.73 (1.72)	17
abundance	624	6.59 (2.01)	5.51 (2.63)	5.80 (2.16)	13
abuse	1	1.80 (1.23)	6.83 (2.70)	3.69 (2.94)	18
acceptance	625	7.98 (1.42)	5.40 (2.70)	6.64 (1.91)	49
accident	2	2.05 (1.19)	6.26 (2.87)	3.76 (2.22)	33
ace	626	6.88 (1.93)	5.50 (2.66)	6.39 (2.31)	15
ache	627	2.46 (1.52)	5.00 (2.45)	3.54 (1.73)	4
achievement	3	7.89 (1.38)	5.53 (2.81)	6.56 (2.35)	65
activate	4	5.46 (0.98)	4.86 (2.56)	5.43 (1.84)	2
addict	581	2.48 (2.08)	5.66 (2.26)	3.72 (2.54)	1
addicted	628	2.51 (1.42)	4.81 (2.46)	3.46 (2.23)	3
admired	5	7.74 (1.84)	6.11 (2.36)	7.53 (1.94)	17
adorable	6	7.81 (1.24)	5.12 (2.71)	5.74 (2.48)	3
adult	546	6.49 (1.50)	4.76 (1.95)	5.75 (2.21)	25
advantage	629	6.95 (1.85)	4.76 (2.18)	6.36 (2.23)	73
adventure	630	7.60 (1.50)	6.98 (2.15)	6.46 (1.67)	14
affection	7	8.39 (0.86)	6.21 (2.75)	6.08 (2.22)	18
afraid	8	2.00 (1.28)	6.67 (2.54)	3.98 (2.63)	57

The multidimensional view on emotions

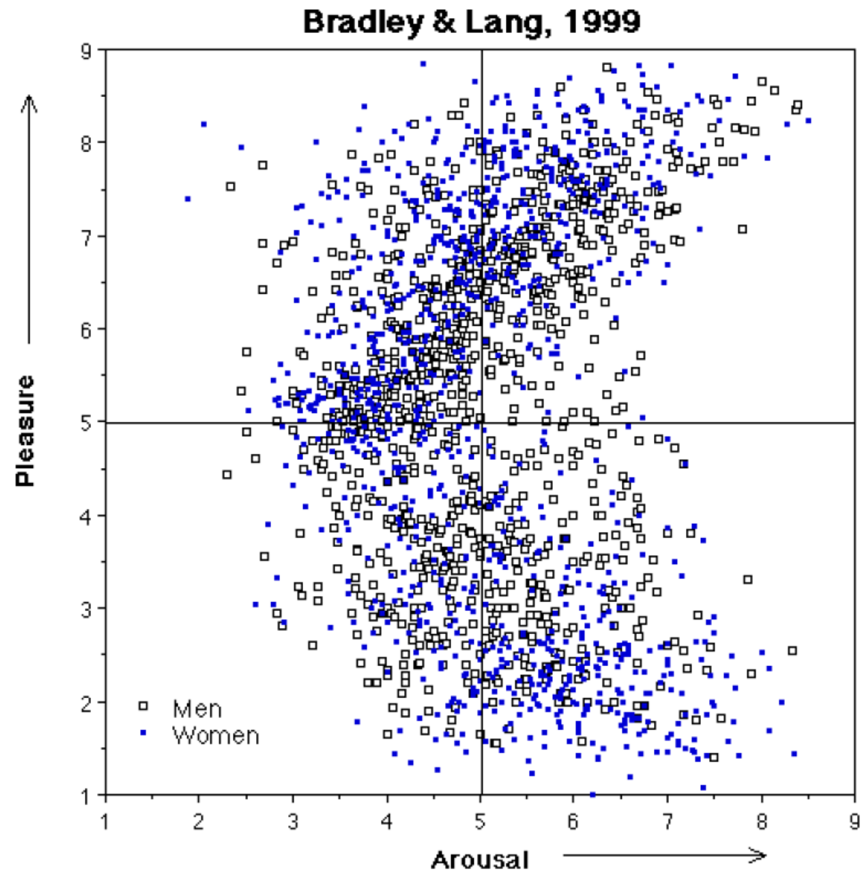

Corpus-based approaches

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.

Corpus-based approaches – A LSA Example

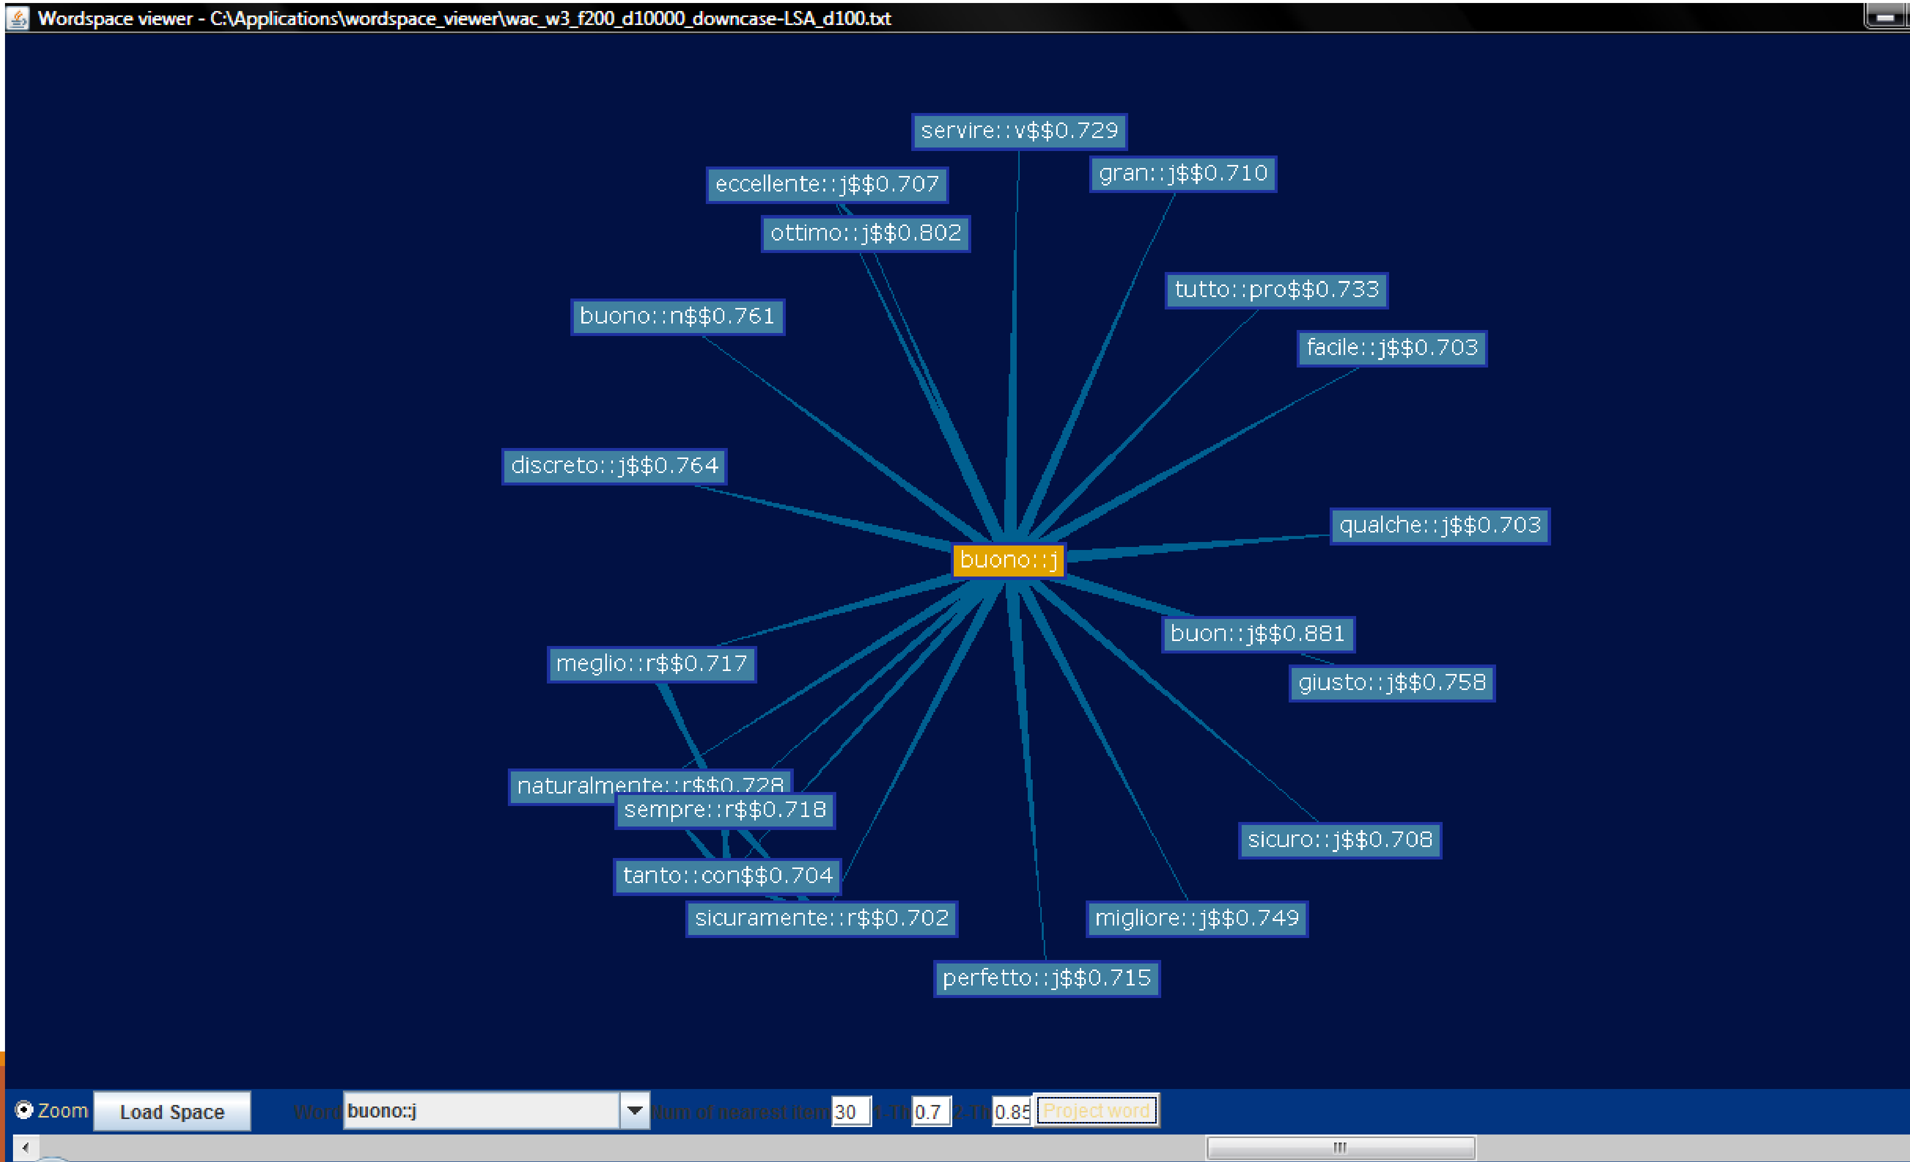

Corpus-based approaches – A LSA Example

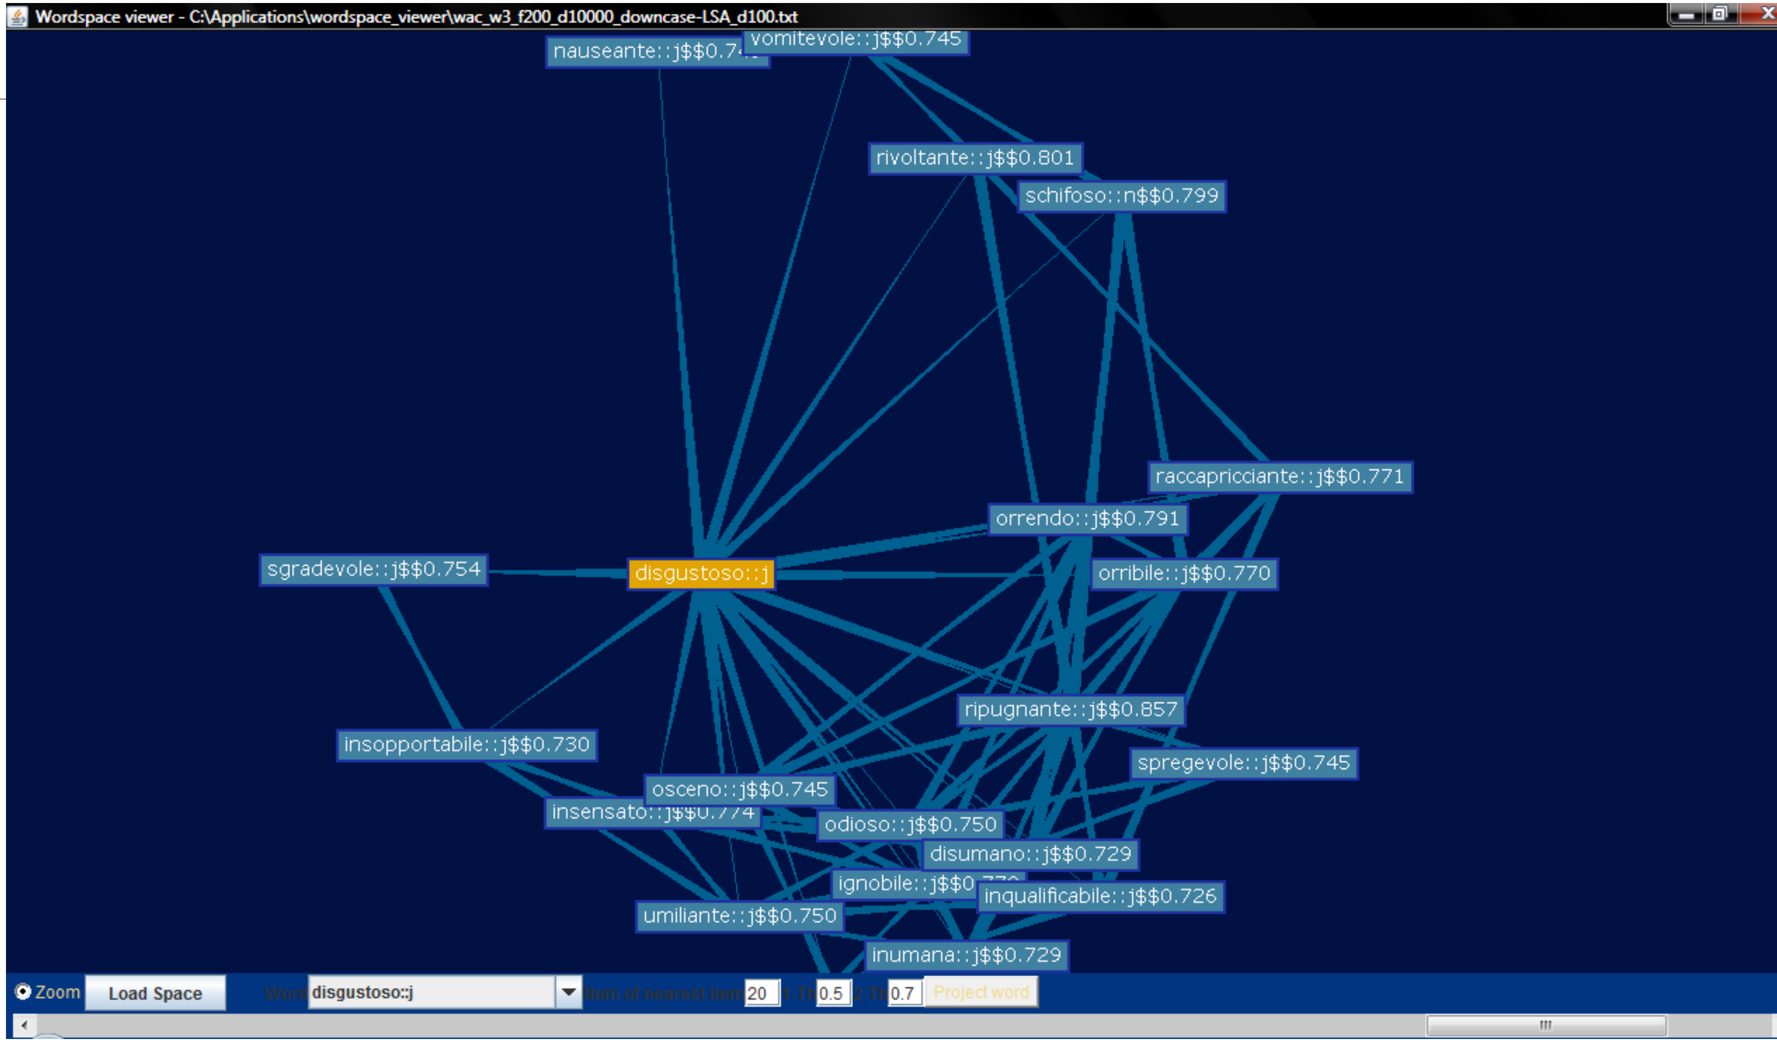

Dictionary-based approaches

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 (Esuli and Sebastiani, CIKM-05).

Weakness of the approach: Do not find domain and/or context dependent opinion words, e.g., small, long, fast.

MPQA Lexicon (Wilson et al., HLT-EMNLP 2005)

- Annotated corpus
- Annotated Lexicon
- Relatively rich theory of appraisal behind sentiment annotations

MPQA Lexicon (Wilson et al., HLT-EMNLP 2005)

MPQA Subjectivity Cues Lexicon

Theresa Wilson, Janyce Wiebe, and Paul Hoffmann (2005). Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. Proc. of HLT-EMNLP-2005.

Riloff and Wiebe (2003). Learning extraction patterns for subjective expressions. EMNLP-2003.

- Home page: http://www.cs.pitt.edu/mpqa/subj_lexicon.html
- 6885 words from 8221 lemmas
  - 2718 positive
  - 4912 negative
- Each word annotated for intensity (strong, weak)

MPQA: Overview

Opinions, evaluations, emotions, speculations are private states.

They are expressed in language by subjective expressions.

Private state: state that is not open to objective observation or verification.

Quirk, Greenbaum, Leech, Svartvik (1985). *A Comprehensive Grammar of the English Language*.

MPQ: Overview

Focus on three ways private states are expressed in language

- **Direct subjective expressions**
- **Expressive subjective elements**
- **Objective speech events**

Direct Subjective Expressions

Direct mentions of private states

The United States ***fears*** *a spill-over from the anti-terrorist campaign.*

Private states expressed in speech events

"We foresaw electoral fraud but not daylight robbery," Tsvangirai ***said****.*

Expressive Subjective Elements
[Banfield 1982]

"We foresaw electoral fraud but not ***daylight robbery****,"*
Tsvangirai said

The part of the US human rights report about China is ***full of absurdities and fabrications***

Objective Speech Events

Material attributed to a source, but presented as objective fact

*"The government, it **added**, has amended the Pakistan Citizenship Act 10 of 1951 to enable women of Pakistani descent to claim Pakistani nationality for their children born to foreign husbands."*

MPQA: Attitude Types

Table 7.1: Set of attitude types

Sentiment	**Agreement**
Positive Sentiment	Positive Agreement
Negative Sentiment	Negative Agreement
Arguing	**Intention**
Positive Arguing	Positive Intention
Negative Arguing	Negative Intention
Speculation	**Other Attitude**

MPQA: *Arguing*

Positive Arguing:
(7.8) Iran **insists** ⟨**its nuclear program**⟩ **is purely for peaceful purposes.**

(7.9) Putin remarked that ⟨the events in Chechnia⟩ **“could be interpreted only in the context of the struggle against international terrorism.”**

Negative Arguing:
(7.10) Officials in Panama **denied that** ⟨**Mr. Chavez or any of his family members**⟩ **had asked for asylum.**

(7.11) “⟨It⟩ **is analogous to the US crackdown on terrorists in Afghanistan,”** Ma said.

MPQA: attitude frames

(7.18) Its <u>aim</u> of the 2001 report is to tarnish China's image and exert political pressure on the Chinese Government, human rights experts <u>said</u> at the seminar held by the China Society for Study of Human Rights (CSSHR) on Friday.

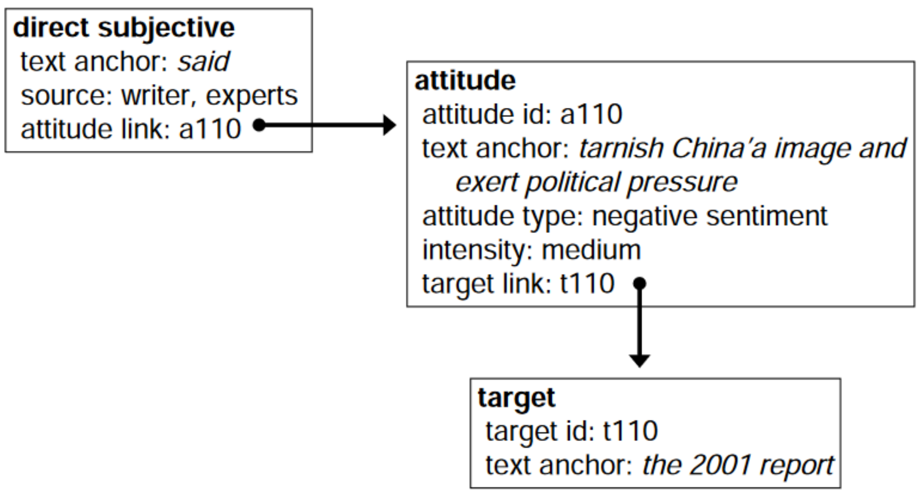

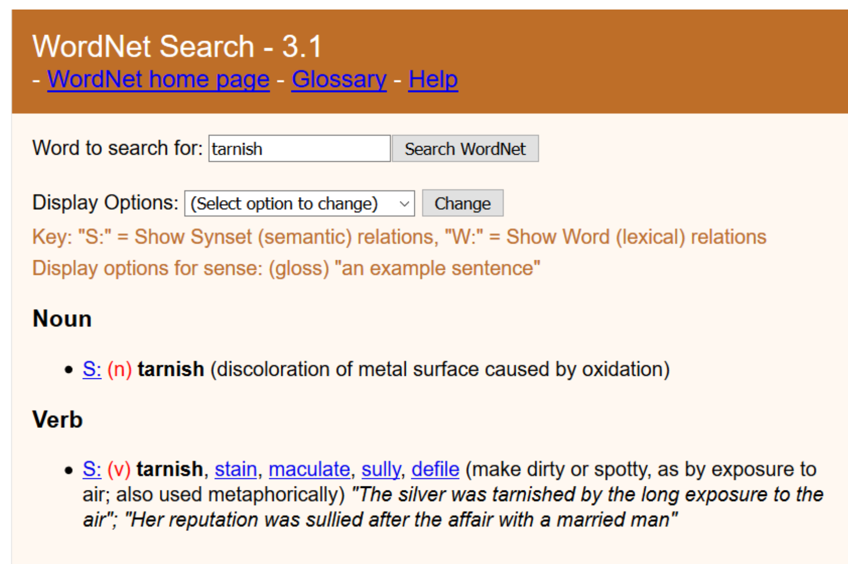

Who does lexicon development ?

Humans

Semi-automatic

Fully automatic

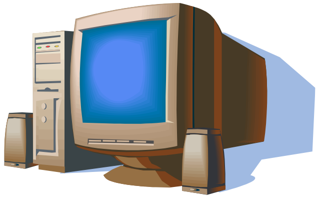

What?

Find relevant words, phrases, patterns that can be used to express subjectivity

Determine the polarity of subjective expressions

Words

Adjectives (e.g. Hatzivassiloglou & McKeown 1997, Wiebe 2000, Kamps & Marx 2002, Andreevskaia & Bergler 2006)

- positive
- negative: harmful hypocritical inefficient insecure
  - It was a macabre and hypocritical circus.
  - Why are they being so inefficient ?
-

Words

Adjectives (e.g. Hatzivassiloglou & McKeown 1997, Wiebe 2000, Kamps & Marx 2002, Andreevskaia & Bergler 2006)

- positive
- negative
- Subjective (but not positive or negative sentiment): curious, peculiar, odd, likely, probable
  - He spoke of Sue as his probable successor.
  - The two species are likely to flower at different times.

Words

Other parts of speech (e.g. Turney & Littman 2003, Riloff, Wiebe & Wilson 2003, Esuli & Sebastiani 2006)

- Verbs
  - positive: **praise, love**
  - negative: **blame, criticize**
  - subjective**: predict**
- Nouns
  - positive: **pleasure, enjoyment**
  - negative: **pain, criticism**
  - subjective: **prediction, feeling**

Attitude Intensity

Table 6.2: Measures of intensity for different attitude types.

Attitude Type	Measure of Intensity	Example
Positive Sentiment	degree of positiveness	*like* $<$ *love*
Negative Sentiment	degree of negativeness	*criticize* $<$ *excoriate*
Positive Agreement	degree of agreement	*mostly agree* $<$ *agree*
Negative Agreement	degree of disagreement	*mostly disagree* $<$ *completely disagree*
Positive Arguing	degree of certainty/strength of belief	*critical* $<$ *absolutely critical*
Negative Arguing	degree of certainty/strength of belief	*should not* $<$ *really should not*
Positive Intention	degree of determination	*promise* $<$ *promise with all my heart*
Negative intention	degree of determination	*no intention* $<$ *absolutely no intention*
Speculation	degree of likelihood	*might win* $<$ *really might win*

Bootstrapping by pattern acquisition
[Riloff & Wiebe 2003]

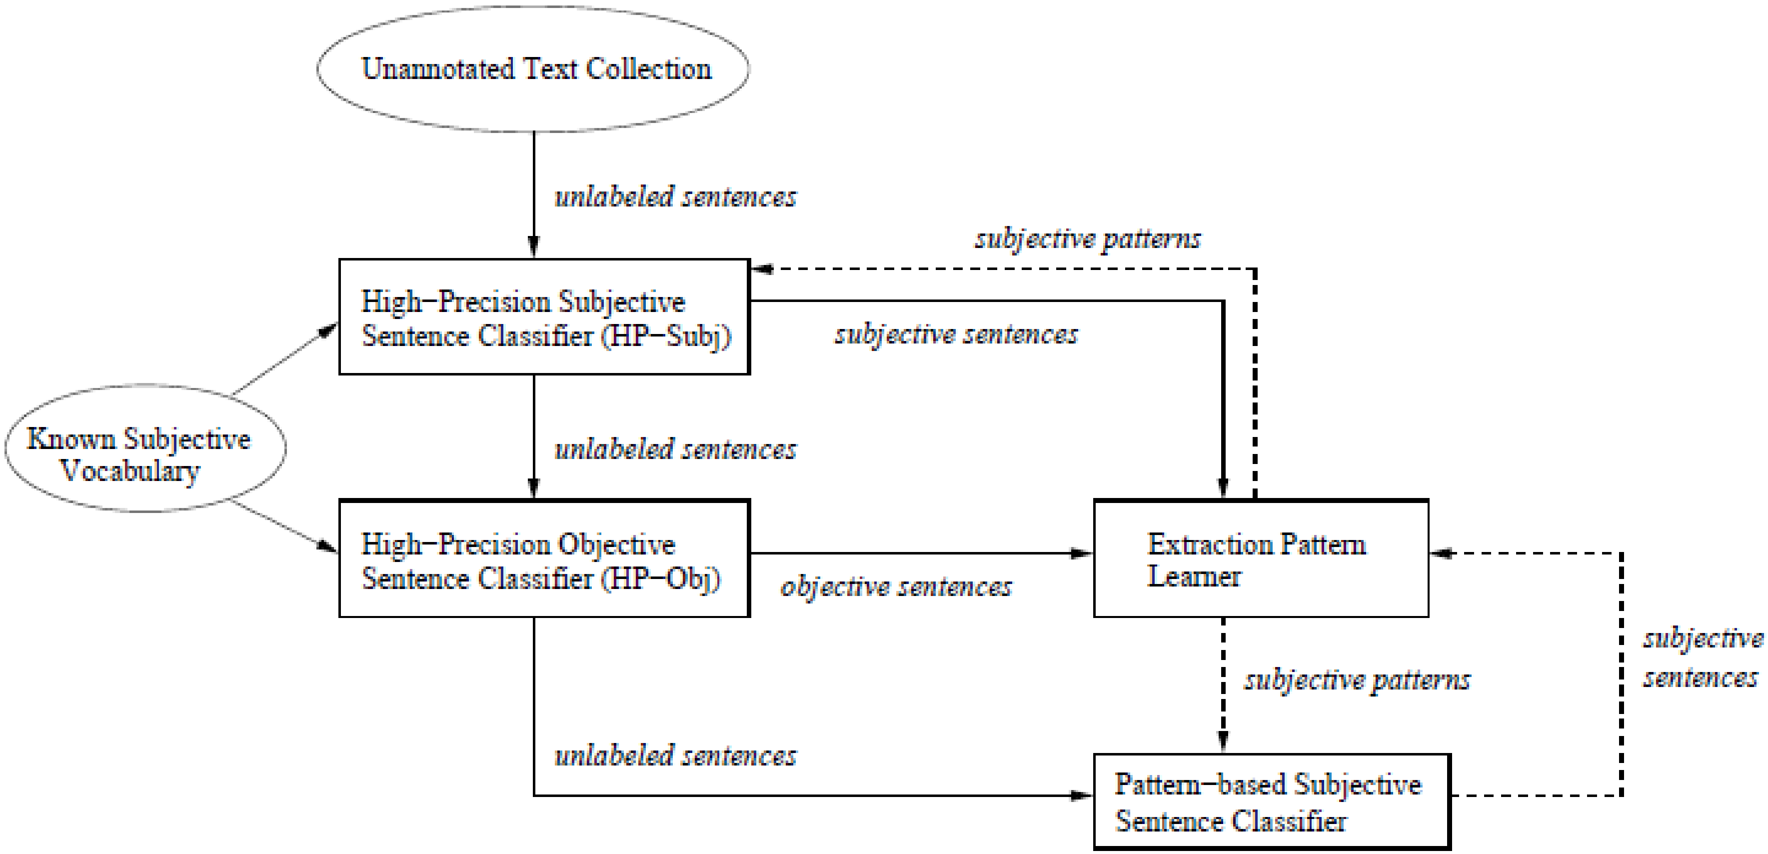

Figure 1: Bootstrapping Process

Bing Liu’s Opinion Lexicon

Minqing Hu and Bing Liu. Mining and Summarizing Customer Reviews. ACM SIGKDD-2004.

- http://www.cs.uic.edu/~liub/FBS/opinion-lexicon-English.rar
- 6786 words
  - 2006 positive
    - … abound, abounds, abundance, abundant, accessable, accessible, acclaim, acclaimed, acclamation, accolade, accolades, accommodative, accomodative, accomplish, accomplished, accomplishment, accomplishments, accurate, ...
  - 4783 negative
    - …., abnormal, abolish, abominable, abominably, abominate, abomination, abort, aborted, aborts, abrade, abrasive, …

OM resources: SentiWordnet

SentiWN (Sebastiani & Esuli, 2008)

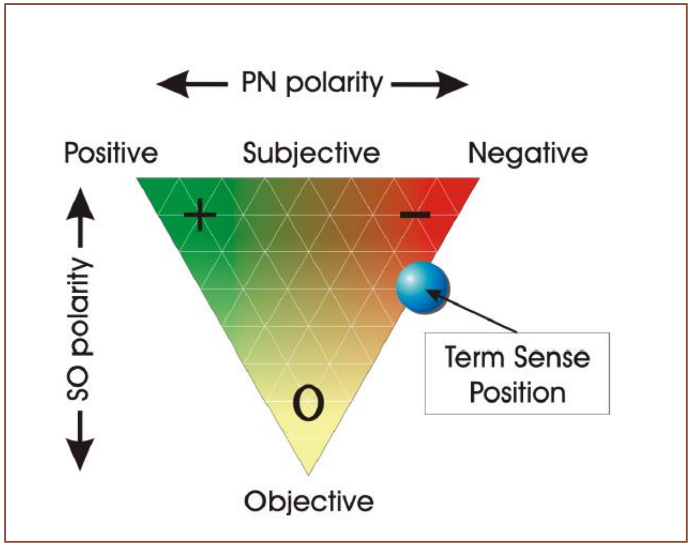

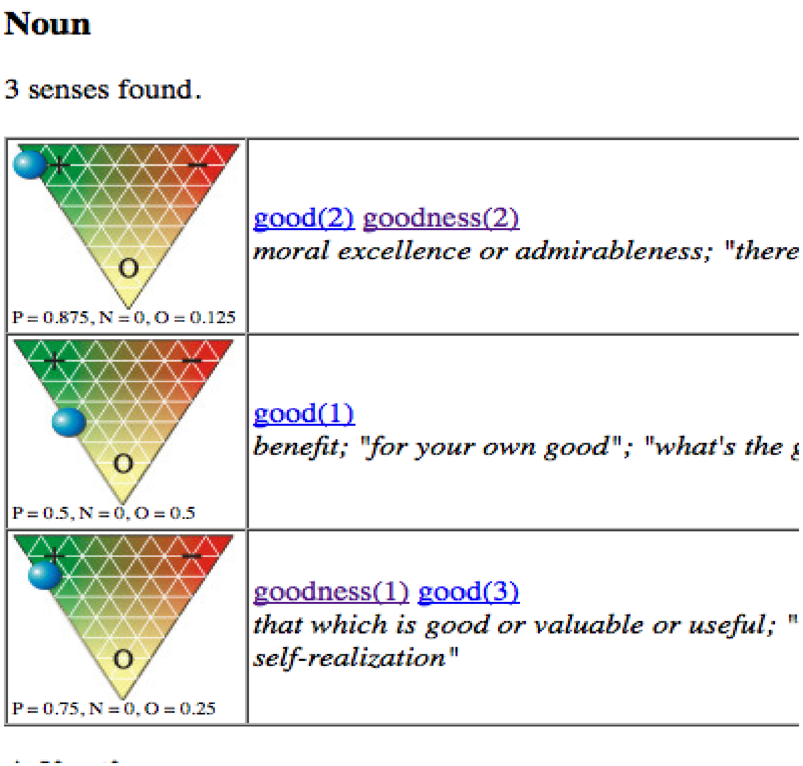

Sentiwordnet

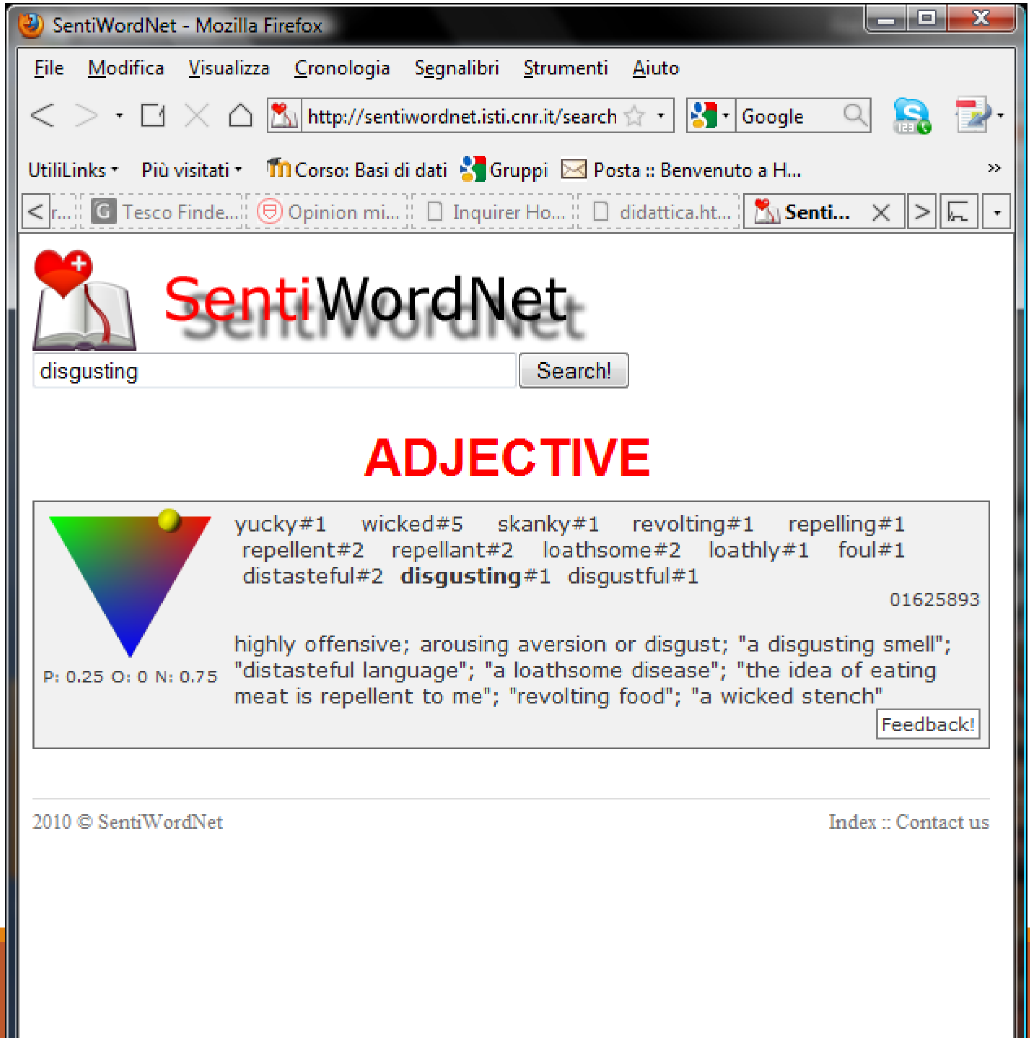

Semi-automatic approach to the design

A SNA methods applied to lexical semantics (Sebastani & Esuli, 2008)

- PageRank over word senses

Creating affective lexicons: using WordNet

WordNet: A lexical database for the English language, that provides various semantic relations between tokens (e.g., synonyms, antonyms)

Can be used to classify positive/negative tokens, based on distance from seed words

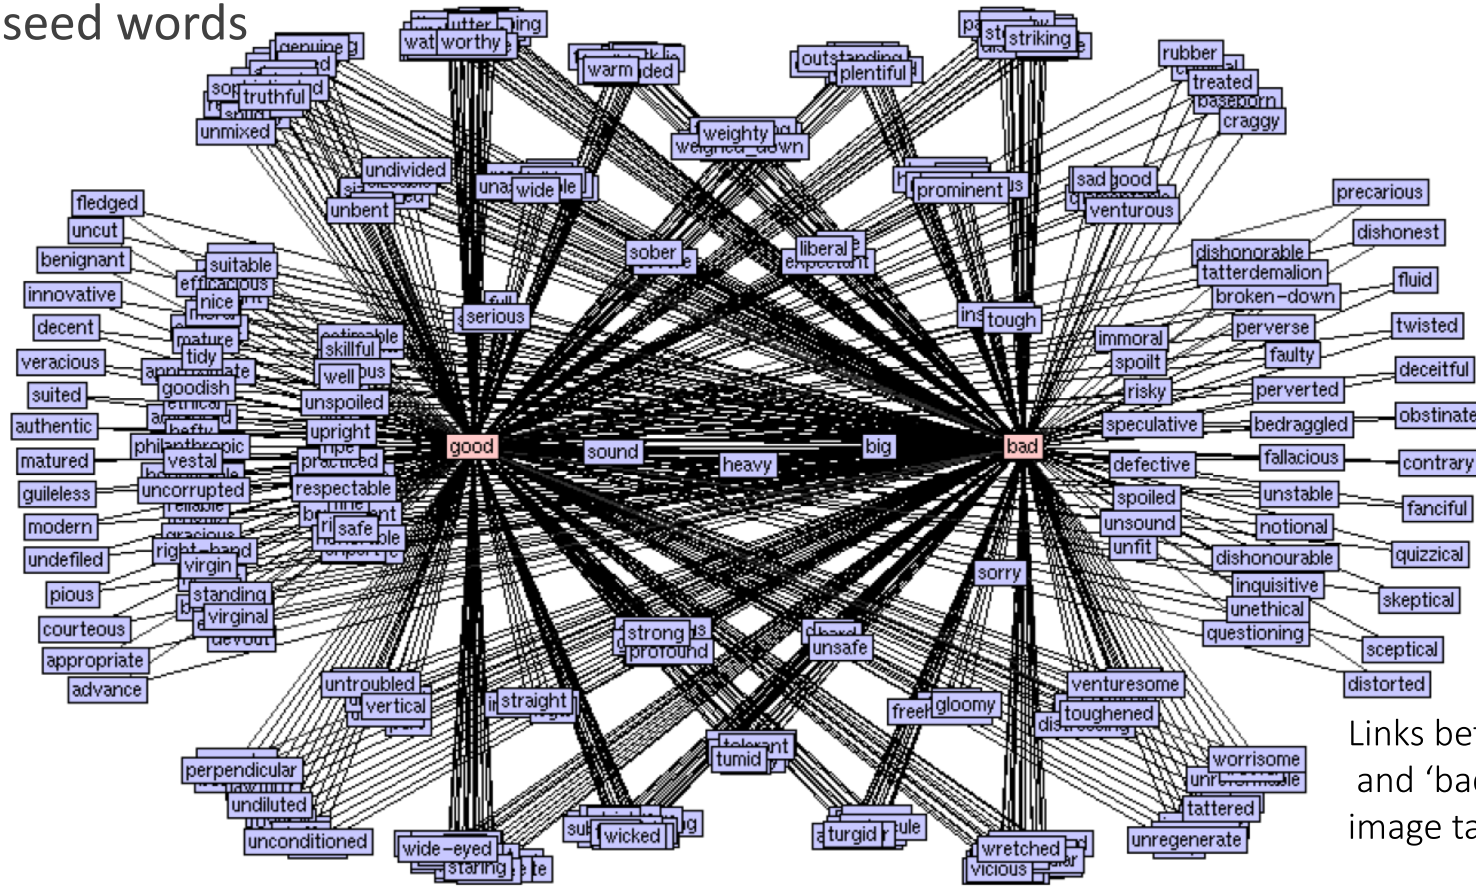

Links between ‘good’ and ‘bad’ in WordNet image taken from [5]

OM Resources: Sentiful DB

Presented by (Neviarouskaya et al., 2010)

Positive seeds:

- ‘interest’, ‘joy’ and ‘surprise’

Negative seeds

- ‘anger’, ‘disgust’, ‘fear’, ‘guilt’, ‘sadness’ ‘shame’

Table 1: Examples of words with sentiment annotations from SentiFul.

Affective word	**POS**	**Non-zero-intensity emotions from Affect database emotional vector**	**Polarity scores**		**Polarity weights**	
			Pos_score	***Neg_score***	***Pos_weight***	***Neg_weight***
tremendous	adjective	‘surprise:1.0’, ‘joy:0.5’, ‘fear:0.1’	0.75	0.1	0.67	0.33
pensively	adverb	‘sadness:0.2’, ‘interest:0.1’	0.1	0.2	0.5	0.5
success	noun	‘joy:0.9’, ‘interest:0.6’, ‘surprise:0.5’	0.67	0.0	1.0	0.0
regret	verb	‘guilt:0.2’, ‘sadness:0.1’	0.0	0.15	0.0	1.0

Sentiful DB

Table 8: Emotional states and relevant expressive means (data partially taken from [53]).

Emotion	Expressive means
Anger	widely open eyes, fixated; pupils contracted; stare gaze; ajar mouth; teeth usually clenched tightly; rigidity of lips and jaw; lips may be tightly compressed, or may be drawn back to expose teeth
Disgust	narrowed eyes, may be partially closed as result of nose being drawn upward; upper lip drawn up; pressed lips; wrinkled nose; turn of the head to the side quasi avoiding something
Fear	widely open eyes; pupils dilated; raised eyebrows; open mouth with crooked lips; trembling chin
Guilt	downcast or glancing gaze; inner corners of eyebrows may be drawn down; lips drawn in, corners depressed; head lowered
Interest	eyes may be exaggeratedly opened and fixed; lower eyelids may be raised as though to sharpen visual focus; increased pupil size; sparkling gaze; mouth slightly smiling; head is slightly inclined to the side
Joy	"smiling" and bright eyes; genuinely smiling mouth
Sadness	eyelids contracted; partially closed eyes; downturning mouth
Shame	downcast gaze; blushing cheeks; head is lowered
Surprise	widely open eyes; slightly raised upper eyelids and eyebrows; the mouth is opened by the jaw drop; the lips are relaxed

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

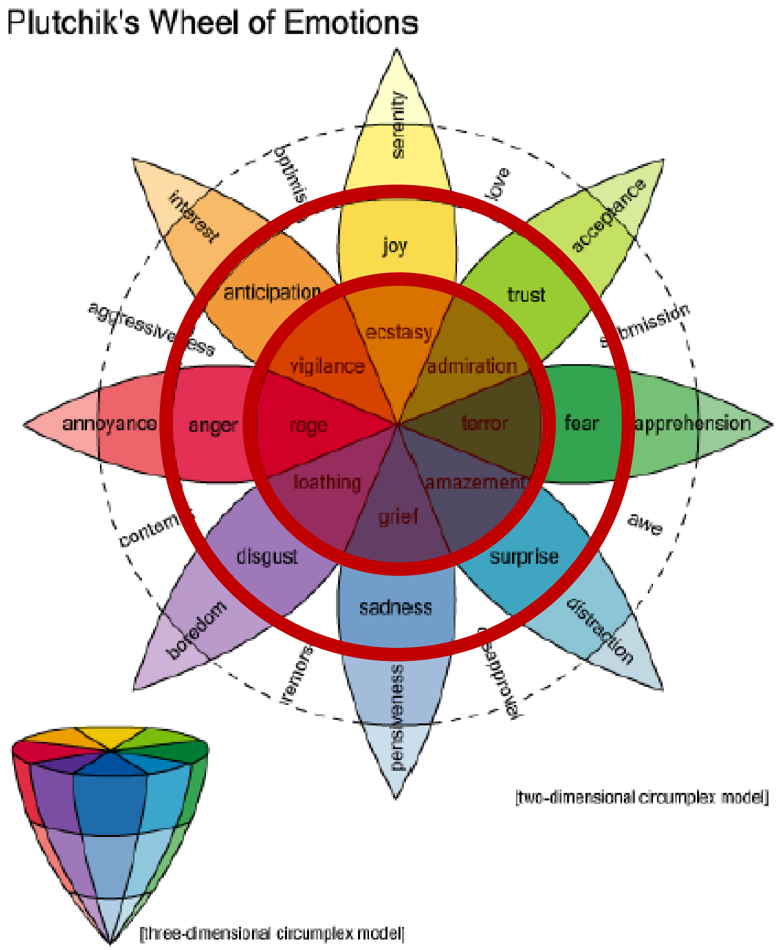

FIGURE 1. Plutchik's wheel of emotions. Similar emotions are placed next to each other. Contrasting emotions are placed diametrically opposite to each other. Radius indicates intensity. White spaces in between the basic emotions represent primary dyads—complex emotions that are combinations of adjacent basic emotions. (The image file is taken from Wikimedia Commons.)

SenticNet (3)

Eric Cambria, 2010

URL: http://sentic.net/

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**

…

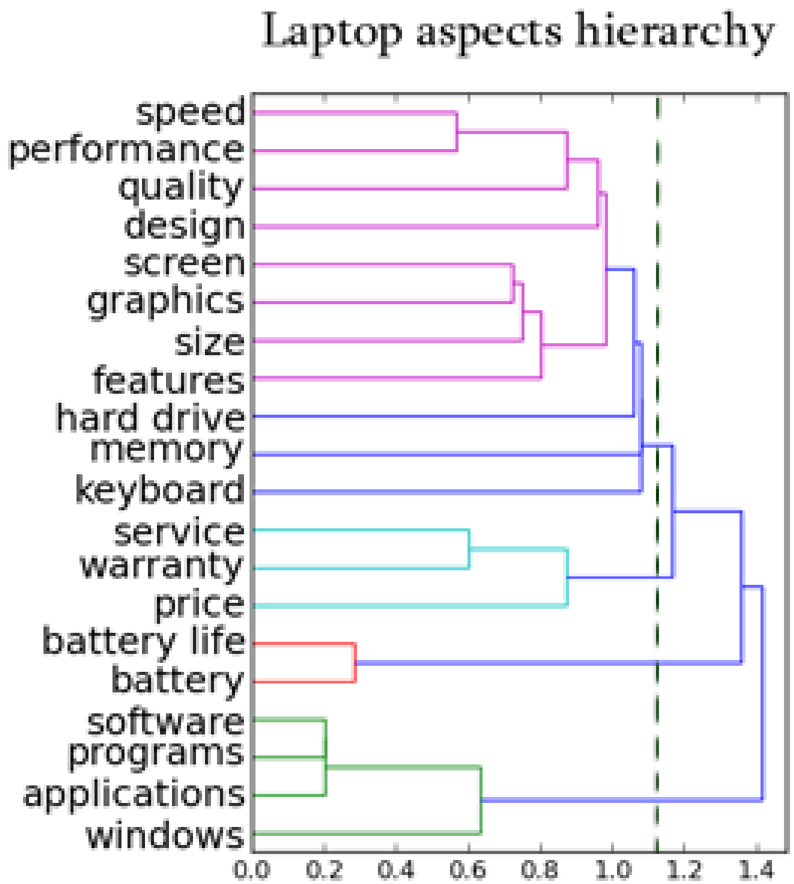
Laptop aspects hierarchy
speed
performance
quality
design
screen
graphics
size
features
hard drive
memory
keyboard
service
warranty
price
battery life
battery
software
programs
applications
windows
0.0 0.2 0.4 0.6 0.8 1.0 1.2 1.4

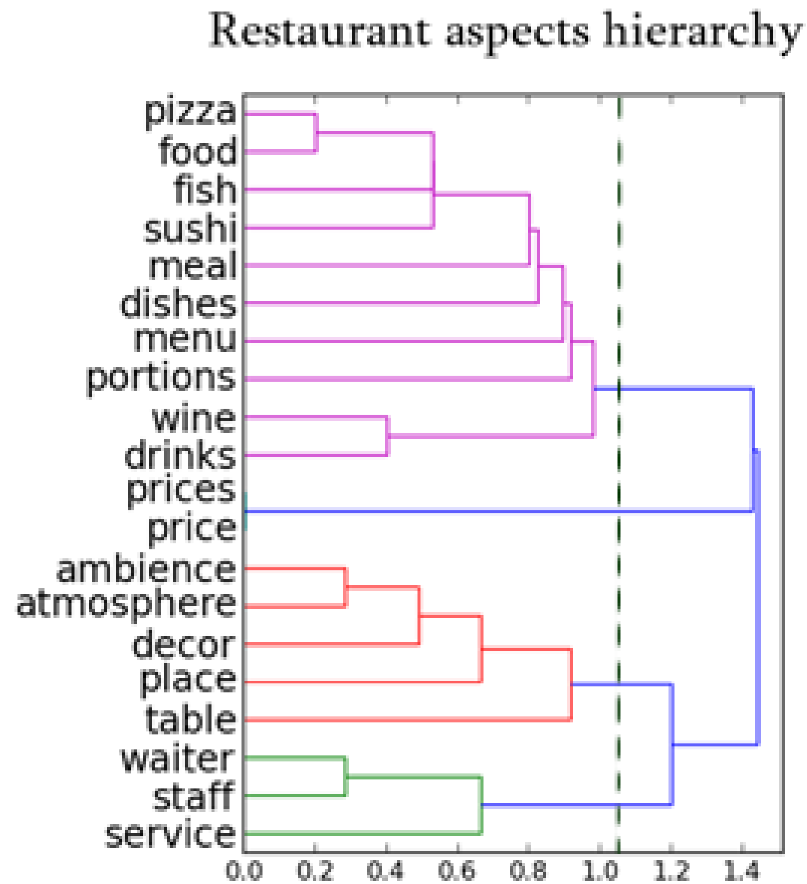
Restaurant aspects hierarchy
pizza
food
fish
sushi
meal
dishes
menu
portions
wine
drinks
prices
price
ambience
atmosphere
decor
place
table
waiter
staff
service
0.0 0.2 0.4 0.6 0.8 1.0 1.2 1.4

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> ⇒ <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.

Creating affective lexicons: using conjunction

Web Results **1** - **10** of about **762,000** for **"was very nice and"**.

The Homestay Experience - Cultural Kaleidoscope 2006
My host's home **was very nice and** comfortable. I got to try all types of food; Malaysian, Chinese, Indonesian and I loved it all. My host's parents were very **...**
www.gardenschool.edu.my/studentportal/aec/Kaleidoscope06/experience.asp - 10k -
Cached - Similar pages - Note this

PriceGrabber User Rating for Watch Your Budget - PriceGrabber.com
Reviews, Camera I purchased **was very nice and** a bargain. There was a problem with shipping, but was resolved quickly. Buy with confidence from this vendor. **...**
www.pricegrabber.com/rating_getreview.php/retid=5821 - Similar pages - Note this

Testimonials
"Everybody **was very nice and** service was as fast as they possibly could. **...** "Staff member who helped me **was very nice and** easy to talk to." **...**
www.sa.psu.edu/uhs/news/testimonials.cfm - 22k - Cached - Similar pages - Note this

Naxos Villages - Naxos Town or Chora Reviews: Very nice and very **...**
-Did you enjoy the trip to Naxos Town: Yes it **was very nice and** very scenic. -In order to get to the village were there enough signs in order to find it: It **...**

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).

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 features: 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).

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."

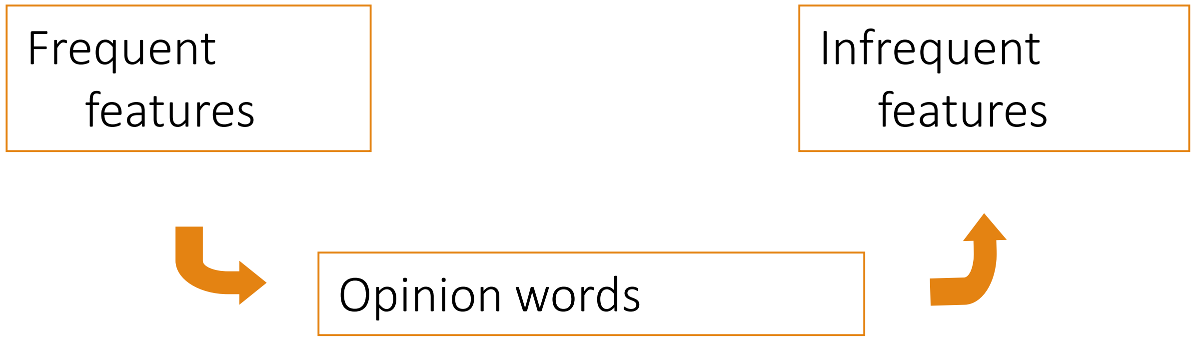

Identify feature synonyms

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.

USE CASES

COVID study (2020): https://mdpi-res.com/d_attachment/applsci/applsci-12-03709/article_deploy/applsci-12-03709.pdf?version=1649318517

SURVEY on DNNs for SA (2020): https://arxiv.org/ftp/arxiv/papers/2006/2006.03541.pdf

The ENEL case: Opinion_Mining_Rbas_the_ENEL_case_v1.0.pptx

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))
  - Neural nets (e.g. (Kim, 2014)
- Social Dynamics through
  - Complex architectures
  - Models of Social profiles and comunications

Twitter Sentiment Analysis@RTV

ACL SemEval champaigns:

- Example 2016, Task 5: http://alt.qcri.org/semeval2016/task5/

EvalIta champaigns:

- Example, 2016, ABSITA: http://sag.art.uniroma2.it/absita/

Further References

Bo Pang and Lillian Lee. 2008. Opinion Mining and Sentiment Analysis. *Found. Trends Inf. Retr.* 2, 1-2 (January 2008), 1-135. DOI=http://dx.doi.org/10.1561/1500000011

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

An Example Use case

See slides on «SA on Twitter at Semeval 2013»

More information in:

“***Injecting sentiment information in context-aware convolutional neural networks***” (Croce et Al, 2016), SocialNLP 2016 Proceedings, IJCAO 2016, New York. URL: https://sites.google.com/site/socialnlp2016/ .

References

NLP, IR & ML:

- «*Speech and Language Processing*", D. Jurafsky and J. H .Martin, Prentice-Hall, 2009.
- "*Introduction to Information Retrieval*", Manning, Raghavan & Schutze, Cambridge University Press 2008.

Opinion Mining

- *Opinion Mining and Sentiment Analysis* (by Bo Pang and Lillian Lee)
- *Sentiment Analysis and Opinion Mining, by Bing Liu, 20*

Sitografia:

- SAG, Univ. Roma Tor Vergata: http://sag.art.uniroma2.i