Opinion Mining

Feiyu XU & Xiwen CHENG

feiyu@dfki.de
DFKI, Saarbruecken, Germany
January 4, 2010

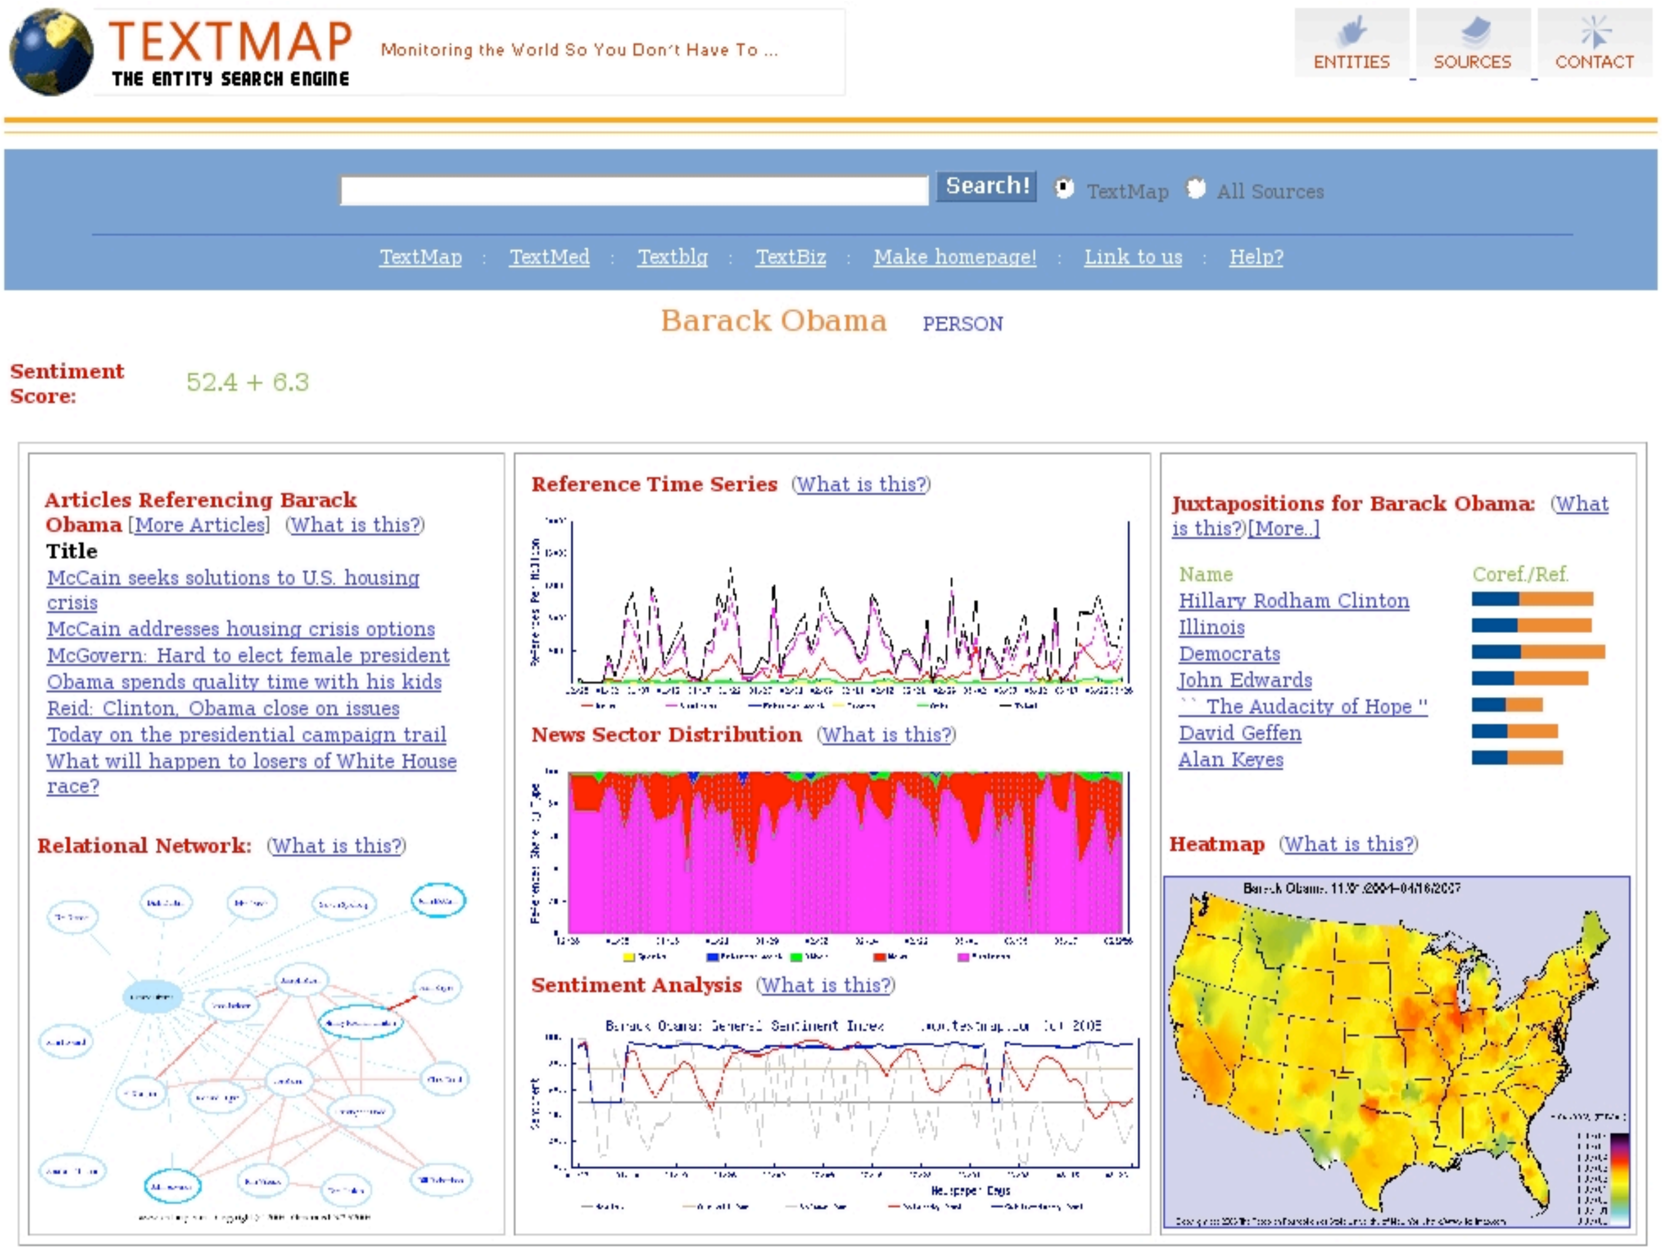
TEXTMAP
THE ENTITY SEARCH ENGINE
Monitoring the World So You Don't Have To ...
ENTITIES
SOURCES
CONTACT
Search!
TextMap
All Sources
TextMap : TextMed : Textblg : TextBiz : Make homepage! : Link to us : Help?
Barack Obama PERSON
Sentiment Score: 52.4 + 6.3
Articles Referencing Barack Obama [More Articles] (What is this?)
Title
McCain seeks solutions to U.S. housing crisis
McCain addresses housing crisis options
McGovern: Hard to elect female president
Obama spends quality time with his kids
Reid: Clinton, Obama close on issues
Today on the presidential campaign trail
What will happen to losers of White House race?
Relational Network: (What is this?)
Reference Time Series (What is this?)
News Sector Distribution (What is this?)
Sentiment Analysis (What is this?)
Juxtapositions for Barack Obama: (What is this?)[More..]
Name
Coref./Ref.
Hillary Rodham Clinton
Illinois
Democrats
John Edwards
`` The Audacity of Hope ''
David Geffen
Alan Keyes
Heatmap (What is this?)

Outline

- Introduction
  - Opinion Mining
  - Linguistic Perspectives
  - Applications
- Opinion Mining
  - Abstraction
  - Linguistic Resources of OM
  - Document, Sentence, Clause Level Sentiment Analysis
  - Feature-based Opinion Mining and Summarization
  - Comparative Sentence and Relation Extraction
- Conclusion
  - Resources
  - Challenges

Introduction – What is an opinion?

- **[Quirk et al., 1985]**

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

- **Wikipedia**

  a person's ideas and thoughts towards something.

  It is an assessment, judgment or evaluation of something.

  An opinion is not a fact, because opinions are either not falsifiable, or the opinion has not been proven or verified.

  If it later becomes proven or verified, it is no longer an opinion, but a fact. Accordingly, all information on the web, from a surfer's perspective, is better described as opinion rather than fact.

Introduction – What is Opinion Mining

- A recent discipline at the crossroads of information retrieval, text mining and computational linguistics which tries to detect the opinions expressed in the natural language texts.

- Opinion Extraction is a specified method of information extraction, delivering inputs for opinion mining

- Sentiment analysis and sentiment classification are sub-areas of opinion extraction and opinion mining

Introduction – Examples

- *John is successful at tennis*.
- *John is never successful at tennis*.
- *Mary is a terrible person. She is mean to her dogs.*
- *It is sufficient*.
- *It is barely sufficient*.

Introduction – More Examples

- Tense
  - E.g. *This is my favorable car*.
  - E.g. *This was my favorable car*.
- Collocation
  - E.g. *It is expensive*. (about prize)
  - E.g. *It looks expensive*. (about appearance)
- Irony
  - E.g. *The very brilliant organizer failed to solve the problem*.
  - E.g. *Terrorists deserve no mercy!*

Introduction – More Examples

- Discourse-level opinions
  - Connectors
    - E.g. Although Boris is brilliant at math, he is a horrible teacher.
  - Discourse Structure: Lists and elaborations
    - E.g. *The 7 Series is a large, well-furnished luxury sedan. The iDrive control system, which uses a single knob to control the audio, navigation, and phone systems, is meant to streamline the cabin, but causes frustration. A midcycle freshening brought revised styling, a 4.8-liter, 360-hp V8, and a new name: the 750i. The six-speed automatic shifts smoothly.*
  - Multi-entity Evaluation
    - E.g. *Coffee is expensive, but Tea is cheap.*
  - Comparative
    - E.g. *In market capital, Intel is way ahead of AMD.*

Introduction – More Examples

- Discourse-level opinions
  - Reported Speech
    - E.g. *Mary was a slob*. Vs. *John said that Mary was a slob*.
  - Subtopics
    - E.g. The economic situation is more than satisfactory. The leading indicators show a rosy picture. When one looks at the human rights picture, one is struck by the increase in arbitrary arrests, by needless persecution of helpless citizens and increase of police brutality.
  - Genre Constraints
    - E.g. This film should be brilliant. The characters are appealing. Stallone plays a happy, wonderful man. His sweet wife is beautiful and adores him. He has a fascinating gift for living life fully. It sounds like a great story, however, the film is a failure.

Introduction – Applications [Liu, 2007]

- Market Intelligence: product, event and service benchmarking
  - Consumer opinion summarization
    - *E.g. Which groups among our customers are unsatisfied? Why?*
  - Public opinion identification and direction
    - *E.g. What are the opinions of the Americans about the European style cars?*
  - Recommendation
    - *E.g. New Beetles is the favorite car of the young ladies.*
  - Consultants
  - Virtual sale experts
  - Marketing predication

- Opinion retrieval / search
  - Opinion-oriented search engine
  - Opinion-based question answering
    - *E.g. What is the general opinion on the proposed tax reform?*
  - Sentiment-enhanced machine translation

Outline

- Introduction
  - Opinion Mining
  - Linguistic Perspective
  - Application
- Opinion Mining
  - ⇨ Abstraction
  - ⇨ Acquisition of sentiment words and their orientation
  - ⇨ Document, Sentence, Clause Level Sentiment
  - ⇨ Analysis
  - ⇨ Feature-based Opinion Mining and Summarization
  - Comparative Sentence and Relation Extraction
- Conclusion
  - Resource
  - Challenges

Opinion Mining – Basic components

[Liu, Web Data Mining book 2007]

- Opinion holder: a person, a group or an organization that holds a specific opinion on a particular object
- Object: a product, person, event, organization, topic or even an opinion.
- Opinion: a view, attitude, or appraisal on an object from an opinion holder. An opinion contains often sentiment words which can be classified into polarities such as *Positive, Negative, Neutral.*

E.g. *John said that Mary was a slob*.

E.g. Gas mileage of VW Golf is great !

Opinion Mining – Model of a review [Liu, Web Data Mining book 2007]

- An object O 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
  - Another word, 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 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

OM – Research topics

- Development of linguistic resources for OM
  - Automatically build lexicons of sentiment terms and determine their orientations
- At the document/sentence/clause level
  - Simple opinion extraction (one holder, one object, one opinion)
  - Subjective / objective classification
  - Sentiment classification: positive, negative and neutral
  - Strength detection of opinions from clauses
- At the feature level
  - Identify and extract commented features
  - Group feature synonyms
  - Determine sentiments towards these features
- Comparative opinion mining
  - Identify comparative sentences
  - Extract comparative relations from these sentences

OM – Automatic Acquisition of Sentiment Lexicon
[Esuli, 2006]

- Linguistic resource of OM are **opinion words or phrases** which are used as instruments for sentiment analysis. It also called polar words, opinion bearing words, subjective element, etc.

- Research words on this topic deal with three main tasks:
  - Determining term **orientation**, as in deciding if a given Subjective term has a Positive or a Negative slant
  - Determining term **subjectivity**, as in deciding whether a given term has a Subjective or an Objective (i.e. neutral, or factual) nature.
  - Determining the **strength** of term attitude (either orientation or subjectivity), as in attributing to terms (real-valued) degrees of positivity or negativity.

- Example
  - Positive terms: *good, excellent, best*
  - Negative terms: *bad, wrong, worst*
  - Objective terms: *vertical, yellow, liquid*

Orientation of terms [Esuli, 2006]

The problem:

Determining if a subjective term has a Positive or a Negative orientation.

[Hatzivassiloglou and McKeown, 1997]

Hypothesis: *adjectives* in and conjunctions usually have similar orientation, though but is used with opposite orientation.

Example (conjuction of adjectives)

1. The tax proposal was *simple and well received...*
2. The tax proposal was *simplistic but well received...*
3. * The tax proposal was *simplistic and well received...*

Method: a weighted graph of similarity of orientation is defined by analyzing conjunctions of adjectives in unprocessed text, then a minimum-cut method is applied to the graph.

Orientation of terms [Esuli, 2006]

[Turney and Littman, 2003]

Hypothesis: terms with similar orientation tend to co-occur in documents.

The *Semantic Orientation* (SO) of a term is estimated by combining a pointwise mutual information (PMI) measure of the term against some *paradigmatic* terms.

Pos = {`good`, `nice`, `excellent`, `positive`, `fortunate`, `correct`, `superior`}
Neg = {`bad`, `nasty`, `poor`, `negative`, `unfortunate`, `wrong`, `inferior`}

PMI is measured using the number of results returned by the AltaVista search engine.

$$\begin{aligned} PMI(t, t_i) &= \log \frac{\#(\text{“}t \ \texttt{NEAR} \ t_i\text{”})}{\#(\text{“}t\text{”})\#(\text{“}t_i\text{”})} \\ SO(t) &= \Sigma_{t_i \in Pos} PMI(t, t_i) - \Sigma_{t_i \in Neg} PMI(t, t_i) \end{aligned}$$

OM – Polarity word lexicon acquisition

- Application:
  - Naive solution to achieve prior polarities

- Problems:
  - Mixture of subjective & objective words
    - E.g. *long & excellent*
  - Conflict
    - E.g. Nice and Nasty (the first hit from Google for “Nice and *”)
  - Context dependent
    - E.g. It looks cheap. It is cheap.
    - E.g. It is expensive. It looks expensive.

Orientation of terms [Esuli, 2006]

[Esuli and Sebastiani, 2005]

Hypothesis: terms with similar orientation have similar glosses.

Example (glosses for terms with similar orientation)

`good`: "*that which is pleasing or valuable or useful*";
"*agreeable or pleasing*".

`beautiful`: "*aesthetically pleasing*".

`pretty`: "*pleasing by delicacy or grace; not imposing*".

Each term is represented by its gloss.

A binary classifier is learned, in a semi-supervised process, using the glosses of the Positive and Negative terms in the training set.

OM – Research topics

- Development of linguistic resources for OM
  - Automatically build lexicons of subjective terms
- At the document/sentence/clause level
  - Simple opinion extraction (one holder, one object, one opinion)
  - Subjective / objective classification
  - Sentiment classification: positive, negative and neutral
  - Strength Detection of opinions from clauses
  - * Less information, more challenges
- At the feature level
  - Identify and extract commented features
  - Determine the sentiments towards these features
  - Group feature synonyms
- Comparative opinion mining
  - Identify comparative sentences
  - Extract comparative relations from these sentences

Document Level Sentiment Analysis - Approaches

- Unsupervised review classification
  - Turyney, 2003
- Sentiment classification using machine learning methods
  - Pang et al., 2002, Pang and Lee, 2004, Whitelaw et al., 2005
- Review classification by scoring features
  - Dave, Lawrence and Pennock, 2005

OM – Document Level Sentiment Analysis

- Motivation: determination of the overall sentiment properties of a text
- Advantage
  - Coarse-grained Analysis
  - Detection of a general sentiment trend of a document
- Problems
  - Different polarities, different topics and different opinion holders in one document, e.g.,

    This film should be brilliant. The characters are appealing. Stallone plays a happy, wonderful man. His sweet wife is beautiful and adores him. He has a fascinating gift for living life fully. It sounds like a great story, however, the film is a failure.

Unsupervised review classification

- Hypothesis: the orientation of the whole document is the sum of the orientation of all its parts
- Three steps
  - POS Tagging and Two consecutive word extraction (e.g. JJ NN)
  - Semantic orientation estimation (AltaVisata near operator)
    - 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(phrase) = PMI(phrase, “excellent”) – PMI(phrase, “poor”)
  - Average SO Computation of all phrases
    - The review is recommended if average SO is positive, not recommended otherwise
- The average accuracy on 410 reviews is 74%, ranging from 84% for automobile reviews to 66% for movie reviews

Other Approaches

- [Pang et al., 2002]
  - Apply some standard supervised automatic text classification methods to classify orientation of movie reviews
    - Learners: Naive Bayes, MaxEnt, SVM
    - Features: unigrams, bigrams, adjective, POS, position
  - 82.9% accuracy, on a 10-fold cross validation experiments on 1,400 movie reviews (best from SVM, unigrams, binary)

- [Pang and Lee, 2004]
  - A sentence subjectivity classifier is applied, as preprocessing, to reviews, to filter out Objective sentences.
  - Accuracy on movie reviews classification raises to 86.4%

- [Whitelaw et al. 2005]
  - Appraisal features are added to the Movie Review Corpus, which obtained a 90.2% classification accuracy.

OM – Sentence-level Sentiment Classification

- Advantage:
  - Even though the analysis is still coarse, it is more specific and precise than document-level analysis
  - The results can be reused as input for document-level classification

- Problem:
  - Multiple sentiment expressions with different polarities, e.g.,

The very brilliant organizer failed to solve the problem.

OM – Sentence Level Sentiment Analysis (cont.)

- [Rilloff and Wiebe, 2003]: subjective / objective classification
  - Taking advantages of Information Extraction techniques
  - Manually collected opinion words + AutoSlog-TS

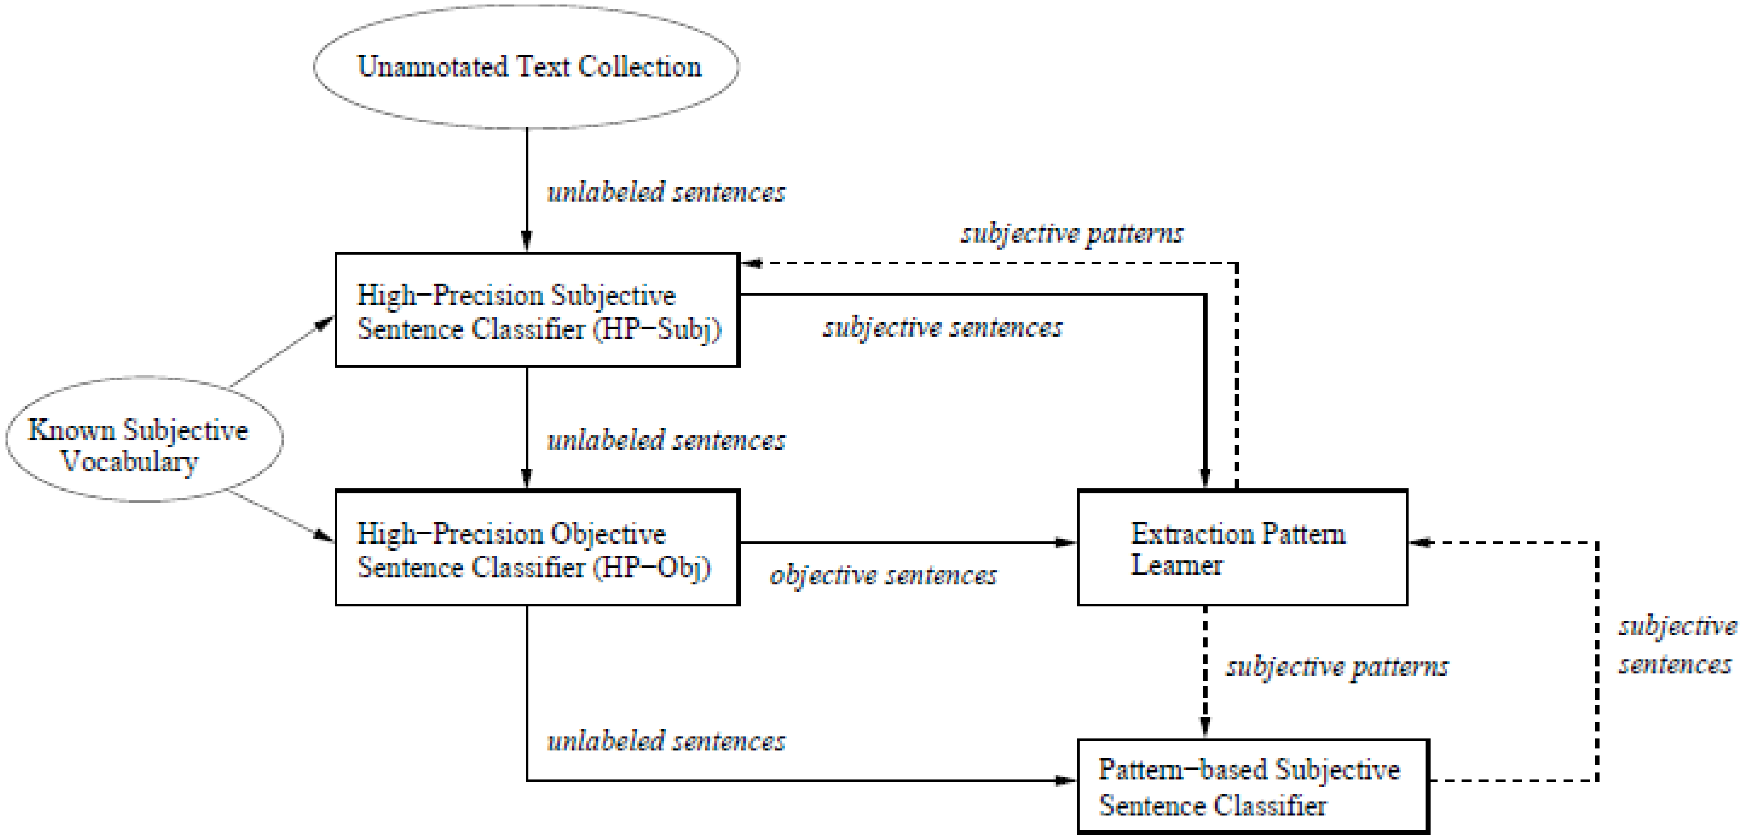

Figure 1: Bootstrapping Process

<subject> passive-vp	*<subj> was satisfied*
<subject> active-vp	*<subj> complained*
<subject> active-vp dobj	*<subj> dealt blow*
<subject> active-vp infinitive	*<subj> appears to be*
<subject> passive-vp infinitive	*<subj> was thought to be*
<subject> auxiliary dobj	*<subj> has position*
active-vp <dobj>	*endorsed <dobj>*
infinitive <dobj>	*to condemn <dobj>*
active-vp infinitive <dobj>	*get to know <dobj>*
passive-vp infinitive <dobj>	*was meant to show <dobj>*
subject auxiliary <dobj>	*fact is <dobj>*
passive-vp prep <np>	*opinion on <np>*
active-vp prep <np>	*agrees with <np>*
infinitive prep <np>	*was worried about <np>*
noun prep <np>	*to resort to <np>*

OM – Feature-based OM and Summarization

[Hu and Liu, 2004]

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. …

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.

Feature extraction:

- Explicit & Implicit
  - E.g. great photos: <photo>
  - E.g. something smaller: <size>
  - E.g. is expensive: <prize>
- Frequent & Infrequent

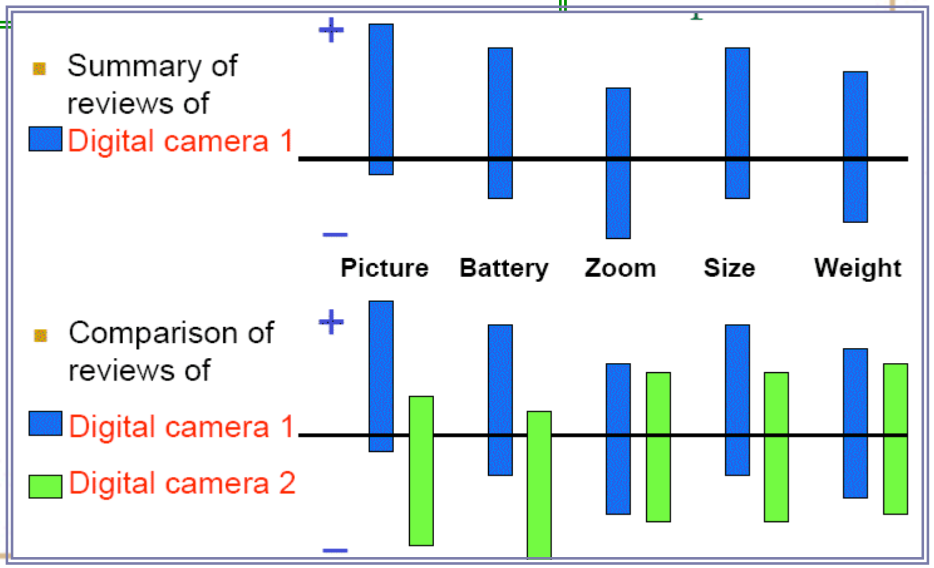

Featured-based – Feature Extraction

- Frequent & Infrequent features
  - Frequent feature
    - Feature extraction by applying lexico-syntactic patterns, e.g.,

      "Included memory is stingy"

      <{included, VB}{$feature, NN}{is, VB}{stingy, JJ}>

  - Infrequent feature
    - Observation: the same *opinion word* can be used to describe different features and objects
      - E.g. The pictures (high-freq) are absolutely amazing.
      - E.g. The software (low-freq) that comes with it is amazing.

Featured-based – Group Feature Synonyms

- Identify part-of relationship
  - [Popescu and Etzioni, 2005]: Each noun phrase is given a PMI score with meronymy discriminators (e.g. “*of scanner*”, “*scanner has*”) associated with the product class, (e.g. *a scanner class*)
  - [Liu et al., 2005] use WordNet

- Camera
  - Lens
    - Digital Zoom
    - Optical Zoom
    - ...
  - Editing/Viewing
    - Viewfinder
    - ...
  - Flash
  - ...
- Image
  - Image Type
    - TIFF
    - JPEG
    - ...
  - Resolution
    - Effective Pixels
    - Aspect Ratio
  - ...

Feature Extraction and Group

- Advantage
  - Precise sentiment analysis about explicit features
- Challenges
  - Multiple relations: part-of, sentiment-feature
    - *Gas Mileage of VW Golf is great.*
      - *Entity: VW Golf*
      - *Attribute: Gas Mileage*
  - Domain knowledge intensive:
    - *V12 8000CC is pretty powerful. (<automobile engine version>)*
    - *V6 4000CC is not a real good engine.*
  - WordNet is too general

OM – Research topics

- Development of linguistic resources for OM
  - Automatically build lexicons of subjective terms
- At the document/sentence/clause level
  - Assumption: each document, sentence or clause focuses on a single object and contains opinion (positive, negative and neutral) from a single opinion holder
  - Subjective / objective classification
  - Sentiment classification: positive, negative and neutral
  - Strength Detection of opinions from clauses
  - * **Less information, more challenges**
- At the feature level
  - Identify and extract commented features
  - Group feature synonyms
  - Determine the sentiments towards these features
- Comparative opinion mining
  - Identify comparative sentences
  - Extract comparative relations from these sentences

Featured-based Sentiment Orientation [Popescu and Etzioni, 2005]

- Context-dependent Semantic Orientation
  - <word, SO>, <word, feature, SO>, <word, feature, sentence, SO>
    - E.g. **SEN:*“I am not happy with this sluggish driver.”***
    - ***<sluggish, ?>, <sluggish, driver, ?>, <sluggish, driver,* SEN, *?>***

- Relaxation labeling: sentiment assignment to words satisfying local constraints.
  - *Constraints*:
    - conjunctions, disjunctions, syntactic dependency rule, morphological relationships, WordNet-supplied synonymy and antonymy, etc.
  - *Neighborhood*: a set of words connected the word through constraints.
    - E.g. “***hot(?) room and broken(-) fan***” → **hot**(-)

OM – Research topics

- Development of linguistic resources for OM
  - Automatically build lexicons of subjective terms
- At the document/sentence/clause level
  - Assumption: each document, sentence or clause focuses on a single object and contains opinion (positive, negative and neutral) from a single opinion holder
  - Subjective / objective classification
  - Sentiment classification: positive, negative and neutral
  - Strength Detection of opinions from clauses
  - * **Less information, more challenges**
- At the feature level
  - Identify and extract commented features
  - Group feature synonyms
  - Determine the sentiments towards these features
- Comparative opinion mining
  - Identify comparative sentences
  - Extract comparative relations from these sentences

OM – Comparative Sentence and Relation Extraction [Jinal and Liu, SIGIR-2006]

- Morphological and syntactic properties
  - Comparative sentences use morphemes like
    - *More/most, -er/-est, less/least, than*
  - Other cases
    - Preferring
      - *E.g. I prefer Intel to AMD.*
    - Non-comparatives with comparative words
      - *E.g. In the context of speed, faster means better.*
- Gradable
  - Non-Equal Gradable: greater or less
    - *E.g. Optics of camera A is better than that of camera B.*
  - Equative
    - *E.g. Camera A and camera B both come in 7MP.*
  - Superlative
    - *E.g. Camera A is the cheapest camera available in market.*
- Non-gradable
  - *E.g. Object A has feature F, but object B does not have.*

OM – Comparative Sentence and Relation Extraction

- Definition: A gradable comparative relation captures the essence of a gradable comparative sentence and is represented with the following:
  (relation word, features, entity S1, entity S2, type)
  - Relation word: The keyword used to expressed a comparative relation in a sentence. *E.g. better, ahead, most, better than*
  - Features: a set of features being compared
  - Entity S1 and Entity S2: sets of entities being compared
  - Type: non-equal gradable, equative or superlative

- Examples
  - *Car X has better controls than car Y.*
    - (better, controls, car X, car Y, non-equal-gradable)
  - *Car X and car Y have equal mileage.*
    - (equal, mileage, car X, car Y, equative)
  - *Car X is cheaper than both car Y and car X.*
    - (cheaper, null, car X, car Y car Z, non-equal-gradable)
  - *Company X produces a variety of cars, but still best cars come from company Y.*
    - (best, cars, company Y, null, superlative)

OMINE – Opinion Mining System

Fine-grained Opinion Topic and Polarity Identification (Cheng & Xu, 2008)

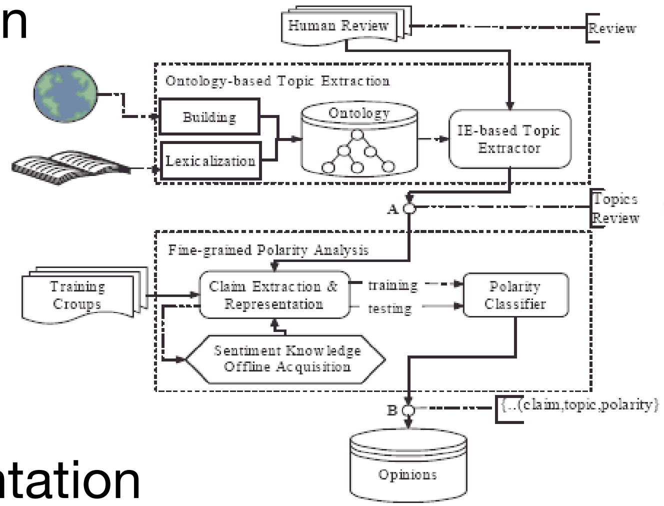

- Ontology-based Topic Extraction
  - Offline Ontology Building
  - Ontology Lexicalization
  - IE-based Topic Extraction

- Fine-grained Polarity Analysis
  - Claim Extraction & Representation
  - Offline Acquisition of Sentiment Knowledge
  - Polarity Analysis

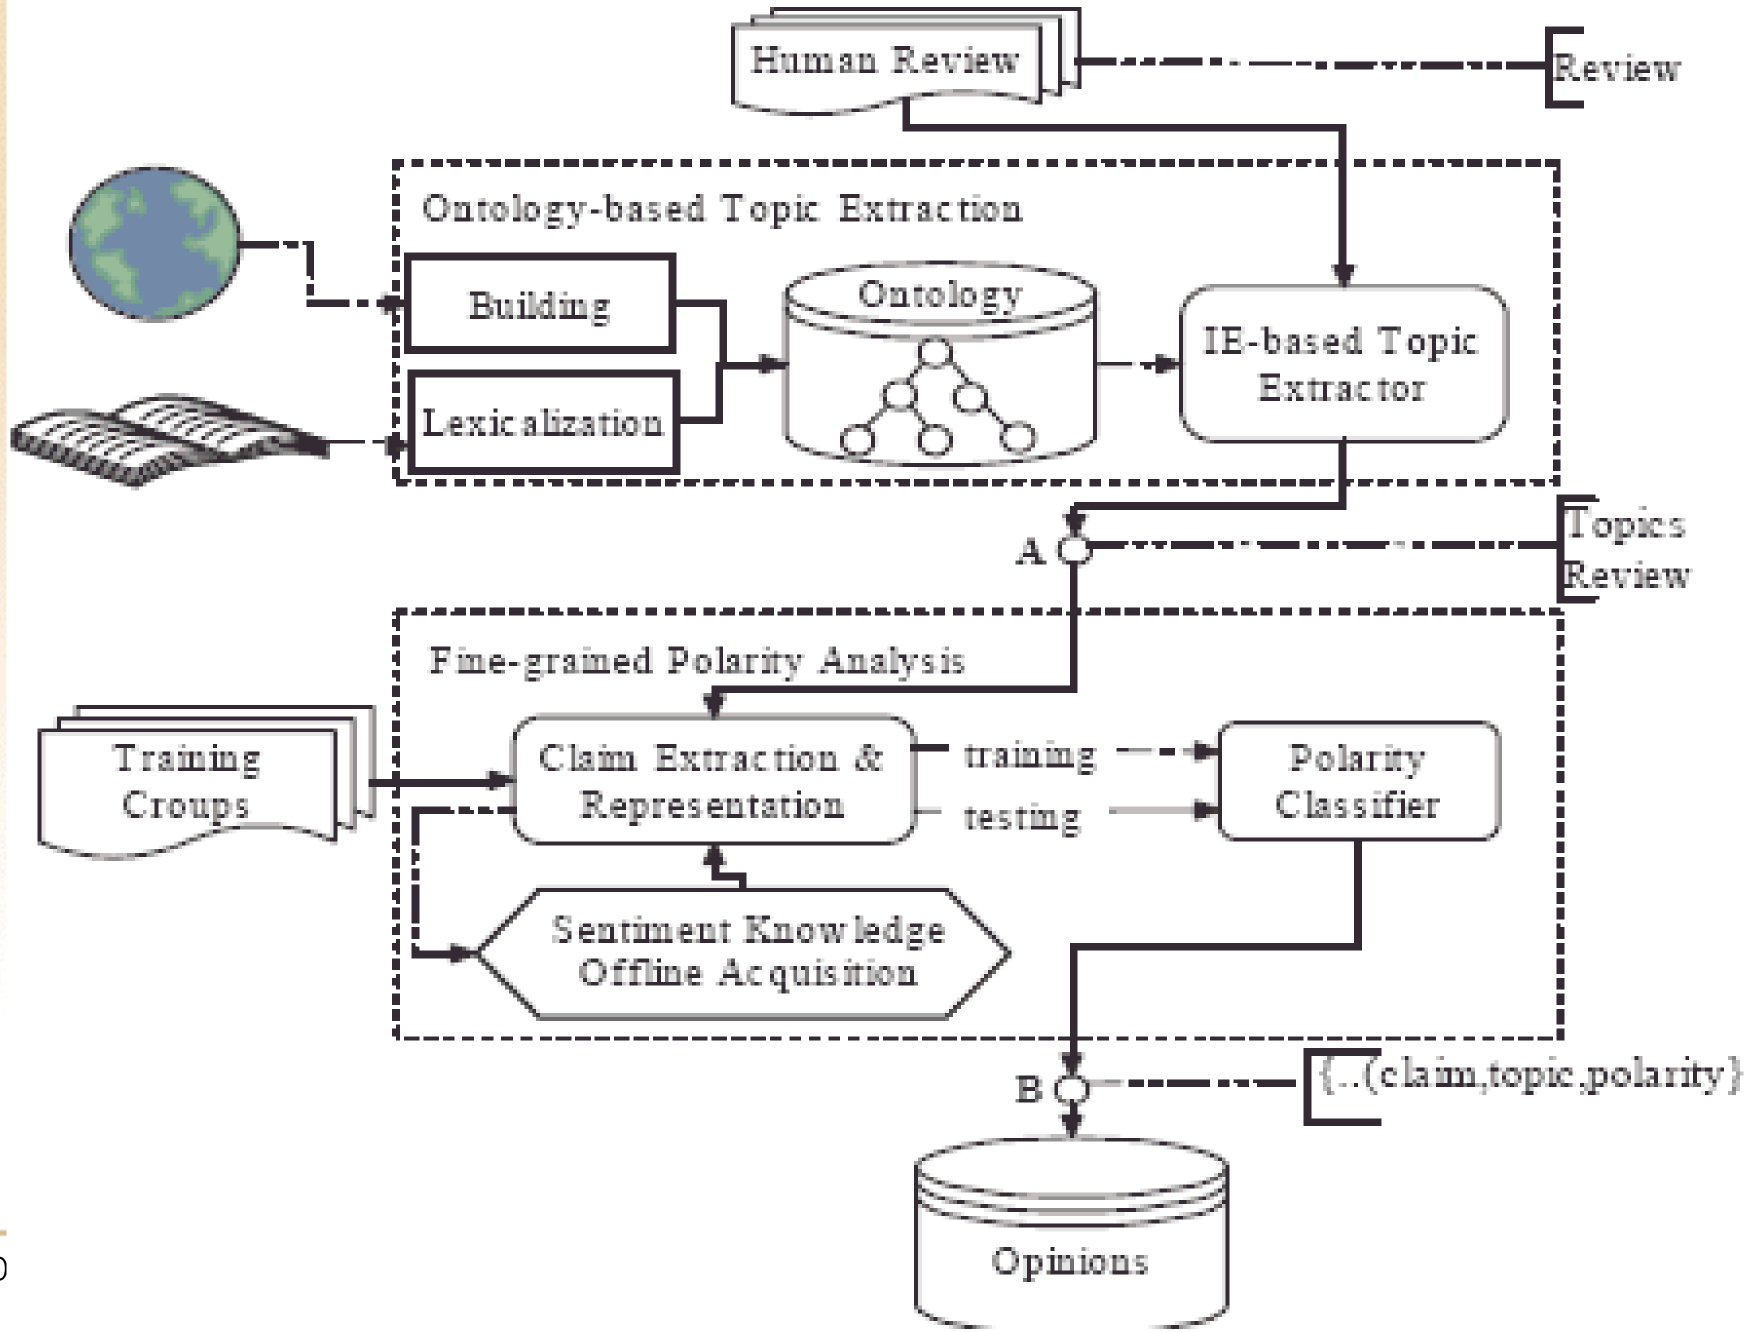

Human Review
Review
Ontology-based Topic Extraction
Building
Lexicalization
Ontology
IE-based Topic Extractor
A
Topics
Review
Fine-grained Polarity Analysis
Training Croups
Claim Extraction & Representation
training
testing
Polarity Classifier
Sentiment Knowledge Offline Acquisition
B
{..(claim,topic,polarity}
Opinions

Topic Extraction - Experiment

- Data
  - Taxonomy Resource: eBay http://www.ebay.com and AutoMSN http://autos.msn.com
  - Automobile glossary: http://www.autoglossary.com, around 10,000 terms
  - Data for topic extraction: 1000 sentences from UserReview of AutoMSN
  - Golden standard: 2038 terms identified manually

- CarOnto
  - 363 concepts (e.g. Air Intake & Fuel Delivery)
  - 1233 instances (e.g. 5- speed automatic overdrive)
  - 145 values (e.g. wagon for Style, 250@5800 RPM for Horsepower)
  - 803 makes and models (e.g. BMW, Z4)
  - Ontology lexicalization is applied to 363 concepts and retrieves 9033 lexicons.
  - 11214 domain-specific lexicon instances as total

- Topic Extraction
  - TermExtractor (Sclano and Velardi, 2007)
  - OPINE (Popescu and Etzioni, 2005)

OntoTpcEx	Recall	Precision
Before Enrichment	20.97%	88.12%
After Enrichment	**89.35%**	**94.44%**
TermExtractor	15.72%	97.46%
OPINE	79.44%	93.12%

Polarity Analysis- Experiment

- Data
  - Resource: UserReview From AutoMSN
  - The polarities of these reviews have already been annotated by reviewers in two classes: pro and con.
  - Around 20 thousand sentences, and 50% of them are positive and the other 50% are negative.
  - 19600 sentences are used to train the classifier, and 200 positive and 147 negative sentences are applied as a test corpus
- Acquisition of Sentiment Knowledge

Type	Precision	Num
Sentiment word	95.0%	623
Negation word	73.8%	22

POSITIVE	NEGATIVE
awesome, cute, speedy	unimpressive, awful, terrible
excellent, well, standard	useless, tremendous, costly
great, strong, comfortable	expensive, troublesome, tight
sporty, super, adorable	cumbersome, ugly, squeaky

POS	Negation words
aux	*doesn't, didn't, wouldn't, shouldn't couldn't, don't, can't, won't*
det	*no, little, least*
mod	*never, barely, not, less*

Challenges

- Interaction between Pattern and Slot
  - <holder> **would like better** <object>
    - I would like better fuel mileage.
  - <object -1> drives like <**object-2**>
    - This car drives like a Porsche/a Nissan.

- Anaphoric resolution for summarization
  - E.g. “The turbo engine is a must-have, which provide a very decent acceleration.”

- Others (context or semantic implication)
  - He is not the sharpest knife in the drawer.
  - Stephanie McMahon is the next Stalin.
  - No one would say that John is smart.
  - My little brother could have told you that.
  - You are no Jack Kennedy.
  - They have not succeeded, and will never succeed, in breaking the will of this valiant people.

- More …

“Who is the fairest one of all, and state your sources!”

Reference

- Slides
  - http://medialab.di.unipi.it/web/Language+Intelligence/OpinionMining06-06.pdf
  - http://www.cs.uic.edu/~liub/opinion-mining-and-search.pdf
- References

  - *Hatzivassiloglou, Vasileios and Kathy McKeown. 1997. Predicting the semantic orientation of adjectives. In Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics (ACL-97), pages 174–181, Madrid, Spain.*
  - *Hu, Minqing and Bing Liu. 2004. Mining and summarizing customer reviews. In Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining 2004 (KDD 2004), pages 168–177, Seattle, Washington.*
  - *A. Popescu, "Extracting Product Features and Opinions from Reviews", Oren Etzioni, Proceedings of HLT-EMNLP, 2005*
  - *Wilson, Theresa, Janyce Wiebe, and Paul Hoffman. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In*
  - *Proceedings of the Human Language Technologies Conference/Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP-2005), pages 347–354, Vancouver, Canada.*
  - *X. Cheng, OMINE: Automatic Topic Term Detection and Sentiment Classification for Opinion Mining, Master thesis, 2007*
  - *…*

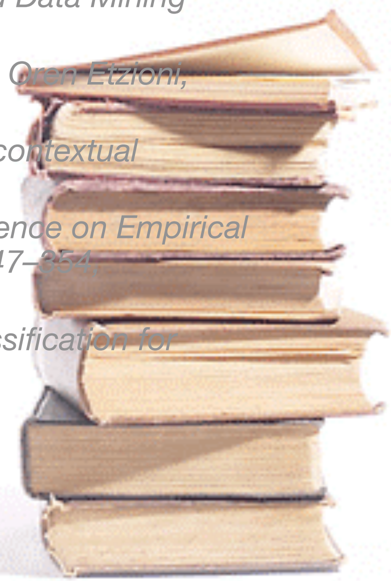


köszönöm
תודה!
děkuji
mahalo
고맙습니다
thank you
merci
谢谢
danke
Ευχαριστώ
شكرا
どうもありがとう
gracias