

01 – Opinion mining, sentiment analysis

IA161 Advanced Techniques of Natural Language Processing

Z. Nevěřilová

NLP Centre, FI MU, Brno

September 21, 2016

Opinion mining, sentiment analysis

So boring. I enjoyed the first book but this one really didn't work for me. The story, characters, and relationships all fell flat.

Lair of Dreams like everything else Miss Bray writes is mind-boggling. It's big. It's insanely atmospheric and it's creepastic.

–goodreads.com

Opinion mining, sentiment analysis

So *boring*. I *enjoyed* the first book but this one really *didn't work* for me. The *story*, *characters*, and *relationships* all fell *flat*.

Lair of Dreams like everything else Miss Bray writes is *mind-boggling*. It's *big*. It's *insanely atmospheric* and it's *creeptastic*.

–goodreads.com

this book: boring
first book: enjoyed
this book: did not work
story: flat
characters: flat
relationships: flat

Lair of Dreams: mind-boggling
LoD: big
LoD: insanely atmospheric
LoD: creepstastic

1 Opinion mining, sentiment analysis

2 Applications of opinion mining

3 Problem definition

4 Methods

Opinion mining, sentiment analysis

Opinion mining / sentiment analysis:

*Given a set of **subjective** texts that express opinions about a certain **object**, the purpose is to extract those **attributes** (features) of the object that have been commented on in the given texts and to **determine** whether these texts are positive, negative or neutral. [Dinu and Iuga, 2012]*

Automatic opinion mining: why?

- many subjective texts exist
- mostly because of social media
 - ▶ people express their opinions in texts
 - ▶ one's opinions influence others' opinions
- emotions make part of a decision process (see [Minsky, 2007])

“Opinions” are key influencers of our behaviors. [Liu, 2012]

Opinion mining: applications



Opinion mining: applications

The screenshot shows the SEMANTRIA application running within Microsoft Excel. The ribbon includes 'SEMANTRIA' with sub-sections for 'Text and Sentiment Analysis configuration' and 'Reporting Options'. The main workspace displays two pie charts and a table of processed documents.

Entity Sentiment

- negative
- positive
- neutral

Categories Breakdown

- Hotels
- Video Games
- Marriage
- Automotive
- Beverages
- Advertising

Semantria for Excel

Configuration: English (English)

Name	Configuration	Documents	Created
Clients feed...	English	(101/101)	2014-05-13 15:31
Clients feed...	English	(101/101)	2014-05-13 13:29
Clients feed...	English	(101/101)	2014-05-10 14:02

Input text: thanks! have a great weekend. rt @tabishb: #ff amfunderburk1

Opinion mining: applications

The screenshot displays a web application interface for sentiment analysis. The main content area shows the following information:

- MALL.cz** (Logo)
- ★ ★ ★ ★ ★ 91% (5 stars)
- 4295 hodnocení (4295 reviews)
- Hodnocení** (Reviews)
- Přidat hodnocení** (Add review)
- jirichott** ★ ★ ★ ★ ★
Kritika: Nakupovat na MALL.cz už nikdy! Jednou jsem to zkusil a byla to katastrofa. NEDOPORUČUJI!!!!
- Lucie.P** ★ ★ ★ ★ ★ (ověřený zákazník)
Chvála: velice spokojená
- Jerry** ★ ★ ★ ★ ★ (ověřený zákazník)
Kritika: Místo objednaného zboží došlo něco zcela jiného, což jsem zjistil po několika týdnech, že zboží vůbec nemají a tudíž jsem musel objednat znovu. :(

The sidebar on the left contains a pie chart titled "Entity Sentiment" with a legend for "negative" (red) and "positive" (green). Below the chart is a table with the following data:

Entity	Count
CENA A TERMÍN DODÁNÍ	4237
KOMUNIKACE ESHOPU	4163
OBSAH ZÁSILKY	4113
ZKUŠENOST S VRÁCENÍM ZBOŽÍ	579
ZKUŠENOST S REKLAMACÍ	518

Below the table, there is a section for "Platba" (Payment) with the following text:

Platba: On-line platby (PaySec, Raiffeisen banka),
Platba kartou (Euro Card, Maestro, Master Card, VISA,

Opinion mining: related applications

- document sentiment classification:
This document contains a lot of negative statements.
- sentence subjectivity classification:
This sentence is objective.
- aspect-based opinion summarization:
Most customers of your company think that the communication is not good.
- mining comparative opinions:
Many people think that iPhone is better than SG.
- utility or helpfulness of reviews:
This review is useless.
- cross-lingual opinion mining

Problem definition

What is an opinion?

- an evaluating proposition: *Linux is great.*
- a comparative proposition: *Linux is better than Windows.*

Problem definition

What is an opinion?

- an evaluating proposition: *Linux is great.*
- a comparative proposition: *Linux is better than Windows.*

*An opinion is simply a **positive or negative** sentiment, view, attitude, emotion, or appraisal about an **entity** or an **aspect of the entity** from an **opinion holder**. [Liu, 2012]*

Problem definition

What is an opinion?

- an evaluating proposition: *Linux is great.*
- a comparative proposition: *Linux is better than Windows.*

*An opinion is simply a **positive or negative** sentiment, view, attitude, emotion, or appraisal about an **entity** or an **aspect of the entity** from an **opinion holder**. [Liu, 2012]*

entity e is a product, person, event, organization, or topic: iPhone, Madonna, Microsoft ...

aspect a (feature) is a component of e or attribute of e: battery, price, appearance, communication skills ...

Problem definition

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

- e_j is a target entity.
- a_{jk} is an aspect/feature of the entity e_j .
- so_{ijkl} is the sentiment value of the opinion from the opinion holder h_i on feature a_{jk} of entity e_j at time t_l .
- h_i is an opinion holder.
- t_l is the time when the opinion is expressed.

Problem definition

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

- e_j is a target entity.
- a_{jk} is an aspect/feature of the entity e_j .
- so_{ijkl} is the sentiment value of the opinion from the opinion holder h_i on feature a_{jk} of entity e_j at time t_l .
- h_i is an opinion holder.
- t_l is the time when the opinion is expressed.

not just **one** problem

Problem definition

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

- e_j is a target entity.
named entity recognition
- a_{jk} is an aspect/feature of the entity e_j .
information extraction
- so_{ijkl} is the sentiment value of the opinion from the opinion holder h_i on feature a_{jk} of entity e_j at time t_l .
sentiment identification
- h_i is an opinion holder.
information extraction
- t_l is the time when the opinion is expressed.
information extraction

not just **one** problem

anaphora resolution + synonym matching

Problem definition

Generally, find structure in **unstructured** data (text)

- document level opinion mining: *The document is negative.*
- sentence level: *The sentence is negative.*
- object/entity and feature/aspect level: *iPhone is expensive.*

Problem definition

Generally, find structure in **unstructured** data (text)

- document level opinion mining: *The document is negative.*
- sentence level: *The sentence is negative.*
- object/entity and feature/aspect level: *iPhone is expensive.*

Classification task:

- 2-classes: positive/negative
- 3-classes: positive/negative/neutral
- 5-classes ...


A hard problem (sometimes)

- opinion mining in tweets is relatively easy (short texts, hashtags) usually 3-classes classification for each tweet
- opinion mining in reviews is harder but still the form contains aspects and the reviewer has to mark the review positive/negative usually 2-classes classification for each aspect (e.g. high price)
- opinion mining in discussions, comments, blogs is very hard

sentiment lexicon

evaluative words: nice, cool, shit, bad. . .

SentiWordNet [Baccianella et al., 2010]



Positive: 0 Objective: 0.125 Negative: 0.875

blue = filled with melancholy and despondency

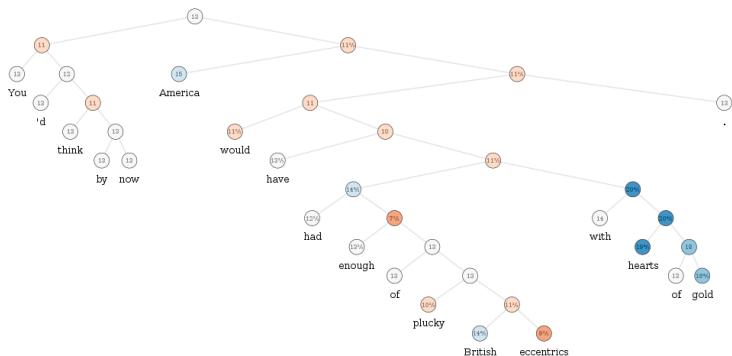
A hard problem (sometimes)

- opinion mining in tweets is relatively easy (short texts, hashtags) usually 3-classes classification for each tweet
- opinion mining in reviews is harder but still the form contains aspects and the reviewer has to mark the review positive/negative usually 2-classes classification for each aspect (e.g. high price)
- opinion mining in discussions, comments, blogs is very hard

evaluative word	aspect	sentiment
thin	phone	good
thin	steak	bad
high	value	good
high	price	bad

A hard problem (sometimes)

- opinion mining in tweets is relatively easy (short texts, hashtags) usually 3-classes classification for each tweet
- opinion mining in reviews is harder but still the form contains aspects and the reviewer has to mark the review positive/negative usually 2-classes classification for each aspect (e.g. high price)
- opinion mining in discussions, comments, blogs is very hard



Opinion mining methods: supervised machine learning

- 1 get example data with labels
- 2 extract features from the data, i.e. convert the documents to feature vectors
- 3 train the parameters (choose an algorithm: SVM, Naive Bayes, Neural Networks ...)
- 4 test the model

Opinion mining methods: supervised machine learning

- 1 get example data with labels
- 2 extract features from the data, i.e. convert the documents to feature vectors
- 3 train the parameters (choose an algorithm: SVM, Naive Bayes, Neural Networks ...)
- 4 test the model

Opinion mining methods: supervised machine learning

[Dinu and Iuga, 2012] report best results on Naive-Bayes with tokens as features and bigrams as features

[Liu, 2012] reports best results with SVM on balanced (English) data
what features, feature extraction methods, training algorithm, parameters of training algorithm to use for Czech data?

let's see during the workshop

Opinion mining methods: state-of-the-art results

- OM on political tweets, [Maynard and Funk, 2012] report **78% precision** and **47% recall**
- on document level OM (movie reviews), [Richa Sharma and Jain, 2014] report **63% accuracy** and **70% recall**

References I



Baccianella, S., Esuli, A., and Sebastiani, F. (2010).

Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining.

In Chair), N. C. C., Choukri, K., Maegaard, B., Mariani, J., Odijk, J., Piperidis, S., Rosner, M., and Tapias, D., editors, *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta. European Language Resources Association (ELRA).



Dinu, L. P. and Iuga, I. (2012).

The Naive Bayes classifier in opinion mining: In search of the best feature set.

In Gelbukh, A., editor, *Computational Linguistics and Intelligent Text Processing*, volume 7181 of *Lecture Notes in Computer Science*, pages 556–567. Springer Berlin Heidelberg.

References II



Liu, B. (2012).

Sentiment analysis and opinion mining.

Synthesis Lectures on Human Language Technologies, 5(1):1–167.



Maynard, D. and Funk, A. (2012).

Automatic detection of political opinions in tweets.

In García-Castro, R., Fensel, D., and Antoniou, G., editors, *The Semantic Web: ESWC 2011 Workshops*, volume 7117 of *Lecture Notes in Computer Science*, pages 88–99. Springer Berlin Heidelberg.



Minsky, M. (2007).

The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind.

SIMON & SCHUSTER.

References III



Richa Sharma, S. N. and Jain, R. (2014).

Opinion mining of movie reviews at document level.

International Journal of Information Theory, 3(3):13–21.