Sentiment Analysis

Mining Opinions, Sentiments and Emotions

Van-Duyet LE (me@duyet.net)

06/2018

Outline

Introduction
- Sentiment Analysis Application
- Sentiment Analysis Research

Document Sentiment Classification

Sentence Sentiment Classification
- Overview
- Subjectivity classification
- Sentence Sentiment Classification

Aspect Sentiment Classification
- Overview
- Aspect Extraction
- Aspect Sentiment Classification

Introduction

Sentiment

- Sentiment = feelings
  - Attitudes
  - Emotions
  - Opinions
- Subjective impressions, not facts.
- For/against, like/dislike, good/bad, etc.

Introduction

Sentiment

- Sentiment = feelings
  - Attitudes
  - Emotions
  - Opinions
- Subjective impressions, not facts.
- For/against, like/dislike, good/bad, etc.

Sentiment analysis

is contextual mining of text which identifies and extracts subjective information in source material.
Using NLP, statistics, or machine learning methods to extract, identify, or otherwise characterize the sentiment content of a text unit.

Outline

Introduction
- Sentiment Analysis Application
- Sentiment Analysis Research

Document Sentiment Classification

Sentence Sentiment Classification
- Overview
- Subjectivity classification
- Sentence Sentiment Classification

Aspect Sentiment Classification
- Overview
- Aspect Extraction
- Aspect Sentiment Classification

Sentiment Analysis Application

- In the past, an organization or a business ~~conducted surveys~~, ~~opition polls~~ and ~~focus groups~~.
  - Helping a **business** to understand the social sentiment of their brand, product or service while monitoring online conversations.

Sentiment Analysis Application

- In the past, an organization or a business ~~conducted surveys~~, ~~opition polls~~ and ~~focus groups~~.
  - Helping a **business** to understand the social sentiment of their brand, product or service while monitoring online conversations.
- Governments can easily obtain public opitions about their policies and measure the pulses of other nations.

Sentiment Analysis Application

- In the past, an organization or a business ~~conducted surveys~~, ~~opition polls~~ and ~~focus groups~~.
  - Helping a **business** to understand the social sentiment of their brand, product or service while monitoring online conversations.
- Governments can easily obtain public opitions about their policies and measure the pulses of other nations.
- Opinionated documents of internal data: customer feedback, email, call centers, results of surveys, etc.

Sentiment Analysis Application

- In the past, an organization or a business ~~conducted surveys~~, ~~opition polls~~ and ~~focus groups~~.
  - Helping a **business** to understand the social sentiment of their brand, product or service while monitoring online conversations.
- Governments can easily obtain public opitions about their policies and measure the pulses of other nations.
- Opinionated documents of internal data: customer feedback, email, call centers, results of surveys, etc.
- Consumer products [1], healthcare [2], tourism, and financial services[3] to social events and political elections[4][5].

Sentiment Analysis Application

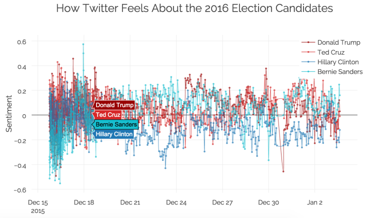

Figure: US Election 2016

food	positive
reservations	positive
value	positive
desserts	negative
location	positive
ambience	neutral
staff	positive
drinks	positive

Figure: SA for customer reviews

Outline

Introduction
- Sentiment Analysis Application
- Sentiment Analysis Research

Document Sentiment Classification

Sentence Sentiment Classification
- Overview
- Subjectivity classification
- Sentence Sentiment Classification

Aspect Sentiment Classification
- Overview
- Aspect Extraction
- Aspect Sentiment Classification

Different Levels of Analysis

Sentiment Analysis research has been mainly carried out at three levels of granularity:

- **Document level**
  - e.g. given a product review, overall positive or negative.
  - Assuming that each document expresses opinions on a single entity.
- **Sentence level**
- **Aspect level**
  - discover sentiment on entities and/or their aspects.
  - e.g. *"I like the iPhone X"*, *"Although the service is not great, I still love this hotel."*

Document Sentiment Classification

- *Document sentiment classification* detects the **overall opinion or sentiment** expressed in a document.
- It is perhaps the most extensively studied topic in the field of SA especially in its early days (see surveys by Pang and Lee, 2008 [6]; Liu, 2012 [7])
- It treat sentiment classification as **a traditinal text classification problem**.
- It **not** concerned the **targets** of sentiment or opinion.

Assumption

The opinion document `d` expresses opinions on a single entity `e` and contains opinions from a single opinion holder `h`

Document Sentiment Classification

- **Holds well for** online reviews of products or services (usually focus on single product or serice).
- **No meaningful for** blog posts, forum discussion (multiple opinions, multiple entities or compares).

Document Sentiment Classification

1. Supervised Sentiment Classification
   1.1 Using Machine Learning Algorithms
   1.2 Using a Custom Score Function
2. Unsupervised Sentiment Classification
   2.1 Using Syntactic Patterns and Web Search
   2.2 Using Sentiment Lexicons
3. Sentiment Rating Prediction
4. Cross-Domain Sentiment Classification
5. Cross-Language Sentiment Classification
6. Emotion Classification of Documents

1.1 Classification Using Machine Learning Algorithms

Any existing supervised learning method can be directly applied, such as Naive Bayes or SVM [8] [9] [10].

1.1 Classification Using Machine Learning Algorithms

Any existing supervised learning method can be directly applied, such as Naive Bayes or SVM [8] [9] [10].

Features engineering:

- *Terms and their frequency* highly effective, TFIDF weighting can be applied too.
- *Part of speech (POS)*
- *Sentiment words annd phrases* e.g. *good, wonderful, ...* are positive sentiment words.
- *Rules of opinion* using other constructs or language compositions.
- *Sentiment shiters* expressions that are used to change sentiment orientations (e.g. *"I don't like you"* is neg, although the word *like* is pos.
- *Syntactic dependency*

1.2 Classification Using a Custom Score Function

- Customized techniques specifically for sentiment classification or reviews.
- Example is the score function of Dave et al [11].

Step 1. Score each term (unigram or n-gram) in the training set using the following equation:

$$\mathrm{score}(t_i) = \frac{\Pr(t_i|C) - \Pr(t_i|C')}{\Pr(t_i|C) + \Pr(t_i|C')}, \tag{3.2}$$

Step 2. Classify a new document $d_i = t_1 \ldots t_n$ by summing up the scores of all terms and using the sign of the total to determine the class:

$$\mathrm{class}(d_i) = \begin{cases} C & \mathrm{eval}(d_i) > 0 \\ C' & \text{otherwise}, \end{cases} \tag{3.3}$$

where

$$\mathrm{eval}(d_i) = \sum_j \mathrm{score}(t_j). \tag{3.4}$$

2. Unsupervised Sentiment Classification

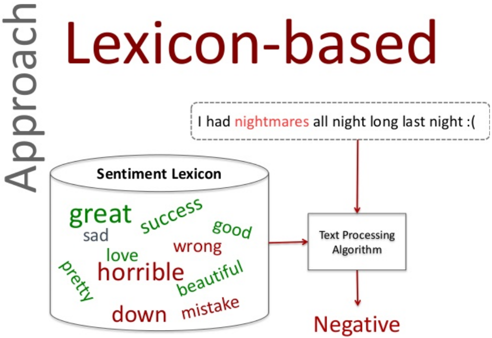

Figure: Staano[1]

[1] https://www.slideshare.net/Staano/senticircles-for-contextual-and-conceptual-semantic-sentiment-analysis-

2. Unsupervised Sentiment Classification (continue)

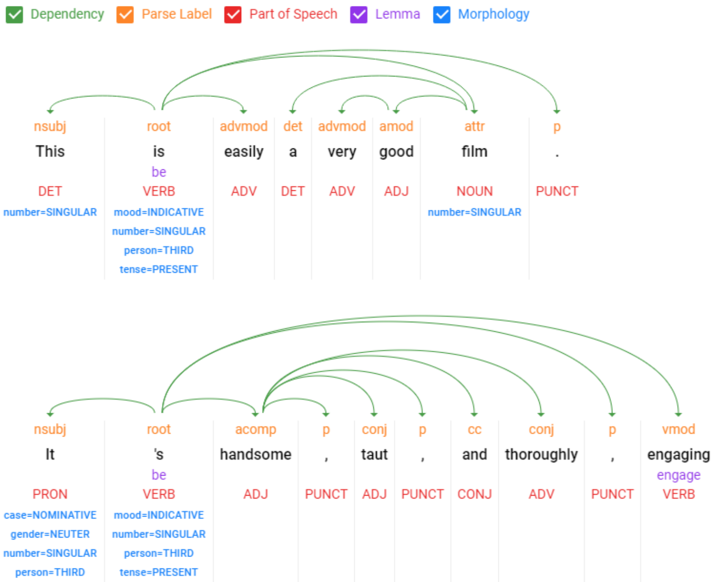

2. Unsupervised Sentiment Classification (continue)

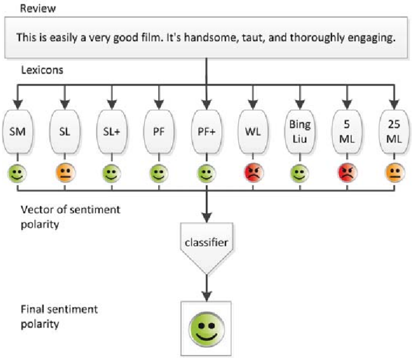

Figure: Augustyniak, Lukasz et al [12]. Simpler is better? Lexicon-based ensemble sentiment classification beats supervised methods. 2014 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2014) (2014): 924-929.

3. Classification Using Rating Prediction

- Using rating score (e.g., 1-5 stars) of reviews
  - Pang and Lee, 2005 [13] using SVM regression and SVM multiclass OVA.
  - Long et al. 2010 [14]: Bayesian network classifier

4. Cross-Domain Sentiment Classification

- Words and even langage constructs used to expressing opinions in different domains can be quite different.
- Existing research is mainly based on two settings:
  - a small amount of labeled training data for the new domain (Aue and Gamon, 2005 [15]).
  - No labeled data for the new domain (Blitzer et al., 2007[16]; Tan et al., 2007[17])

5. Cross-Language Sentiment Classification

Motivations:

- Apply existing SA (done and good) to another languages.
- Many apps, companies want to know and compare consumer opinions in different countries.

5. Cross-Language Sentiment Classification

Motivations:

- Apply existing SA (done and good) to another languages.
- Many apps, companies want to know and compare consumer opinions in different countries.

Some approaches:

- Wan (2008)[18] translate each Chinese review into English using multiple translators, classify and sums up sentiment score.
- Wan (2009)[19] using co-training method (SVM) and Wan (2013)[20] based on co-training idea, using co-regresion method.
- Boyd-Graber and Resnik (2010)[21] extended SLDA to MLSLDA.

6. Emotion Classification of Documents

Emotion Classification of Documents

Outline

Introduction
Sentiment Analysis Application
Sentiment Analysis Research

Document Sentiment Classification

Sentence Sentiment Classification
Overview
Subjectivity classification
Sentence Sentiment Classification

Aspect Sentiment Classification
Overview
Aspect Extraction
Aspect Sentiment Classification

Sentence Sentiment Classification - Overview

- Same with document level.
- The goal is to classify positive, negative or neutral*.

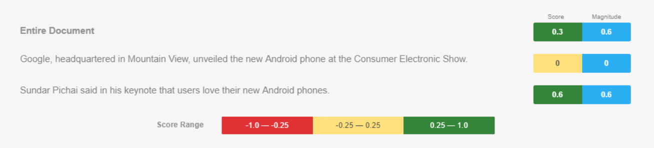

Sentence Sentiment Classification - Overview

- Same with document level.
- The goal is to classify positive, negative or neutral*.

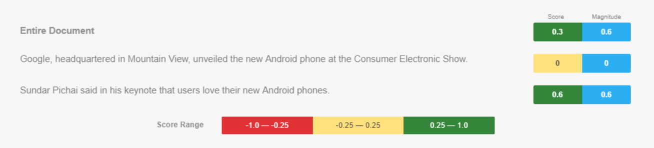

- Can be solved either as **(1) a three-class** classification or as **(2) two separate two-class** classification.

Outline

Introduction
- Sentiment Analysis Application
- Sentiment Analysis Research

Document Sentiment Classification

Sentence Sentiment Classification
- Overview
- Subjectivity classification
- Sentence Sentiment Classification

Aspect Sentiment Classification
- Overview
- Aspect Extraction
- Aspect Sentiment Classification

Subjectivity classification

(2) two separate two-class classification:

1. First step: classify whether a sentence expresses an opinion (*subjectivity classification*).
2. Second step: classifies those opinion sentences into positive and negative classes.

Subjectivity Analysis	**Sentiment Analysis**
Subjective	Positive
	Negative
Objective	Neutral

Subjectivity classification

Subjectivity classification classifies sentences into two classes, subjective and objective (Wiebe et al., 1999 [22]).
Most approaches are based on supervised or unsupervised learning:

- Weibe et al. (1999) [22] Naive Bayes
- Yu and Hatzivassiloglou (2003) [23] sentence similarity and Naive Bayes
- Pang and Lee (2004) [24] mincut-based

Outline

Introduction
- Sentiment Analysis Application
- Sentiment Analysis Research

Document Sentiment Classification

Sentence Sentiment Classification
- Overview
- Subjectivity classification
- Sentence Sentiment Classification

Aspect Sentiment Classification
- Overview
- Aspect Extraction
- Aspect Sentiment Classification

Sentence Sentiment Classification

Supervised learning again can be applied to solve the problem, and so can lexicon-based methods.

- Dealing with Conditional Sentences
- Dealing with Sarcastic Sentences
- Using Discourse Information for Sentiment Classification
- Emotion Classification of Sentences

Dealing with Conditional Sentences

Conditional sentences are sentences that describe implications or hypothetical situations and their consequences.

Example

"If your phone is not good, buy this iPhone"

Narayanan et al. (2009) [25] using a set of linguistic features.

POSITIVE	:: =	ENTITY is for you
	|	ENTITY is it
	|	ENTITY is the one
	|	ENTITY is your baby
	|	go (with | for) ENTITY
	|	ENTITY is the way to go
	|	this is it
	|	(search | look) no more
	|	CHOOSE ENTITY
	|	check ENTITY out
NEGATIVE	:: =	forget (this | it | ENTITY)
	|	keep looking
	|	look elsewhere
	|	CHOOSE (another one | something else)
CHOOSE	:: =	select | grab | choose | get | buy | purchase | pick | check | check out
ENTITY	: =	this | this ENTITY_TYPE | ENTITY_NAME

Dealing with Sarcastic Sentences

Sarcasm is a sophisticated form of speech act in which the speakers or the writers say or write the opposite of what they mean.

Example

"The Earth is full. Go home."
"Don't bother me. I'm living happily ever after."

- Tsur et al. (2010) [26] uses a small set of labeled sentences (seeds) and expands through web search.
- Gonzlez-Ibnez (2011) studied in Twitter data to distinguish sarcastic and nonsarcastic tweets (SVM, LR), they used unigrams and some dict-based information.

Discussion

- Sentence-level classifcation only suiable for simple sentences with a single opinion.
- Cannot deal with opinions in comparative sentences.
  - E.g. *"X is better than Y"*

Outline

Introduction
- Sentiment Analysis Application
- Sentiment Analysis Research

Document Sentiment Classification

Sentence Sentiment Classification
- Overview
- Subjectivity classification
- Sentence Sentiment Classification

Aspect Sentiment Classification
- Overview
- Aspect Extraction
- Aspect Sentiment Classification

Aspect Sentiment Classification

Aspect Sentiment (or entity-based sentiment analysis).

Example

- *"iPhone is great"*
  - *iPhone* is entity, aspect is *GENERAL*
- *"iPhone's voice is great"*
  - *iPhone* is entity, aspect is *voice quality*

Aspect Sentiment Classification

Two tasks:

1. Aspect extraction
2. Aspect sentiment classification

Outline

Introduction
- Sentiment Analysis Application
- Sentiment Analysis Research

Document Sentiment Classification

Sentence Sentiment Classification
- Overview
- Subjectivity classification
- Sentence Sentiment Classification

Aspect Sentiment Classification
- Overview
- Aspect Extraction
- Aspect Sentiment Classification

Aspect Extraction

Example

The sound from this iPhone X phone is great
Girls from A is so beautiful

The entities are *iPhone X* and *A*, the aspects are *sound* and *girls*.

Aspect Extraction

Example

The sound from this iPhone X phone is great
Girls from A is so beautiful

The entities are *iPhone X* and *A*, the aspects are *sound* and *girls*.

Approaches

There are four main approaches to extracting explicit aspects:

1. Extraction by finding frequent nouns and noun phrases.
2. Extraction by exploiting syntactic relations:
   2.1 Syntactic dependencies depicting opinion and target relations.
   2.2 Lexico-syntactic patterns recoding entity and part/attribute relations.
3. Extracting using supervised learning.
4. Extracting using topic models.

Outline

Introduction
- Sentiment Analysis Application
- Sentiment Analysis Research

Document Sentiment Classification

Sentence Sentiment Classification
- Overview
- Subjectivity classification
- Sentence Sentiment Classification

Aspect Sentiment Classification
- Overview
- Aspect Extraction
- Aspect Sentiment Classification

Aspect Sentiment Classification (ASC)

ASC has two main approaches:

1. The supervised learning approach
2. The unsupervised lexicon-based approach

1. The supervised learning approach

- Jiang et al. (2011) [27] uses syntactic parse tree to generate a set of target-dependent features.
- Boiy and Moens (2009) [28] computed the feature weight for each word feature based on `distance(word, target_aspect)`

2. The unsupervised lexicon-based approach

TBD

Ref I

[1] Eldar Sadikov, Aditya G Parameswaran, Petros Venetis, et al. Blogs as predictors of movie success. In ICWSM, 2009.

[2] Taimoor Khan and Shehzad Khalid. Sentiment analysis for health care. 3:78–91, 06 2015.

[3] Ann Devitt and Khurshid Ahmad. Sentiment polarity identification in financial news: A cohesion-based approach. In Proceedings of the 45th annual meeting of the association of computational linguistics, pages 984–991, 2007.

[4] Brendan O'Connor, Ramnath Balasubramanyan, Bryan R Routledge, Noah A Smith, et al. From tweets to polls: Linking text sentiment to public opinion time series. Icwsm, 11(122-129):1–2, 2010.

[5] Andranik Tumasjan, Timm Oliver Sprenger, Philipp G Sandner, and Isabell M Welpe. Predicting elections with twitter: What 140 characters reveal about political sentiment. Icwsm, 10(1):178–185, 2010.

Ref II

[6] Bo Pang, Lillian Lee, et al. Opinion mining and sentiment analysis. Foundations and Trends® in Information Retrieval, 2(1–2):1–135, 2008.

[7] Bing Liu. Sentiment analysis and opinion mining. Synthesis lectures on human language technologies, 5(1):1–167, 2012.

[8] Thorsten Joachims. Transductive inference for text classification using support vector machines. In ICML, volume 99, pages 200–209, 1999.

[9] Nello Cristianini and John Shawe-Taylor. An introduction to support vector machines and other kernel-based learning methods. Cambridge university press, 2000.

[10] Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. Thumbs up?: sentiment classification using machine learning techniques. In Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10, pages 79–86. Association for Computational Linguistics, 2002.

Ref III

[11] Kushal Dave, Steve Lawrence, and David M Pennock. Mining the peanut gallery: Opinion extraction and semantic classification of product reviews. In Proceedings of the 12th international conference on World Wide Web, pages 519–528. ACM, 2003.

[12] Lukasz Augustyniak, Tomasz Kajdanowicz, Piotr Szymanski, Wlodzimierz Tuliglowicz, Przemyslaw Kazienko, Reda Alhajj, and Boleslaw K. Szymanski. Simpler is better? lexicon-based ensemble sentiment classification beats supervised methods. 2014 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2014), pages 924–929, 2014.

[13] Bo Pang and Lillian Lee. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In Proceedings of the 43rd annual meeting on association for computational linguistics, pages 115–124. Association for Computational Linguistics, 2005.

Ref IV

[14] Chong Long, Jie Zhang, and Xiaoyan Zhut. A review selection approach for accurate feature rating estimation. In Proceedings of the 23rd International Conference on Computational Linguistics: Posters, pages 766–774. Association for Computational Linguistics, 2010.

[15] Anthony Aue and Michael Gamon. Customizing sentiment classifiers to new domains: A case study. In Proceedings of recent advances in natural language processing (RANLP), volume 1, pages 2–1. Citeseer, 2005.

[16] John Blitzer, Mark Dredze, and Fernando Pereira. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In Proceedings of the 45th annual meeting of the association of computational linguistics, pages 440–447, 2007.

Ref V

[17] Songbo Tan, Gaowei Wu, Huifeng Tang, and Xueqi Cheng. A novel scheme for domain-transfer problem in the context of sentiment analysis. In Proceedings of the sixteenth ACM conference on Conference on information and knowledge management, pages 979–982. ACM, 2007.

[18] Xiaojun Wan. Using bilingual knowledge and ensemble techniques for unsupervised chinese sentiment analysis. In Proceedings of the conference on empirical methods in natural language processing, pages 553–561. Association for Computational Linguistics, 2008.

[19] Xiaojun Wan. Co-training for cross-lingual sentiment classification. In Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 1-volume 1, pages 235–243. Association for Computational Linguistics, 2009.

Ref VI

[20] Xiaojun Wan. Co-regression for cross-language review rating prediction. In Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), volume 2, pages 526–531, 2013.

[21] Jordan Boyd-Graber and Philip Resnik. Holistic sentiment analysis across languages: Multilingual supervised latent dirichlet allocation. In Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing, pages 45–55. Association for Computational Linguistics, 2010.

[22] Janyce M Wiebe, Rebecca F Bruce, and Thomas P O'Hara. Development and use of a gold-standard data set for subjectivity classifications. In Proceedings of the 37th annual meeting of the Association for Computational Linguistics on Computational Linguistics, pages 246–253. Association for Computational Linguistics, 1999.

Ref VII

[23] Hong Yu and Vasileios Hatzivassiloglou. Towards answering opinion questions: Separating facts from opinions and identifying the polarity of opinion sentences. In Proceedings of the 2003 conference on Empirical methods in natural language processing, pages 129–136. Association for Computational Linguistics, 2003.

[24] Bo Pang and Lillian Lee. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In Proceedings of the 42nd annual meeting on Association for Computational Linguistics, page 271. Association for Computational Linguistics, 2004.

[25] Ramanathan Narayanan, Bing Liu, and Alok Choudhary. Sentiment analysis of conditional sentences. In Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1-Volume 1, pages 180–189. Association for Computational Linguistics, 2009.

Ref VIII

[26] Dmitry Davidov, Oren Tsur, and Ari Rappoport. Semi-supervised recognition of sarcastic sentences in twitter and amazon. In Proceedings of the fourteenth conference on computational natural language learning, pages 107–116. Association for Computational Linguistics, 2010.

[27] Long Jiang, Mo Yu, Ming Zhou, Xiaohua Liu, and Tiejun Zhao. Target-dependent twitter sentiment classification. In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1, pages 151–160. Association for Computational Linguistics, 2011.

[28] Erik Boiy and Marie-Francine Moens. A machine learning approach to sentiment analysis in multilingual web texts. Information retrieval, 12(5):526–558, 2009.