

Sentiment Analysis of Social Media Texts

Saif M. Mohammad and Xiaodan Zhu

National Research Council Canada

1200 Montreal Road

Ottawa, K1A 0R6, ON, Canada

{Saif.Mohammad,Xiaodan.Zhu}@nrc-cnrc.gc.ca

1 Tutorial Description

Automatically detecting sentiment of product reviews, blogs, tweets, and SMS messages has attracted extensive interest from both the academia and industry. It has a number of applications, including: tracking sentiment towards products, movies, politicians, etc.; improving customer relation models; detecting happiness and well-being; and improving automatic dialogue systems. In this tutorial, we will describe how you can create a state-of-the-art sentiment analysis system, with a focus on social media posts.

We begin with an introduction to sentiment analysis and its various forms: term level, message level, document level, and aspect level. We will describe how sentiment analysis systems are evaluated, especially through recent SemEval shared tasks: Sentiment Analysis of Twitter (SemEval-2013 Task 2, SemEval 2014-Task 9) and Aspect Based Sentiment Analysis (SemEval-2014 Task 4).

We will give an overview of the best sentiment analysis systems at this point of time, including those that are conventional statistical systems as well as those using deep learning approaches. We will describe in detail the NRC-Canada systems, which were the overall best performing systems in all three SemEval competitions listed above. These are simple lexical- and sentiment-lexicon features based systems, which are relatively easy to re-implement.

We will discuss features that had the most impact (those derived from sentiment lexicons and negation handling). We will present how large tweet-specific sentiment lexicons can be automatically generated and evaluated. We will also show how negation impacts sentiment differently depending on whether the scope of the negation is positive or negative. All along, we will provide pointers

to how the relevant resources can be obtained. Finally, we will flesh out limitations of current approaches and promising future directions.

2 Tutorial Outline

- Introduction
 - Sentiment analysis
 - Document-level, sentence-level, term-level, aspect-level
 - Quirks of social media texts that are relevant to sentiment
 - Applications
- Recent SemEval sentiment tasks (task description, data, results)
 - Message-level (tweets): SemEval-2013 Task 2, SemEval-2014 Task 9
 - Term-level (within tweets): SemEval-2013 Task 2, SemEval-2014 Task 9
 - Aspect-level (in reviews): SemEval-2014 Task 4
- Features for sentiment analysis
 - Sentiment lexicons
 - Manually created lexicons
 - Automatically generated lexicons
 - Handling negation
 - Impact of negation on sentiment
 - Creating negation-specific sentiment lexicons
 - Ngrams, emoticons, hashtags, elongated words, punctuations
- Intrinsic evaluation of automatically generated sentiment lexicons
 - Creating a gold dataset using MaxDiff method of annotation
- Overview of sentiment analysis systems
 - Rule-based systems
 - Conventional statistical systems
 - Deep learning based systems
- Detailed description of a supervised statistical system (The NRC-Canada System)
 - Classifier, preprocessing, feature list, results
 - Message-level (tweets): SemEval-2013 Task 2, SemEval-2014 Task 9

- Term-level (within tweets): SemEval-2013 Task 2, SemEval-2014 Task 9
- Aspect-level (in reviews): SemEval-2014 Task 4
 - Detecting terms
 - Detecting term sentiment
 - Detecting categories
 - Detecting category sentiment
- Key ideas from other SemEval submissions
- Future directions

3 The Instructors

Saif M. Mohammad, Researcher, National Research Council Canada
saif.mohammad@nrc-cnrc.gc.ca
<http://www.saifmohammad.com>

Saif Mohammad is a Research Officer at the National Research Council Canada. His research interests are in Computational Linguistics, especially Lexical Semantics. He develops computational models for sentiment analysis, emotion detection, semantic distance, and lexical-semantic relations such as word-pair antonymy.

Xiaodan Zhu, Researcher, National Research Council Canada
xiaodan.zhu@nrc-cnrc.gc.ca
<https://sites.google.com/site/zhu2048/>

Xiaodan Zhu is a Research Officer at the National Research Council Canada. His research interests are in Natural Language Processing, Spoken Language Understanding, and Machine Learning. His recent work focuses on sentiment analysis, emotion detection, speech summarization, and deep learning.

The instructors along with Svetlana Kiritchenko developed the NRC-Canada Sentiment Analysis System, which was the top-performing system in recent SemEval shared-task competitions (SemEval-2013, Task 2, SemEval-2014 Task 9, and SemEval-2014 Task 4).

Acknowledgments

We thank Svetlana Kiritchenko for help with the tutorial development.

References

- Apoorv Agarwal, Boyi Xie, Ilia Vovsha, Owen Rambow, and Rebecca Passonneau. 2011. Sentiment Analysis of Twitter Data. In *Proceedings of the Workshop on Languages in Social Media*, Portland, Oregon.
- Yejin Choi and Claire Cardie. 2008. Learning with Compositional Semantics as Structural Inference for Subsentential Sentiment Analysis. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Edinburgh, UK.
- Alec Go, Richa Bhayani, and Lei Huang. 2009. Twitter sentiment classification using distant supervision. Technical report, Stanford University.
- Vasileios Hatzivassiloglou and Kathleen R. McKeown. 1997. Predicting the Semantic Orientation of Adjectives. In *Proceedings of the Conference of European Chapter of the Association for Computational Linguistics*, Madrid, Spain.
- Alistair Kennedy and Diana Inkpen. 2005. Sentiment Classification of Movie and Product Reviews Using Contextual Valence Shifters. In *Proceedings of the Workshop on the Analysis of Informal and Formal Information Exchange during Negotiations*, Ottawa, Ontario, Canada.
- Efthymios Kouloumpis, Theresa Wilson, and Johanna Moore. 2011. Twitter Sentiment Analysis: The Good the Bad and the OMG!. In *Proceedings of the 5th International AAAI Conference on Weblogs and Social Media*, Barcelona, Spain.
- Svetlana Kiritchenko, Xiaodan Zhu, and Saif M. Mohammad. 2014. State-of-the-Art in Sentiment Analysis of Short Informal Texts. (*To Appear*) *Journal of Artificial Intelligence Research (JAIR)*.
- Svetlana Kiritchenko, Xiaodan Zhu, Colin Cherry, and Saif M. Mohammad. 2014. NRC-Canada-2014: Detecting Aspects and Sentiment in Customer Reviews. In *Proceedings of the 2014 International Workshop on Semantic Evaluation*, Dublin, Ireland.
- Bing Liu and Lei Zhang. 2012. A Survey of Opinion Mining and Sentiment Analysis. In *Charu C. Aggarwal and ChengXiang Zhai, editors, Mining Text Data*, 415–463. Springer US.
- Eugenio Martínez-Cámara, M. Teresa Martín-Valdivia, L. Alfonso Ureña-López, and A Rtuero Montejo-Ráez. 2012. Sentiment analysis in Twitter. *Natural Language Engineering*, 20(1), 1-28.

- Saif M. Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu. 2013. NRC-Canada: Building the State-of-the-Art in Sentiment Analysis of Tweets. In *Proceedings of the 2013 International Workshop on Semantic Evaluation*, Atlanta, USA.
- Saif M. Mohammad. 2012. #Emotional Tweets. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics*, Montreal, Canada.
- Saif M. Mohammad and Peter D. Turney. 2010. Emotions Evoked by Common Words and Phrases: Using Mechanical Turk to Create an Emotion Lexicon. In *Proceedings of the NAACL-HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, LA, California, USA
- Bo Pang and Lillian Lee. 2008. Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval*, 2(1–2):1–135.
- Maria Pontiki, Dimitrios Galanis, John Pavlopoulos, Haris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. SemEval-2014 Task 4: Aspect Based Sentiment Analysis. In *Proceedings of the 2014 International Workshop on Semantic Evaluation*, Dublin, Ireland.
- Sara Rosenthal, Alan Ritter, and Preslav Nakov, and Veselin Stoyanov. 2014. SemEval-2014 Task 9: Sentiment Analysis in Twitter. In *Proceedings of the 2014 International Workshop on Semantic Evaluation*, Dublin, Ireland.
- Christina Sauper and Regina Barzilay. 2013. Automatic Aggregation by Joint Modeling of Aspects and Values. *Journal of Artificial Intelligence Research*, 46, 89-127.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Chris Manning, Andrew Ng, and Chris Potts. 2013. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. In *Proceedings of Conference on Empirical Methods in Natural Language Processing*, Seattle, USA.
- Peter D. Turney. 2002. Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews. In *Proceedings of Association For Computational Linguistics*, Philadelphia, PA.
- Duyu Tang, Furu Wei, Nan Yang, Ming Zhou, Bing Qin and Ting Liu. Learning Sentiment-Specific Word Embedding for Twitter Sentiment Classification. In *Proceedings of the 52th Annual Meeting of the Association for Computational Linguistics (ACL-2014)*, Baltimore, USA.
- Maite Taboada, Julian Brooke, Milan Tofiloski, Kimberly Voll, and Manfred Stede. 2011. Lexicon-based Methods for Sentiment Analysis. *Computational Linguistics*, 37 (2), 267-307.
- Michael Wiegand, Alexandra Balahur, Benjamin Roth, Dietrich Klakow, and Andres Montoyo. 2010. A survey on the Role of Negation in Sentiment Analysis. In *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*, Stroudsburg, PA, USA.
- Theresa Wilson, Zornitsa Kozareva, Preslav Nakov, Sara Rosenthal, Veselin Stoyanov, and Alan Ritter. 2013. SemEval-2013 task 2: Sentiment analysis in twitter. In *Proceedings of the International Workshop on Semantic Evaluation*, Atlanta, Georgia, USA.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing Contextual Polarity in Phrase-level Sentiment Analysis. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing*, Stroudsburg, PA, USA.
- Xiaodan Zhu, Svetlana Kiritchenko, and Saif M. Mohammad. 2014. NRC-Canada-2014: Recent Improvements in the Sentiment Analysis of Tweets. In *Proceedings of the 2014 International Workshop on Semantic Evaluation*, Dublin, Ireland.
- Xiaodan Zhu, Hongyu Guo, Svetlana Kiritchenko, and Saif M. Mohammad. An Empirical Study on the Effect of Negation Words on Sentiment. 2014. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Baltimore, USA.