

Aspect-Level Sentiment Analysis in Czech

Josef Steinberger

Department of Computer
Science and Engineering,
Faculty of Applied Sciences,
University of West Bohemia,
Univerzitní 8, 306 14 Plzeň
Czech Republic
jstein@kiv.zcu.cz

Tomáš Brychcín

NTIS – New Technologies
for the Information Society,
Faculty of Applied Sciences,
University of West Bohemia,
Univerzitní 8, 306 14 Plzeň
Czech Republic
brychcin@kiv.zcu.cz

Michal Konkol

NTIS – New Technologies
for the Information Society,
Faculty of Applied Sciences,
University of West Bohemia,
Univerzitní 8, 306 14 Plzeň
Czech Republic
konkol@kiv.zcu.cz

Abstract

This paper presents a pioneering research on aspect-level sentiment analysis in Czech. The main contribution of the paper is the newly created Czech aspect-level sentiment corpus, based on data from restaurant reviews. We annotated the corpus with two variants of aspect-level sentiment – aspect terms and aspect categories. The corpus consists of 1,244 sentences and 1,824 annotated aspects and is freely available to the research community. Furthermore, we propose a baseline system based on supervised machine learning. Our system detects the aspect terms with F-measure 68.65% and their polarities with accuracy 66.27%. The categories are recognized with F-measure 74.02% and their polarities with accuracy 66.61%.

1 Introduction

The interest in sentiment analysis (SA) is increasing with the amount of easily accessible content on the web, especially from the social media. Sentiment polarity is one of the critical information needed for many analysis of the data. Its use ranges from analysing product reviews (Stepanov and Riccardi, 2011) to predicting sales and stock markets using social media monitoring (Yu et al., 2013).

The majority of current approaches tries to detect the overall polarity of a sentence (or a document) regardless of the target entities (e.g., restaurants, laptops) and their aspects (e.g., food, price, battery, screen). By contrast, the aspect-driven sentiment analysis identifies the aspects of a given target entity and estimates the sentiment polarity for each mentioned aspect. This opens up completely new possibilities how to analyse the data.

The most of the research in automatic sentiment analysis has been devoted to English. There were

several attempts in Czech (Steinberger et al., 2011; Veselovská, 2012; Habernal et al., 2013; Brychcín and Habernal, 2013), but all were focused on the global (sentence- or document-level) sentiment. Although Czech is not a widely-spoken language on the global scale, it is in many ways similar to other Slavic languages and their speakers altogether represent an important group. The rich morphology and the free word order also makes it interesting from the linguistic perspective.

Our main goal is the creation of a aspect-level corpus as there is no such resource for Czech. We would like to support the beginning of aspect-level sentiment analysis for Czech and a human-annotated corpus is the first step in this direction. In addition, we want to provide results of a baseline system (based on machine learning techniques). This creates an easily reproducible starting point and allows anyone to quickly join the research of this task.

The rest of the paper is organised as follows. Section 2 is devoted to related work. It covers the aspect-level SA and sentiment analysis in Czech. Then we introduce the aspect-level architecture (Section 3) used for both the annotation of the corpus (Section 4) and for the automatic supervised approach (Section 5). In Section 6 we summarize our contribution and reveal our future plans.

2 Related work

The impact of SA can be seen in many practical applications, The users' opinions are mostly extracted either on a certain polarity scale, or binary (positive, negative). From the point of view of the granularity, the polarity has been assigned to a document or to a sentence. However, classifying opinions at the document level or the sentence level is often insufficient for applications because they do not identify opinion targets or assign sentiments to such targets (Liu, 2012). Even if we recognize the target entity (as the entity-centered

approaches do (e.g. Steinberger et al. (2011)), a positive opinion about the entity does not mean that the author has positive opinions about all aspects of the entity. Aspect-based sentiment analysis, which has been also called ‘feature-based’ (Hu and Liu, 2004), goes even deeper as it attempts to identify (and assign the polarity to) aspects of the target entity within a sentence (Hajmohammadi et al., 2012). Whenever we talk about an aspect, we must know which entity it belongs to. In the further discussion, we often omit the entity as we analysed restaurant reviews and thus our target entities are the reviewed restaurants.

2.1 Aspect-based sentiment analysis

The aspect scenario can be decomposed into two tasks: aspect extraction and aspect sentiment classification (Liu, 2012).

2.1.1 Aspect extraction

The task of aspect extraction, which can also be seen as an information extraction task, is to detect aspects that have been evaluated. For example, in the sentence, *The voice quality of this phone is amazing*, the aspect is *voice quality* of the entity represented by *this phone*.

The basic approach is finding frequent nouns and noun phrases. In (Liu et al., 2005), a specific method based on a sequential learning method was proposed to extract aspects from pros and cons, Blair-Goldensohn et al. (2008) refined the frequent noun and noun phrase approach by considering mainly those noun phrases that are in sentiment-bearing sentences or in some syntactic patterns which indicate sentiments. Moghaddam and Ester (2010) augmented the frequency-based approach with an additional pattern-based filter to remove some non-aspect terms. Long et al. (2010) extracted aspects (nouns) based on frequency and information distance.

Using supervised learning is another option. Aspect extraction can be seen as a special case of the general information extraction problem. The most dominant methods are based on sequential learning. Since these are supervised techniques, they need manually labeled data for training. One needs to manually annotate aspects and non-aspects in a corpus. The current state-of-the-art sequential learning methods are Hidden Markov Models (HMM) (Rabiner, 2010) and Conditional Random Fields (CRF) (Lafferty et al., 2001).

The last group of methods use topic models (Mei et al., 2007; Titov and McDonald, 2008; Blei et al., 2003). There are two main basic models, pLSA (Probabilistic Latent Semantic Analysis) (Hofmann, 1999) and LDA (Latent Dirichlet allocation) (Blei et al., 2003). In the SA context, one can design a joint model to model both sentiment words and topics at the same time, due to the observation that every opinion has a target.

2.1.2 Aspect sentiment classification

This task is to determine whether the opinions on different aspects are positive, negative, or neutral.

The classification approaches can be divided to supervised learning approaches and lexicon-based approaches. Supervised learning performs better in a particular application domain but it has difficulty to scale up to a large number of domains. Lexicon-based techniques often lose the fight against the learning but they are suitable for open-domain applications (Liu, 2012).

The key issue for learning methods is to determine the scope of each sentiment expression, i.e., whether it covers the aspect of interest in the sentence. In (Jiang et al., 2011), a dependency parser was used to generate a set of aspect dependent features for classification. A related approach was also used in (Boiy and Moens, 2009), which weights each feature based on the position of the feature relative to the target aspect in the parse tree.

Lexicon-based approaches use a list of sentiment phrases as the core resource. The method in (Ding et al., 2008) has four steps to assign a polarity to an aspect: mark sentiment words and phrases, apply sentiment shifters, handle but-clauses and aggregate opinions using an aggregation function (e.g. Hu and Liu (2004)).

2.2 Sentiment analysis for Czech

Pilot study of Czech sentiment analysis was shown in (Steinberger et al., 2012) where sentiment dictionaries for many languages (including Czech) were created using semi-automatic “triangulation” method.

Veselovská (2012) created a small corpus containing polarity categories for 410 news sentences and used the Naive Bayes and lexicon-based classifiers.

Three large labeled corpora (10k Facebook posts, 90k movie reviews, and 130k product reviews) were introduced in (Habernal et al.,

2013). Authors also evaluate three different classifiers, namely Naive Bayes, SVM (Support Vector Machines) and Maximum Entropy on these data.

Recently, Habernal and Bryhcín (2013) experimented with building word clusters, obtained from semantic spaces created on unlabeled data, as an additional source of information to tackle the high flexion issue in Czech.

These results were later outperformed by another unsupervised extension (Bryhcín and Habernal, 2013), where the global target context was shown to be very useful source of information.

3 The task definition

The aspect-level sentiment analysis firstly identifies the aspects of the target entity and then assigns a polarity to each aspect. There are several ways to define aspects and polarities. We use the definition based on the Semeval2014's Aspect-based SA task, which distinguishes two types of aspect-level sentiment – aspect terms and aspect categories.

The task is decomposed into the following 4 subtasks. We briefly describe each subtask and give some examples of source sentences and the expected results of the subtask.

3.1 Subtask 1: Aspect term extraction

Given a set of sentences with pre-identified entities (e.g., restaurants), the task is to identify the aspect terms present in the sentence and return a list containing all the distinct aspect terms. An aspect term names a particular aspect of the target entity.

Examples:

Děti dostaly naprosto krvavé maso.
(They brought a totally bloody meat to the kids.)
maso (meat)

Tlačenka se rozpadla, polévka ušla.
(The porkpie broke down, the soup was ok.)
Tlačenka (porkpie), polévka (soup)

3.2 Subtask 2: Aspect term polarity

For a given set of aspect terms within a sentence, the task is to determine the polarity of each aspect term: positive, negative, neutral or bipolar (i.e., both positive and negative).

Examples:

Děti dostaly naprosto krvavé maso.
(They brought a totally bloody meat to the kids.)

maso (meat): negative

Tlačenka se rozpadla, polévka ušla.
(The porkpie broke down, the soup was ok.)
Tlačenka (porkpie): negative, polévka (soup): positive

3.3 Subtask 3: Aspect category detection

Given a predefined set of aspect categories (e.g., price, food), the task is to identify the aspect categories discussed in a given sentence. Aspect categories are typically coarser than the aspect terms of Subtask 1, and they do not necessarily occur as terms in the given sentence.

For example, given the set of aspect categories food, service, price, ambience:

Přivítala nás velmi příjemná servírka, ale také místnost s ošuntělým nábytkem.
(We found a very nice waitress but also a room with time-worn furniture.)
service, ambience

Tlačenka se rozpadla, polévka ušla.
(The porkpie broke down, the soup was ok.)
food

3.4 Subtask 4: Aspect category polarity

Given a set of pre-identified aspect categories (e.g., food, price), the task is to determine the polarity (positive, negative, neutral or bipolar) of each aspect category.

Examples:

Přivítala nás velmi příjemná servírka, ale také místnost s ošuntělým nábytkem.
(We found a very nice waitress but also a room with time-worn furniture.)
service: positive, ambience: negative

Tlačenka se rozpadla, polévka ušla.
(The porkpie broke down, the soup was ok.)
food: bipolar

4 Building the aspect-level corpus

Aspect-level annotations are strictly connected to the analysed domain. As our final goal is going multilingual, we work on the domains selected for the Semeval2014's Aspect-based SA task (restaurants, laptop) which will allow us to compare approaches for both English and Czech on the same domains.

We started with the restaurants and in the future, we would also like to cover the laptops.

We downloaded restaurant reviews from www.nejezto.cz. Ten restaurants with the largest number of reviews were selected. The reviews were splitted into sentences. Average number of sentences per restaurant was 223.

4.1 Guidelines

The purpose of this annotation was to detect aspects and their sentiment polarity within sentences. The target entities were particular restaurants. For a given restaurant, the annotator had following tasks:

1. **Identify irrelevant sentences:** Sentences that do not contain any information relevant to the topic of restaurants. They were later filtered out of the corpus. Example: *Urážet někoho pro jeho názor je nedůstojné dospělého člověka. (Offencing somebody for his opinion is discreditable for an adult.)*
2. **Identify aspect terms:** Single or multiword terms naming particular aspects of the target entity. These are either full nominal phrases (*špíz a restované brambory – skewer with fried potatoes*) or verbs (*stojt – priced*). References, names or pronouns should not be annotated.
3. **Aspect term polarity:** Each aspect term has to be assigned one of the following polarities based on the sentiment that is expressed in the sentence about it: positive, negative, bipolar (both positive and negative sentiment) and neutral (neither positive nor negative sentiment).
4. **Aspect category:** The task of the annotator is to identify the aspect categories discussed in a sentence given the following five aspect categories: food, service, price, ambience, general (sentences mentioning the restaurant as a whole). Example: *Celkově doporučuji a vrátím se tam – Overall I would recommend it and go back again. general.*
5. **Aspect category polarity:** Each aspect category discussed by a particular sentence has to be assigned one of the following polarities based on the sentiment that is expressed in the sentence about it: positive, negative, bipolar, neutral.

4.2 Annotation statistics

Three native Czech speakers annotated in total 1,532 sentences. 18.8% of the sentences were marked as irrelevant, leaving 1,244 sentences for further analysis. Their average agreement for the task of aspect terms' identification was 82.6% (measured by F-measure). Only strict matches were considered correct. In the case of identifying the categories, their average agreement (F-measure) was 91.8%. The annotators agreed on 85.5% (accuracy) in the task of assigning polarity to terms and on 82.4% (accuracy) in the case of the category polarity assignment. It corresponds to Cohen's κ of 0.762, resp. 0.711, which represents a substantial agreement level (Pustejovsky and Stubbs, 2013), therefore the task can be considered as well-defined.

There were several reasons of disagreement. The annotators did not always identify the same terms, mainly in the cases with general meaning. In the case of polarity, the annotators did not agree on the most difficult cases to which bipolar class could be assigned:

Trochu přesolená omáčka, ale jinak luxus. (Too salted sauce, but luxury otherwise.)
food: bipolar vs. positive

The cases, on which the two annotators did not agree, were judged by the third super-annotator and golden standard data were created. The final dataset¹ contains 1244 sentences. The sentences contain 1824 annotated aspect terms (679 positive, 725 negative, 403 neutral, 17 bipolar) and 1365 categories (521 positive, 569 negative, 246 neutral, 28 bipolar).

5 Results of the supervised approach

5.1 Overview

We use machine learning approach in all subtasks. For aspect term extraction we use Conditional Random Fields (CRF). For the other three tasks we use the Maximum Entropy classifier. We use the Brainy² implementation of these algorithms.

During the data preprocessing, we use simple word tokenizer based on regular expressions. All tokens are lowercased for tasks 3 and 4. Due to the complex morphology of Czech we also use the un-

¹We will provide the dataset at <http://liks.fav.zcu.cz/sentiment>.

²Available at <http://home.zcu.cz/~konkol/brainy.php>

supervised stemmer called HPS³, that has already proved to be useful in sentiment analysis (Habernal et al., 2013; Habernal and Brychcín, 2013; Brychcín and Habernal, 2013).

All particular subtasks share following features:

Bag of words: The occurrence of a word.

Bag of bigrams: The occurrence of a bigram.

Bag of stems: The occurrence of a stem.

Bag of stem bigrams: The occurrence of a stem bigram.

5.2 Aspect term extraction

The system for aspect term extraction is based on CRF. The choice of CRF is based on a current state of the art in named entity recognition (see for example (Konkol and Konopík, 2013)) as it is a very similar task. We use the BIO (Ramshaw and Marcus, 1999) model to represent aspect terms. In addition to the previously mentioned features we use affixes and learned dictionaries. Affixes are simply prefixes and suffixes of length 2 to 4. Learned dictionaries are phrases that are aspect terms in the training data.

Our system achieved 58.14 precision, 83.80 recall and 68.65 F-measure.

5.3 Aspect term polarity

During the detection of the aspect term polarities, the words affecting the sentiment of the aspect term are assumed to be close in most of cases. Thus we use a small window (10 words in both directions) around the target aspect term. We assume the further the word or bigram is from the target aspect term, the lower impact it has on sentiment label. To model this assumption we use a weight for each word and bigram feature taken from the Gaussian distribution according to distance from aspect term. The mean is set to 0 and variance is optimized on training data. The classifier uses only the features presented in section 5.1. The results are presented in table 1.

5.4 Aspect category detection

Aspect category detection is based on the Maximum Entropy classifiers. We use one binary classifier for each category. Each classifier then decides whether the sentence has the given category

Table 1: Aspect term polarity results. p , r and f denote the precision, recall and F-measure. The results are expressed by percentages.

label	[%]	[%]	[%]
negative	76.41	63.31	69.25
neutral	33.75	50.18	40.36
positive	74.78	76.82	75.78
Accuracy: 66.27%			

or not. For this task we use only the bag of stems and Tf-Idf features.

Our system achieved 68.71 precision, 80.21 recall and 74.02 F-measure.

5.5 Aspect category polarity

For the category polarity detection we use the same features as for aspect term polarity detection. However in this case, we always take the whole sentence into account. We cannot take a limited window as we do not know where exactly the category is mentioned in the sentence. Moreover, it can be at several positions. To distinguish between different categories we use multiple Maximum Entropy classifiers, one for each category. The results are shown in table 2.

Table 2: Aspect category polarity results. p , r and f denote the precision, recall and F-measure. The results are expressed by percentages.

label	[%]	[%]	[%]
negative	74.07	66.04	69.83
neutral	37.80	46.73	41.80
positive	72.12	75.30	73.67
Accuracy: 66.61%			

5.6 Discussion

In section 5 we described our system for aspect-level sentiment analysis and showed the results. We do not use any language-dependent features, everything is learned from the training data. It is thus possible to say that our system is both language and domain independent, i.e. the system is able to work for any domain or language, if the training data are provided.

From another perspective, the already trained model is language and domain dependent (i.e. the model trained on restaurant domain probably will not perform well on laptop domain). The depen-

³Available at <http://liks.fav.zcu.cz/HPS>.

dence on the domain has multiple reasons. First, the categories are defined strictly for one domain (e.g. *food, price, etc.*). Second, many words can have different sentiment polarity in different domains.

In general, the sentiment analysis deals with many problems. These problems are much more evident for Czech as a representative of language with rich morphology and also with almost free word order. Here are two examples, where our system wrongly estimate the sentiment label.

Na nic si nejde stěžovat.
(*There is nothing to complain about.*)
general: positive

The sentence contains words that frequently occur in negative reviews: *nic* - *nothing*, *stěžovat* - *complain*; but the sentence is positive.

O těch labužnických a delikatesních zážitcích si člověk pouze přečte, ale realita je jiná.
(*One can only read about these gourmand and delicious experiences, but the reality is completely different.*)
food: negative

Sentence contains words like *labužnických* - *gourmand* and *delikatesních* - *delicious* that are strictly positive, but in this context it is mentioned negatively.

As we already said, this is the pilot study of aspect-level sentiment analysis in Czech. Several studies about sentence-level sentiment analysis of Czech have been already published, and thus it is worth comparing how these two tasks differ in terms of difficulty. Note that the aspect-level sentiment analysis has to deal with multiple aspects and categories in a given sentence, and thus it is apparently a much more difficult task.

We believe the results of (Brychcín and Habernal, 2013) on Czech movie reviews dataset can be a comparable example of sentence-level sentiment analysis as they also distinguish 3 sentiment labels (positive, negative and neutral) and the data are taken from a closed domain (movies). Their best result (given by the model with all extensions) is 81.53%. Our best results are 66.27% and 66.61% for aspect and category polarity detection, respectively.

6 Conclusion

The aspect level sentiment analysis has not been studied for Czech yet. The main reason for this is

the lack of annotated data. In this paper, we create a high quality gold data for this task, we describe our approach to their annotation and discuss their properties. Corpus is available for free at <http://likes.fav.zcu.cz/sentiment>.

We also propose a baseline model based on state-of-the-art supervised machine learning techniques. Our system is language and domain independent, i.e. it can be easily trained on data from another domain or language. It achieved 68.65% F-measure in the aspect term detection, 74.02% F-measure in the aspect category assigning, 66.27% accuracy in the aspect term polarity classification, and 66.61% accuracy in the aspect category polarity classification.

In the future, we would like to continue the aspect-level research direction in three ways. We would like to extend the currently created restaurant reviews' corpus, to add the second (laptop's) domain to the corpus, and finally, to experiment with extensions to the baseline system. As the corpus for the Semeval2014 aspect-based SA task contains review sentences from the same domains, we will be able to compare the results of the system cross-lingually.

Acknowledgments

This work was supported by grant no. SGS-2013-029 Advanced computing and information systems, by the European Regional Development Fund (ERDF), by project "NTIS - New Technologies for Information Society", European Centre of Excellence, CZ.1.05/1.1.00/02.0090, and by project MediaGist, EU's FP7 People Programme (Marie Curie Actions), no 630786.

References

- Sasha Blair-Goldensohn, Kerry Hannan, Ryan McDonald, Tyler Neylon, George Reis, and Jeff Reynar. 2008. Building a sentiment summarizer for local service reviews. In *Proceedings of WWW-2008 workshop on NLP in the Information Explosion Era*.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022.
- Erik Boiy and Marie-Francine Moens. 2009. A machine learning approach to sentiment analysis in multilingual web texts. *Information retrieval*, 12(5):526–558.
- Tomáš Brychcín and Ivan Habernal. 2013. Unsupervised improving of sentiment analysis using global

- target context. In *Proceedings of the International Conference Recent Advances in Natural Language Processing RANLP 2013*, pages 122–128, Hissar, Bulgaria, September. Incoma Ltd. Shoumen, Bulgaria.
- Xiaowen Ding, Bing Liu, and Philip S. Yu. 2008. A holistic lexicon-based approach to opinion mining. In *Proceedings of the Conference on Web Search and Web Data Mining*.
- Ivan Habernal and Tomáš Brychcín. 2013. Semantic spaces for sentiment analysis. In *Text, Speech and Dialogue*, volume 8082 of *Lecture Notes in Computer Science*, pages 482–489, Berlin Heidelberg. Springer.
- Ivan Habernal, Tomáš Ptáček, and Josef Steinberger. 2013. Sentiment analysis in czech social media using supervised machine learning. In *Proceedings of the 4th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 65–74, Atlanta, Georgia, June. Association for Computational Linguistics.
- Mohammad Sadegh Hajmohammadi, Roliana Ibrahim, and Zulaiha Ali Othman. 2012. Opinion mining and sentiment analysis: A survey. *International Journal of Computers & Technology*, 2(3).
- Thomas Hofmann. 1999. Probabilistic latent semantic indexing. In *Proceedings of Conference on Uncertainty in Artificial Intelligence*.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining, KDD '04*, pages 168–177, New York, NY, USA. ACM.
- Long Jiang, Mo Yu, Ming Zhou, Xiaohua Liu, and Tiejun Zhao. 2011. Target-dependent twitter sentiment classification. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*.
- Michal Konkol and Miloslav Konopík. 2013. Crf-based czech named entity recognizer and consolidation of czech ner research. In Ivan Habernal and Václav Matoušek, editors, *Text, Speech and Dialogue*, volume 8082 of *Lecture Notes in Computer Science*, pages 153–160. Springer Berlin Heidelberg.
- John Lafferty, Andrew McCallum, and Fernando Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of International Conference on Machine Learning*.
- Bing Liu, Minqing Hu, and Junsheng Cheng. 2005. Opinion observer: Analyzing and comparing opinions on the web. In *Proceedings of International Conference on World Wide Web*.
- Bing Liu. 2012. *Sentiment Analysis and Opinion Mining*. Morgan & Claypool Publishers.
- Chong Long, Jie Zhang, and Xiaoyan Zhu. 2010. A review selection approach for accurate feature rating estimation. In *Proceedings of Coling 2010: Poster Volume*.
- Qiaozhu Mei, Xu Ling, Matthew Wondra, Hang Su, and ChengXiang Zhai. 2007. Topic sentiment mixture: modeling facets and opinions in weblogs. In *Proceedings of International Conference on World Wide Web*.
- Samaneh Moghaddam and Martin Ester. 2010. Opinion digger: an unsupervised opinion miner from unstructured product reviews. In *Proceeding of the ACM conference on Information and knowledge management*.
- James Pustejovsky and Amber Stubbs. 2013. *Natural Language Annotation for Machine Learning*. OR-eilly Media, Sebastopol, CA 95472.
- Lawrence Rabiner. 2010. A tutorial on hidden markov models and selected applications in speech recognition. In *Proceedings of the IEEE*, pages 257–286.
- Lance A Ramshaw and Mitchell P Marcus. 1999. Text chunking using transformation-based learning. In *Natural language processing using very large corpora*, pages 157–176. Springer.
- J. Steinberger, P. Lenkova, M. Kabadjov, R. Steinberger, and E. van der Goot. 2011. Multilingual entity-centered sentiment analysis evaluated by parallel corpora. In *Proceedings of the 8th International Conference Recent Advances in Natural Language Processing, RANLP'11*, pages 770–775.
- J. Steinberger, M. Ebrahim, Ehrmann M., A. Hurriyetoglu, M. Kabadjov, P. Lenkova, R. Steinberger, H. Tanev, S. Vzquez, and V. Zavarella. 2012. Creating sentiment dictionaries via triangulation. *Decision Support Systems*, 53:689–694.
- E.A. Stepanov and G. Riccardi. 2011. Detecting general opinions from customer surveys. In *Data Mining Workshops (ICDMW), 2011 IEEE 11th International Conference on*, pages 115–122.
- Ivan Titov and Ryan McDonald. 2008. Modeling online reviews with multi-grain topic models. In *Proceedings of International Conference on World Wide Web*.
- Kateřina Veselovská. 2012. Sentence-level sentiment analysis in czech. In *Proceedings of the 2nd International Conference on Web Intelligence, Mining and Semantics*. ACM.
- Liang-Chih Yu, Jheng-Long Wu, Pei-Chann Chang, and Hsuan-Shou Chu. 2013. Using a contextual entropy model to expand emotion words and their intensity for the sentiment classification of stock market news. *Knowledge Based Syst*, 41:89–97, March.