Exploiting the Focus of the Document for Enhanced Entities' Sentiment Relevance Detection

Zvi Ben-Ami, Ronen Feldman and Benjamin Rosenfeld
The Hebrew University

Extended Abstract

Sentiment Analysis (SA) is driven by the desire to understand people's explicit thoughts and opinions on particular entities or situations, in a given text. A key question in SA is whether existing sentiment expressions are related to particular entities, which appear in the same text. This is an imperative question, since people are typically interested in sentiments on specific entities and not in the overall sentiment articulated in a document or an article. Effective methods are, therefore, needed for adequately relating sentiment expressions with underlining entities, or in other words, methods to determine the sentiment relevance for each entity of interest.

In this work, we argue that a better binding of sentiment expressions with entities can be achieved by exploiting context information about the entity of interest, such as, the type of the entity, whether it is the only entity of its type in the document or one of many, the position of the entity within the document, section or paragraph, the lexical contents, etc. We establish that one of the most valuable context clues is the "*focus of the document on the entity of interest*", i.e., whether the entity is the main topic of the document being analyzed, one of several main topics, or is just mentioned in passing. Hence, a document would likely have different foci with respect to the different entities mentioned within it. In the datasets we use for the experiments, the document type is manually annotated with respect to each entity, this allows us to directly observe the influence of this information on detecting sentiment relevance. We, further, show that this information can be automatically extracted using a supervised classification learning method, by using part of the annotated data as training data for the classifier.

In order to assess the value of such information, we look at various methods for detecting sentiment relevance for entities. We consider both rule-based algorithms that rely on the entity's physical or syntactic proximity to the sentiment expressions as well as more sophisticated machine learning classification algorithms. We demonstrate that the focus of the document on the entities within it is, indeed, an important piece of information, which can be accurately learned with supervised classification means.

Understanding the nature of different entity types and the way they interact with entities of other types are additional properties that should be considered when determining the sentiment relevance for entities. In this work, we distinguish between cases where the interest is in one entity type, such as COMPANY or PRODUCT, and cases where the interest is in the interaction between entities of different types, for example the sentiment on a certain drug when treating a certain disease. This problem is close to Relation Extraction in a sense. In particular, we are interested in examples from the medical domain regarding three entity types: PERSON, DRUG, and DISEASE, where PERSON is restricted to known physicians. While each of the entity types can be the target of a sentiment expression, the more interesting questions in this domain involve multiple entities, specifically, the intersection of DRUG and DISEASE, in order to address questions such as "how effective is this drug for this disease?", and PERSON + DRUG + DISEASE, in order to address questions such as "what does this physician say about using this drug to cure this disease?" We solve the multiple-entity relevance problem by intersecting the relevance ranges of different-type entities, thus reducing the problem to the single-entity relevance detection. As such, the experiments regarding the multiple-entity relevance need only check the accuracy of this reduction.

Existing sentiment analysis systems are typically differentiated on two general dimensions. The first is the level of granularity at which the analysis is being conducted and the second is the level of automatization, i.e., the extent to which ML versus rule-based techniques are applied. In this work, we primarily study the effect of information on the focus of the document in determining sentiment relevance and, therefore, look at a wide spectrum of methods for SA, while providing or withholding this particular information. At the lower-end of complexity, we execute a document-level analysis. This simple model serves as a baseline in evaluating the overall necessity of detecting

sentiment relevance. At the next level, we execute sentence-level analysis in which a sentiment expression is deemed relevant only if the entity of interest is explicitly mentioned (including co-reference) in the same sentence. At a finer granularity level, we look at algorithms of two types, rule-based algorithms that rely on the entity's proximity (physical and syntactic) to the sentiment expressions and ML algorithms (specifically, linear classification and sequence classification), which use a diversified set of features representation pertaining to the context in which the entity is mentioned. Using rule-based algorithms allow us to evaluate whether designing rules that are suited to the different possible foci of documents, can improve the performance of the detection task, which could, in turn, validate the benefit of exploiting this context information. The ML algorithms allow us to evaluate how well this task can be performed with supervised learning means.

In addition to evaluating various algorithms for detecting sentiment relevance with manually annotated data on the focus of the document, we evaluate these algorithms with annotated data that was obtained with supervised classification means. In other words, we first learned the focus of the document for each entity using supervised means, and then evaluated the performances sentiment relevance algorithms using this automatically obtained data. Although many algorithms can be thought to be applicable to this task, comparing between the performances of different possible algorithms for identifying the focus of the document falls out of the scope of this study. The purpose of this particular exercise is to determine the feasibility of automatically identifying the focus of the document and not to find the most optimal methods.

In this study, we pay particular attention to two domains of interest: Financial and Medical, but the results can be expanded to other domains. The main contributions of this study are on three levels, first, we show that sentiment relevance detection is important for the general task of sentiment analysis. Second, we show that the focus of the document is valuable information that should be considered when determining sentiment relevance, and that it can be learned with supervised ML means. Third, we compare between the performances of various methods for detecting sentiment relevance. Overall, we found that classification-based algorithms perform better than the deterministic ones in identifying sentiment relevance, with sequence-classification performing significantly better than direct classification.

Selected References

[1] Engonopoulos, N. et al. (2011).ELS: a word-level method for entity-level sentiment analysis. *Proceedings of the International Conference on Web Intelligence, Mining and Semantics* (2011), 12.

[2] Etzioni, O. et al. (2008).Open information extraction from the web. *Communications of the ACM*. 51, 12 (2008), 68–74.

[3] Feldman, R. (2013).Techniques and applications for sentiment analysis. *Communications of the ACM*. 56, 4 (Apr. 2013), 82.

[4] Feldman, R. et al. (2010).The Stock Sonar — Sentiment Analysis of Stocks Based on a Hybrid Approach. *Proceedings of the Twenty-Third Innovative Applications of Artificial Intelligence Conference* (2010), 1642–1647.

[5] Liu, B. (2012).*Sentiment Analysis and Opinion Mining Synthesis Lectures on Human Language Technologies*. Morgan & Claypool Publishers.

[6] Loughran, T.I.M. and Mcdonald, B. (2010).When is a Liability not a Liability ? Textual Analysis , Dictionaries , and 10-Ks Journal of Finance , forthcoming. *Journal of Finance*. 66, (2010), 35–65.

[7] Pang, B. et al. (2002).Thumbs up ? Sentiment Classification using Machine Learning Techniques. *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP), Philadelphia, July 2002, pp. 79-86. Association for Computational Linguistics.* (2002), 79–86.

[8] Pontiki, M. et al. (2014).Semeval-2014 task 4: Aspect based sentiment analysis. *Proceedings of SemEval, Dublin, Ireland*. (2014).

[9] Popescu, A.-M. and Etzioni, O. (2005).Extracting product features and opinions from reviews. *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing HLT 05* (2005), 339–346.

[10] Scheible, C. and Schütze, H. (2013).Sentiment Relevance. *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (Sofia, Bulgaria, Aug. 2013), 954–963.

[11] Tiedemann, J. and Mur, J. (2008).Simple is best: experiments with different document segmentation strategies for passage retrieval. *Coling 2008: Proceedings of the 2nd workshop on Information Retrieval for Question Answering* (Stroudsburg, PA, USA, 2008), 17–25.

[12] Turney, P. (2002).Thumbs Up or Thumbs Down ? Semantic Orientation Applied to Unsupervised Classification of Reviews. *Proceedings of the Association for Computational Linguistics (ACL)* (2002), 417–424.

[13] Zhang, T. (2001).Regularized winnow methods. Advances in Neural Information Processing Systems. (2001), 703–712.