An Embedding Approach For Microblog Polarity Classification

Andi Rexha, Mark Kröll, Roman Kern, and Mauro Dragoni

Know-Center GmbH Graz, Austria
{arexha, mkroell, rkern}@know-center.at
dragoni@fbk.eu

Abstract. In the last years, forms of communications such as social media have emerged. Short and unstructured messages are used to share, interact and collaborate in different online communities. Identifying the nature of the emotion (positive or negative) expressed in these kind of text is a big challenge for the standard Natural Language Processing (NLP). Tweets are one of the most popular type of short messages. In this paper we try to predict the polarity (positive or negative emotion) expressed by the user for a specific target phrase in a tweet. We try to exploit a Tweeter Word2Vec[1] model in order to classify the polarity of these message refereed to the target phrase. We use to approaches to extract the features: windows based and whole message based. Evaluating with the *SemEval 2016 Task 4* dataset, we show that these simple approaches perform quite well, even though they do not use any polarity of single words. We also show that the performance of considering the whole tweet message is slightly better than the one considering a window around the target phrase.

1 Introduction

With the growing popularity of the online social media, different forms of communication are being used more and more often. The trend of messaging is shifting to microblogging and short texts which usually are unstructured and very informal. While users of these social media aren't limited to specific type of text, they usually express their opinions or emotions about specific interests.

One of the most popular social media providing sharing of short texts is Twitter and it's messages are called Tweets. The language used in these messages is very informal, with creative spelling and punctuation, misspellings, slang, URLs and abbreviations. The difficulty in processing this kind of text is challanging for researchers.

Efforts have been made in tasks for automatically predicting sentiment polarity (whether positive or negative) of tweets. Even more challenging is to predict the opinion of specific target. In order to illustrate this, consider the following example:

New features ***@Microsoft*** *suck. Check them back!* ***#Linux*** *solutions are awesome*

[1] Word2Vec models provide a representation of words in a feature space that reflects their relation to other words in the training corpus

There might be a neutral overall opinion, being that the first part "New features **@Microsoft** suck" expresses a negative emotion meanwhile the last part of the message "#Linux solutions are awesome." expresses a positive one. The two different references of these opinions (in this case *@Microsoft* and *#Linux*) are called the target phrases. In this paper we try to address exactly this challenge. Specifically we try to automatically predict the polarity (whether positive or negative) in a message about a given target. For instance, in the previews example, the algorithm should return a positive aspect about the target *@Microsoft* and a negative one about the target *#Linux*.

In order to tackle the challenge, in this paper we explore the semantic information given by a Word2Vec[21] trained model on twitter messages. Word2Vec models provide a representation of words in a feature space that reflects their relation to other words in the training corpus. We investigate whether the role of the target by using this technique has different outcomes compared to the only use of the words close to the target phrase. To evaluate this approach we use the test and golden standard dataset of the *Semeval 2016 Task #4* challenge about Twitter sentiment mining.

The paper is structure as follow. In the section 2 we present the approach of the challengers of the SemEval task of the 2015 and other related works. In the section 3 we suggest the different way to extract the features and in section 4 we compare the different approaches and learning algorithms. In the 5 section we draw the conclusion and the future work.

2 Related Work

The task of Sentiment Analysis, also known as opinion mining (cf. [22], [20]), is to classify textual content according to expressed emotions and opinions. Sentiment classification has been a challenging topic in Natural Language Processing (cf. [32]). It is commonly defined as a binary classification task to assign a sentence either positive or negative polarity (cf. [23]). Turney's work was among the first ones to tackle automatic sentiment classification ([31]). He employed an information-theoretic measure, i.e. mutual information, between a text phrase and the words "excellent" and "poor" as a decision metric.

The approaches presented above are applied at the document-level[8, 24, 26, 13], i.e., the polarity value is assigned to the entire document content. However, in some case, for improving the accuracy of the sentiment classification, a more fine-grained analysis of a document is needed. Hence, the sentiment classification of the single sentences, has to be performed. In the literature, we may find approaches ranging from the use of fuzzy logic [12, 11, 25] to the use of aggregation techniques [7] for computing the score aggregation of opinion words. In the case of sentence-level sentiment classification, two different sub-tasks have to be addressed: (i) to determine if the sentence is subjective or objective, and (ii) in the case that the sentence is subjective, to determine if the opinion expressed in the sentence is positive, negative, or neutral. The task of classifying a sentence as subjective or objective, called "subjectivity classification", has been widely discussed in the literature [14, 28, 34] and systems implementing the capabilities of identifying opinion's holder, target, and polarity have been presented [4]. Once subjective sentences are identified, the same methods as for sentiment classification may be applied. For example, in [17] the authors consider gradable adjectives for sentiment spotting; while in [18, 27] the authors built models to identify some specific types of opinions.

A particular attention should be given also to the application of sentiment analysis in social networks [9, 10]. Micro-blogging data such as tweets differs from regular text as it is extremely noisy, informal and does not allow for long messages (which might not be a disadvantage (cf. [6]). As a consequence, analyzing sentiment in Twitter data poses a lot of opportunities. Traditional feature representations such as part-of-speech information or the usage of lexicon features such as SentiWordNet have to be re-evaluated in the light of Twitter data. In case of part-of-speech information, Gimpel et al. ([15]) annotated tweets and developed a tagset and features to train an adequate tagger. Kouloumpis et al. ([19]) investigated the usefulness of existing lexical resources and other features including part-of-speech information in the analysis task.

Go et al. ([16]), for instance, used emoticons as additional features, for example, ":)" and ":-)" for the positive class, ":(" and ":-(" for the negative class. They then applied machine learning techniques such as support vector machines to classify the tweets into a positive and a negative class. Agarwal et al. ([3]) introduced POS-specific prior polarity features along with using a tree kernel for tweet classification. Barbosa and Feng ([5]) present a robust approach to Twitter sentiment analysis. The robustness is based on an abstract representation of tweets as well as the usage of noisy/biased labels from three websites to train their model.

Last but not least, recent years have seen a lot of participation in the annual SemEval tasks on Twitter Sentiment Analysis (cf. [33], [30], [29]). This event provides optimal conditions to implement novel ideas and is a good starting point to catch up on the latest trends in this area.

3 Approach

As exlained in the previews sections we intend to experiment two different approaches for extracting mining features. In the first approach we use the sole information of each word without considering the position of the target phrase. On the other hand, in the second approach we consider only the surrounding of the target phrase. As a preprocessing step, we annotate the tweets (words, Part Of Speach Tagging etc.) by using the Tweet NLP library [1]. Further, for each word of the tweet we extract the Word2Vec vector representation by using a Twitter model trained over 400 million tweets [2]. In the postprocessing step we make an average over the considered segment (every word or words within the window). As the last step we use a binary feature which is set to 1 if in the tweet exists any negation word (don't, not, . . .). We believe that this feature can give a hint to the learning algorithm whether the expressed emotion might be negated without considering those kind of words. Below we explain more in detail each of the approaches for extracting the features.

Whole Tweet Run The whole tweet run can be expained in the current steps:

- Preprocess the tweet messages, extracting each word
- For each of the words of the tweet, extract the Word2Vec value

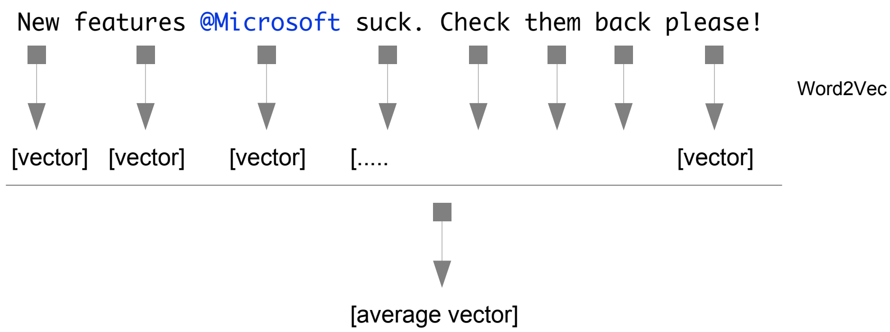

Fig. 1. Whole Tweet Approach: showing how we extract the features without even considering the target phrase

– For each corresponding feature extracted from the Word2Vec, make an average
– a binary feature as negation: if one of the words in the tweet contains a "not" or ends with a "'t", the feature is set to 1, otherwise, to 0

In the figure 1 there is shown the way the Word2Vec features are extracted for the *Whole Tweet approach*

Window Run The window run can be expained in the current steps:

– Preprocess the tweet messages, extracting each word
– Annotate the target of the tweet
– Build a winodow of "n" words from left and right of the target
– Extract the Word2Vec value for each word of the window
– For each corresponding feature extracted from the Word2Vec, make an average
– a binary feature as negation: if one of the words in the tweet contains a "not" or ends with a "'t", the feature is set to 1, otherwise, to 0

We show the extraction of the Word2Vec features for the *Window approach* in the figure 2.

4 Results

To evaluate the two different approaches that we proposed in the section 3 we have trained different classifiers to predict the opinion of the target phrase in a Tweet. We have used the dataset of the *Semeval 2016 Task 4* about sentiment analysis in Tweets. The training set is composed by 3858 entries and the evaluation set by 10551 entries. Both datasets are skewed. The training set contains 17% of negative and 83% of positive and the evaluation set of 22% of negative and 78% of positive examples. For each of the approaches we present the evaluation for the positive and negative classes by desplaying the precision, recall and F1 measure. In the tables 1 and 2 we show the performance of

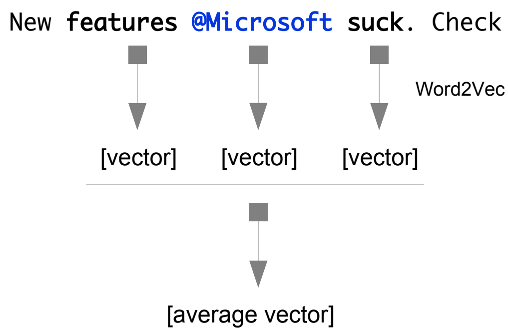

Fig. 2. Window approach: Extracting the features from a window of words close to the target phrase (in the example the size of the window is 1)

the positive and negative classes for the full text approach. On the other hand, in the tables 3 and 4 we present the evaluation for the window approach. The chosen size of the window is set to 3. We believe that this size reflect the idea of chosing related words close to the target phrase.

	Precision	Recall	F1-Measure
Naive Bayes	0.396	**0.733**	0.514
Support Vector Machine	**0.724**	0.347	0.469
Logistic Regression	0.606	0.481	**0.536**
Random Tree	0.321	0.266	0.291

Table 1. Evaluation for the Negative Opinions by using an average of the Word2Vec over all the words composing the tweet

Something to highlight from the tables is that the accuracy of the negative class is lower. We believe that this is due to the skewed nature of the dataset. Another detail to note is the difference between the two approaches. This characteristic might be due to the fact that we throw away some important information that are not in the proximity to the target phrase.

5 Conclusion

In this paper we have evaluate an approach by using the semantic information collected from the Word2Vec for the prediction of the polarity in tweets. Specifically we have addressed the opinion mining for target phrases.

In future work we intend to exploit the effect of the dependency trees in tweets. The text proximity can give just partial information about the semantic proximity of

	Precision	Recall	F1-Measure
Naive Bayes	**0.900**	0.681	0.775
Support Vector Machine	0.838	**0.962**	**0.896**
Logistic Regression	0.860	0.911	0.885
Random Tree	0.801	0.840	0.820

Table 2. Evaluation for the Positive Opinions by using an average of the Word2Vec over all the words composing the tweet

	Precision	Recall	F1-Measure
Naive Bayes	0.303	**0.732**	**0.429**
Support Vector Machine	**0.442**	0.273	0.338
Logistic Regression	0.396	0.391	0.394
Random Tree	0.260	0.281	0.271

Table 3. Evaluation for the Negative Opinions by using a window of size 3

the positive or negative words in the short messages. We believe that exploiting this information can improve the performance of the classifying algoritm.

Acknowledgment

This work is funded by the KIRAS program of the Austrian Research Promotion Agency (FFG) (project number 840824). The Know-Center is funded within the Austrian COMET Program under the auspices of the Austrian Ministry of Transport, Innovation and Technology, the Austrian Ministry of Economics and Labour and by the State of Styria. COMET is managed by the Austrian Research Promotion Agency FFG

	Precision	Recall	F1-Measure
Naive Bayes	**0.872**	0.520	0.652
Support Vector Machine	0.813	**0.902**	**0.855**
Logistic Regression	0.827	0.830	0.829
Random Tree	0.791	0.773	0.781

Table 4. Evaluation for the Positive Opinions by using a window of size 3

References

1. Nlp for tweets. `http://www.cs.cmu.edu/~ark/TweetNLP/`
2. Word2vec for tweets. `http://www.fredericgodin.com/software/`
3. Agarwal, A., Biadsy, F., Mckeown, K.R.: Contextual phrase-level polarity analysis using lexical affect scoring and syntactic n-grams. In: Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics. EACL '09 (2009)
4. Aprosio, A.P., Corcoglioniti, F., Dragoni, M., Rospocher, M.: Supervised opinion frames detection with RAID. In: Gandon, F., Cabrio, E., Stankovic, M., Zimmermann, A. (eds.) Semantic Web Evaluation Challenges - Second SemWebEval Challenge at ESWC 2015, Portorož, Slovenia, May 31 - June 4, 2015, Revised Selected Papers. Communications in Computer and Information Science, vol. 548, pp. 251–263. Springer (2015), `http://dx.doi.org/10.1007/978-3-319-25518-7_22`
5. Barbosa, L., Feng, J.: Robust sentiment detection on twitter from biased and noisy data. In: Proceedings of the 23rd International Conference on Computational Linguistics: Posters. COLING '10 (2010)
6. Bermingham, A., Smeaton, A.F.: Classifying sentiment in microblogs: Is brevity an advantage? In: Proceedings of the 19th ACM International Conference on Information and Knowledge Management. CIKM '10, ACM (2010)
7. da Costa Pereira, C., Dragoni, M., Pasi, G.: A prioritized "and" aggregation operator for multidimensional relevance assessment. In: Serra, R., Cucchiara, R. (eds.) AI*IA 2009: Emergent Perspectives in Artificial Intelligence, XIth International Conference of the Italian Association for Artificial Intelligence, Reggio Emilia, Italy, December 9-12, 2009, Proceedings. Lecture Notes in Computer Science, vol. 5883, pp. 72–81. Springer (2009), `http://dx.doi.org/10.1007/978-3-642-10291-2_8`
8. Dragoni, M.: Shellfbk: An information retrieval-based system for multi-domain sentiment analysis. In: Proceedings of the 9th International Workshop on Semantic Evaluation. pp. 502–509. SemEval '2015, Association for Computational Linguistics, Denver, Colorado (June 2015)
9. Dragoni, M.: A three-phase approach for exploiting opinion mining in computational advertising. IEEE Intelligent Systems 32(3), 21–27 (2017), `https://doi.org/10.1109/MIS.2017.46`
10. Dragoni, M., da Costa Pereira, C., Tettamanzi, A.G.B., Villata, S.: Smack: An argumentation framework for opinion mining. In: Kambhampati, S. (ed.) Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016. pp. 4242–4243. IJCAI/AAAI Press (2016), `http://www.ijcai.org/Abstract/16/641`
11. Dragoni, M., Tettamanzi, A.G.B., da Costa Pereira, C.: A fuzzy system for concept-level sentiment analysis. In: Presutti, V., Stankovic, M., Cambria, E., Cantador, I., Iorio, A.D., Noia, T.D., Lange, C., Recupero, D.R., Tordai, A. (eds.) Semantic Web Evaluation Challenge - SemWebEval 2014 at ESWC 2014, Anissaras, Crete, Greece, May 25-29, 2014, Revised

Selected Papers. Communications in Computer and Information Science, vol. 475, pp. 21–27. Springer (2014), http://dx.doi.org/10.1007/978-3-319-12024-9_2
12. Dragoni, M., Tettamanzi, A.G., da Costa Pereira, C.: Propagating and aggregating fuzzy polarities for concept-level sentiment analysis. Cognitive Computation 7(2), 186–197 (2015), http://dx.doi.org/10.1007/s12559-014-9308-6
13. Federici, M., Dragoni, M.: A knowledge-based approach for aspect-based opinion mining. In: Sack, H., Dietze, S., Tordai, A., Lange, C. (eds.) Semantic Web Challenges - Third SemWebEval Challenge at ESWC 2016, Heraklion, Crete, Greece, May 29 - June 2, 2016, Revised Selected Papers. Communications in Computer and Information Science, vol. 641, pp. 141–152. Springer (2016), https://doi.org/10.1007/978-3-319-46565-4_11
14. Federici, M., Dragoni, M.: Towards unsupervised approaches for aspects extraction. In: Dragoni, M., Recupero, D.R., Denecke, K., Deng, Y., Declerck, T. (eds.) Joint Proceedings of the 2th Workshop on Emotions, Modality, Sentiment Analysis and the Semantic Web and the 1st International Workshop on Extraction and Processing of Rich Semantics from Medical Texts co-located with ESWC 2016, Heraklion, Greece, May 29, 2016. CEUR Workshop Proceedings, vol. 1613. CEUR-WS.org (2016), http://ceur-ws.org/Vol-1613/paper_2.pdf
15. Gimpel, K., Schneider, N., O'Connor, B., Das, D., Mills, D., Eisenstein, J., Heilman, M., Yogatama, D., Flanigan, J., Smith, N.A.: Part-of-speech tagging for twitter: Annotation, features, and experiments. In: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Short Papers - Volume 2. HLT '11 (2011)
16. Go, A., Bhayani, R., Huang, L.: Twitter sentiment classification using distant supervision. Tech. rep., Stanford (2009)
17. Hatzivassiloglou, V., Wiebe, J.: Effects of adjective orientation and gradability on sentence subjectivity. In: COLING. pp. 299–305 (2000)
18. Kim, S.M., Hovy, E.H.: Crystal: Analyzing predictive opinions on the web. In: EMNLP-CoNLL. pp. 1056–1064 (2007)
19. Kouloumpis, E., Wilson, T., Moore, J.D.: Twitter sentiment analysis: The good the bad and the omg! In: Proceedings of the Fifth International Conference on Weblogs and Social Media, Barcelona, Spain (2011)
20. Liu, B., Zhang, L.: Mining Text Data, chap. A Survey of Opinion Mining and Sentiment Analysis, pp. 415–463. Springer US (2012)
21. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 (2013)
22. Pang, B., Lee, L.: Opinion mining and sentiment analysis. Foundations and Trends in Information Retrieval 2(1-2) (2008)
23. Pang, B., Lee, L., Vaithyanathan, S.: Thumbs up?: Sentiment classification using machine learning techniques. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing. pp. 79–86. EMNLP '02, Association for Computational Linguistics, Stroudsburg, PA, USA (2002), http://dx.doi.org/10.3115/1118693.1118704
24. Petrucci, G., Dragoni, M.: An information retrieval-based system for multi-domain sentiment analysis. In: Gandon, F., Cabrio, E., Stankovic, M., Zimmermann, A. (eds.) Semantic Web Evaluation Challenges - Second SemWebEval Challenge at ESWC 2015, Portorož, Slovenia, May 31 - June 4, 2015, Revised Selected Papers. Communications in Computer and Information Science, vol. 548, pp. 234–243. Springer (2015), http://dx.doi.org/10.1007/978-3-319-25518-7_20

25. Petrucci, G., Dragoni, M.: The IRMUDOSA system at ESWC-2016 challenge on semantic sentiment analysis. In: Sack, H., Dietze, S., Tordai, A., Lange, C. (eds.) Semantic Web Challenges - Third SemWebEval Challenge at ESWC 2016, Heraklion, Crete, Greece, May 29 - June 2, 2016, Revised Selected Papers. Communications in Computer and Information Science, vol. 641, pp. 126–140. Springer (2016), https://doi.org/10.1007/978-3-319-46565-4_10
26. Rexha, A., Kröll, M., Dragoni, M., Kern, R.: Exploiting propositions for opinion mining. In: Sack, H., Dietze, S., Tordai, A., Lange, C. (eds.) Semantic Web Challenges - Third SemWebEval Challenge at ESWC 2016, Heraklion, Crete, Greece, May 29 - June 2, 2016, Revised Selected Papers. Communications in Computer and Information Science, vol. 641, pp. 121–125. Springer (2016), https://doi.org/10.1007/978-3-319-46565-4_9
27. Rexha, A., Kröll, M., Dragoni, M., Kern, R.: Polarity classification for target phrases in tweets: A word2vec approach. In: Sack, H., Rizzo, G., Steinmetz, N., Mladenic, D., Auer, S., Lange, C. (eds.) The Semantic Web - ESWC 2016 Satellite Events, Heraklion, Crete, Greece, May 29 - June 2, 2016, Revised Selected Papers. Lecture Notes in Computer Science, vol. 9989, pp. 217–223 (2016), https://doi.org/10.1007/978-3-319-47602-5_40
28. Riloff, E., Patwardhan, S., Wiebe, J.: Feature subsumption for opinion analysis. In: EMNLP. pp. 440–448 (2006)
29. Rosenthal, S., Nakov, P., Kiritchenko, S., Mohammad, S., Ritter, A., Stoyanov, V.: Semeval-2015 task 10: Sentiment analysis in twitter. In: Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015). Association for Computational Linguistics (2015)
30. Rosenthal, S., Ritter, A., Nakov, P., Stoyanov, V.: Semeval-2014 task 9: Sentiment analysis in twitter. In: Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014). Association for Computational Linguistics and Dublin City University (2014)
31. Turney, P.D.: Thumbs up or thumbs down?: Semantic orientation applied to unsupervised classification of reviews. In: Proceedings of the 40th Annual Meeting on Association for Computational Linguistics. Association for Computational Linguistics (2002)
32. Wiebe, J., Wilson, T., Cardie, C.: Annotating expressions of opinions and emotions in language. Language Resources and Evaluation 1(2) (2005)
33. Wilson, T., Kozareva, Z., Nakov, P., Alan, R., Rosenthal, S., Stoyonov, V.: Semeval-2013 task 2: Sentiment analysis in twitter. In: Proceedings of the 7th International Workshop on Semantic Evaluation. Association for Computation Linguistics (2013)
34. Wilson, T., Wiebe, J., Hwa, R.: Recognizing strong and weak opinion clauses. Computational Intelligence 22(2), 73–99 (2006)