

LTG-Oslo Hierarchical Multi-task Network: The Importance of Negation for Document-level Sentiment in Spanish

Jeremy Barnes^[0000–0002–8043–8058]

Language Technology Group
University of Oslo, Norway
jeremycb@ifi.uio.no

Abstract. This paper details LTG-Oslo team’s participation in the sentiment track of the NEGES 2019 evaluation campaign (We make the code available at https://github.com/jbarnesspain/neges_2019). We participated in the task with a hierarchical multi-task network, which used shared lower-layers in a deep BiLSTM to predict negation, while the higher layers were dedicated to predicting document-level sentiment. The multi-task component shows promise as a way to incorporate information on negation into deep neural sentiment classifiers, despite the fact that the absolute results on the test set were relatively low for a binary classification task.

Keywords: Sentiment Analysis · Negation · Multi-task

1 Introduction

Sentiment analysis has improved greatly over the last decade, moving from models trained on hand-engineered features [29,10] to neural models that are trained in an end-to-end fashion [34]. The success of these neural architectures is often attributed to their ability to capture compositionality effects [34,23], of which negation is the most common and influential for sentiment analysis [42]. However, recent research has shown that these models are still not able to fully resolve the effect that negation has on sentence-level sentiment [4].

Explicit negation detection has proven useful to create features for lexicon-based sentiment models [8,9] and machine-learning approaches to sentiment classification [22]. At the same time, these approaches build upon work on negation detection as its own task [40,26,18].

More recent approaches to sentiment, however, have concentrated on learning the effects of negation in an end-to-end fashion. Current state-of-the-art approaches employ neural networks which implicitly learn to resolve negation, by either directly training on sentiment annotated data [34,37], or by pre-training

the model on a language modeling task [31,11]. State-of-the-art neural methods, however, have not attempted to harness explicit negation detection models and annotated negation datasets to improve results. We hypothesize that multi-task learning (MTL) [5,7] is an appropriate framework to incorporate negation information into neural models.

In this paper, we propose a multi-task learning approach to explicitly incorporate negation annotated data into a neural sentiment model. We show that this approach improves the final result on the 2019 NEGES shared task [17], although our model performs weakly in absolute terms.

2 Related Work

In this section, we briefly review previous work that is relevant to *(i)* attempts to use negation information in sentiment analysis, *(ii)* research on negation detection as a separate task, and *(iii)* multi-task learning.

2.1 Negation informed Sentiment Analysis

Negation is a pervasive linguistic phenomenon which has a direct effect on the sentiment of a text [42]. Take the following example from the SFU ReviewSP-Neg training data, where the negation cue is shown in **bold** and the scope is underlined.

Example 1.

El hotel está situado en la puerta de toledo, **no** está lejos del centro.

The English translation is “The hotel is located at the *puerta de toledo*, it is not far from the center.” A sentiment classification model must be able to identify the relevant sentiment words (in this case “lejos del centro”), negation cues (“no”), and resolve the scope in order to correctly predict that this sentence expresses negative polarity. Intuitively, a sentiment model that has access to negation scope information should perform better than a non-informed version.

The first approaches to detecting negation scope for sentiment used heuristics, such as assuming all tokens between a negation cue and the next punctuation mark are in scope [16]. However, this simplification does not work well on noisy text, such as tweets, or texts that use more complex syntax, such as those in the political domain.

Later research showed that using machine-learning techniques to detect the scope of negation could improve both lexicon-based [8,9] and machine learning [22] classification of sentiment.

2.2 Negation detection

Approaches to negation analysis often decompose the task into two sub-tasks, performing (i) negation cue detection, followed by (ii) scope detection.

Much work was done within the biomedical domain [25,27,39] due largely to the availability of the BioScope corpus [40], which is annotated for negation cues and scopes. The *SEM shared task [26] instead focused on detection of negation cues and scopes in a corpus of sentences taken from the works of Aurther Conan Doyle.

Traditional approaches to the task of negation detection have typically employed a wide range of hand-crafted features describing a number of both lexical, morphosyntactic and even semantic properties of the text [33,28,22,41,12]. More recently, research has moved towards using neural models such as CNNs [32], feed-forward networks, or LSTMs [13], finding that these architectures often outperform the previous methods, while requiring less hand-crafting of features.

2.3 Multi-task learning

Multi-task learning (MTL) is an approach to machine learning where a single model is trained simultaneously on two tasks. By restricting the search space of possible representations to those that are predictive for both tasks, we attempt to give the model a useful inductive bias [5].

Hard parameter sharing [5], which assumes that all layers are shared between tasks except for the final predictive layer, is the simplest way to implement a multi-task model. When the main task and auxiliary task are closely related, this approach has been shown to be an effective way to improve model performance [7,30,24,2]. On the other hand, [35] find that it is better to make predictions for low-level auxiliary tasks at lower layers of a multi-layer MTL setup. They also suggest that under the hard-parameter framework auxiliary tasks need to be sufficiently similar to the main task for MTL to improve over the single-task baseline.

In this work, we implement a multi-task learning where the lower layers of a deep neural network are shared for the main and auxiliary tasks (in our case sentiment classification and negation detection, respectively), while higher layers are allowed to adapt to the main task.

3 Model

We propose a *hierarchical multi-task model* (see Figure 1) which relies on a BiLSTM to create a representation for each sentence in a document, and a second BiLSTM to aggregate these sentence representations into a full document representation. In this section, we first describe the negation submodel, then the sentiment submodel, and finally the multi-task model.

3.1 Negation Model

In previous work on negation detection, it is common to model negation scope as a two step process, where first the negation cues are identified, and then negation scope is determined. However, we hypothesize that within a multi-task

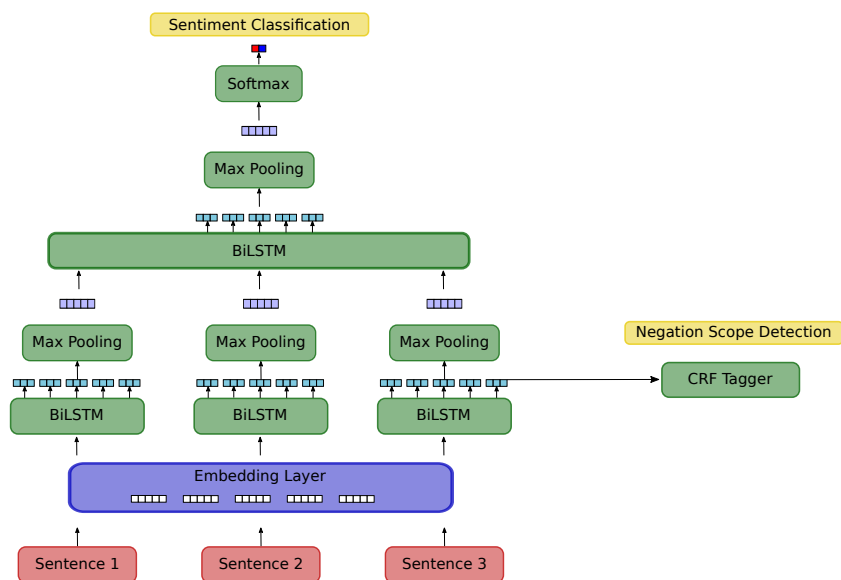


Fig. 1: Hierarchical multi-task model. The lower BiLSTM is used both to perform sequence-tagging of negation, as well as creating sentence-level features. These features are then aggregated using a second BiLSTM layer and used for predicting the sentiment at document-level.

framework, it is more beneficial for a network to learn to both identify cues and resolve scope jointly. Therefore, we model negation as a *sequence labeling task* with BIO tags. In the cases where there are more than one negation scope in a sentence that overlap, we flatten these multiple representations, as shown in Figure 2. The negation model, therefore, attempts to identify all cues and all scopes in a sentence at the same time. Note that scopes can also begin before the negation cue and also be discontinuous. While this is an oversimplification of the full negation scope task, we argue that in order to classify sentiment, it is enough for a model to know which tokens are negated.

The negation model is comprised of an embedding layer which embeds the tokens for each sentence. The embeddings pass to a bidirectional Long Short-Term Memory module (BiLSTM), which creates contextualized representations of each word. A linear chain conditional random field (CRF) uses the output of the BiLSTM layer as features. We use Viterbi decoding and minimize the negative log likelihood of CRF predictions.

	Es	que	no	nos	ayudó	,	y	luego	ni	siquera	llamó
negation labels:	O	O	CUE ¹	N ¹	N ¹	O	O	O	CUE ²	CUE ²	N ²
BIO labels:	O	O	B_cue	B_neg	I_neg	O	O	O	B_cue	I_cue	B_neg

Fig. 2: An example of the negation which has been converted to BIO labels. Although the example here shows two negation structures where the cue is at the beginning of the scope, there are also examples where the scope begins before the cue or is discontinuous.

3.2 Sentiment Model

As mentioned above, the sentiment model uses a hierarchical approach. For each sentence in a document, we first extract features with a BiLSTM. We take the max of the BiLSTM output as a representation for the sentence. This is then passed to a second BiLSTM layer, after which we again take the max. We use a softmax layer to compute the sentiment predictions for each document and minimize the cross entropy loss. As a baseline, we train a single-task sentiment model (STL) on the available sentiment data.

3.3 Multi-task Model

For the hierarchical multi-task model (MTL), we train both tasks simultaneously by sequentially training the negation classification model for one full epoch and then training the sentiment model. We use Adam as an optimizer, and a dropout layer (0.3) after the embedding layer to regularize the model, as this is common for both the main and auxiliary tasks.

4 Experimental Setup

Given that neural models are sensitive to random initialization, we perform five runs for each model on the development data with different random seeds and report both mean accuracy and standard deviation across the five runs. As the final submission required a single prediction for each document, we take a majority vote of the five learned classifiers in order to provide an ensemble prediction.

Besides the proposed STL and MTL models, we also compare with a baseline (BOW) which uses an L2 regularized logistic regression classifier trained on a bag-of-words representation of the documents. We choose the optimal C parameter on the development data.

4.1 Dataset

The SFU ReviewSP-NEG dataset [19] provided in the shared task contains 400 Spanish-language reviews from eight domains (books, cars, cellphones, computers,

hotels, movies, music, and washing machines) which also contain annotations for negation cues, negation scope, and relevance of the negation to sentiment. The participants were provided with the train and dev splits, while the test split was kept from participants until after the final results were posted. Table 1 shows the statistics of the dataset.

Previous work [20] reported Macro F1 score of 75.89 when using a Bayesian logistic regression classifier trained with bag-of-words features plus negation features that indicate that negation changes the polarity of the negated phrase. However, these results are not comparable to those obtained in the shared task, as the authors evaluated their model using 10-fold cross-validation and not on the test set provided by the organizers. Additionally, they had access to negation information in the test set, which participants in the shared task do not.

Table 1: Statistics of the document-level sentiment (number of documents) and negation (number of negation structures) data provided by the organizers of the shared task.

Task	Train	Dev	Test
Document-level Sentiment	264	56	80
Negative Structures	2,733	645	949

4.2 Model performance

As we only had access to the gold labels on the development set, we report the mean and average accuracy of all three models (BOW, STL, MTL) in Table 2. Additionally, we show the official accuracy score of the MTL model on the test set¹. BOW and STL achieve the same performance, with 71.4 accuracy on the dev set. MTL improves 1.1 percentage points over the other two models on the dev set, and reaches 66.2 accuracy on the test set. In absolute terms, the performance of all models is weak for a binary document-level classification task. This is likely due to the small number of training examples available, as well as the number of domains, which has been shown to be more problematic for machine-learning approaches than lexicon-based approaches [36].

4.3 Error Analysis

Given that the classification task is performed at document-level, it is often difficult to determine what exactly was the cause of a change in prediction from one model to another. Instead, Figure 3 shows a relative confusion matrix of the development results, where positive numbers (dark purple) indicate that the MTL

¹ Note that we do not have access to the gold sentiment or negation labels on the test set, so we cannot perform multiple runs, but must rely on the organizers evaluation.

Table 2: Accuracy of the models on the development and test data. Neural models also report mean accuracy and standard deviation on the development data over five runs with different random seeds.

Model	Dev	Test
BOW	71.4	-
STL	71.4 (5.2)	-
MTL	72.5 (1.8)	66.2

model made more predictions in that square than the STL model and negative numbers (white) indicate fewer predictions. On the development data, the MTL model tends to help with the negative class, while adding little to the positive class. The number of negation structures per class (shown in Table 3) shows that there are more negation structures in documents labeled with negative sentiment in the development set, which seems to corroborate the idea that the MTL model is able to use negation information to improve the results on the negative class.




Fig. 3: A relative confusion matrix, where positive numbers (dark purple) indicate that the MTL model made more predictions in that square than the STL model and negative numbers (white) indicate fewer predictions.

Table 3: Number of negation structures per sentiment class found in the training and development data.

	Train	Dev
Positive	1,421	303
Negative	1,312	342

5 Conclusion and Future Work

In this paper, we have detailed our participation in the 2019 Neges shared task. Our approach, the hierarchical multi-task negation model, did not give a strong performance in absolute numbers on the test set (66% accuracy), but does indicate that multi-task learning is an appropriate framework for incorporating negation information into sentiment models, improving from 71.4 to 72.5 accuracy on the development set.

The hierarchical RNN model used in this participation is similar to strong performing approaches at sentence-level. However, it is not clear that it is the most adequate model for document-level classification. Convolutional neural networks [21] or self-attention networks [1] have shown good performance for text classification and may be better models for document-level sentiment tasks.

Additionally, the small training set size for the sentiment task (271 documents) and number of domains (8) complicates the use of deep neural architectures. Lexicon-based and linear machine-learning approaches have shown to perform quite well under these circumstances [36,9]. In the future, it would be interesting to use distant supervision [38,15] to augment the sentiment signal, or cross-lingual approaches [6,3] to improve the results.

In this work we have only explored using a sequence-labeling approach to negation scope. It would be interesting to incorporate state-of-the-art negation scope models [14] into a multi-task setup.

Finally, the SFU ReviewSP-NEG dataset has several additional levels of annotation, *i.e.* if a negation structure changes the polarity of the tokens in scope or the final polarity after negation. Future work should explore the use of this information further.

Acknowledgements

This work has been carried out as part of the SANT project, funded by the Research Council of Norway (grant number 270908).

References

1. Ambartsoumian, A., Popowich, F.: Self-Attention: A Better Building Block for Sentiment Analysis Neural Network Classifiers. In: Proceedings of the 9th Workshop

- on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis. pp. 130–139. Association for Computational Linguistics (2018), <http://aclweb.org/anthology/W18-6219>
2. Augenstein, I., Ruder, S., Søgaard, A.: Multi-Task Learning of Pairwise Sequence Classification Tasks over Disparate Label Spaces. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers). pp. 1896–1906. Association for Computational Linguistics (2018). <https://doi.org/10.18653/v1/N18-1172>, <http://aclweb.org/anthology/N18-1172>
 3. Barnes, J., Klinger, R., Schulte im Walde, S.: Bilingual Sentiment Embeddings: Joint Projection of Sentiment Across Languages. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 2483–2493. Association for Computational Linguistics (2018), <http://aclweb.org/anthology/P18-1231>
 4. Barnes, J., Øvrelid, L., Velldal, E.: Sentiment Analysis is Not Solved!: Assessing and Probing Sentiment Classifiers. In: Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP. p. to appear. Association for Computational Linguistics, Florence, Italy (2019)
 5. Caruana, R.: Multitask Learning: A Knowledge-Based Source of Inductive Bias. In: Proceedings of the Tenth International Conference on Machine Learning. pp. 41–48. Morgan Kaufmann (1993)
 6. Chen, X., Athiwaratkun, B., Sun, Y., Weinberger, K.Q., Cardie, C.: Adversarial Deep Averaging Networks for Cross-Lingual Sentiment Classification. CoRR [abs/1606.01614](https://arxiv.org/abs/1606.01614) (2016), <http://arxiv.org/abs/1606.01614>
 7. Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., Kuksa, P.: Natural Language Processing (almost) from Scratch. *Journal of Machine Learning Research* **12**, 2493–2537 (2011)
 8. Councill, I., McDonald, R., Velikovich, L.: What’s Great and What’s not: Learning to Classify the Scope of Negation for Improved Sentiment Analysis. In: Proceedings of the Workshop on Negation and Speculation in Natural Language Processing. pp. 51–59. University of Antwerp, Uppsala, Sweden (Jul 2010), <https://www.aclweb.org/anthology/W10-3110>
 9. Cruz, N.P., Taboada, M., Mitkov, R.: A Machine-learning Approach to Negation and Speculation Detection for Sentiment Analysis. *Journal of the Association for Information Science and Technology* **67**(9), 2118–2136 (2016). <https://doi.org/10.1002/asi.23533>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/asi.23533>
 10. Das, S.R., Chen, M.Y.: Yahoo! For Amazon: Sentiment Extraction from Small Talk on the Web. *Management Science* **53**(9), 1375–1388 (Sep 2007). <https://doi.org/10.1287/mnsc.1070.0704>, <http://dx.doi.org/10.1287/mnsc.1070.0704>
 11. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). pp. 4171–4186. Association for Computational Linguistics, Minneapolis, Minnesota (Jun 2019), <https://www.aclweb.org/anthology/N19-1423>
 12. Enger, M., Velldal, E., Øvrelid, L.: An Open-source Tool for Negation Detection: a Maximum-margin Approach. In: Proceedings of the EACL workshop on Computational Semantics Beyond Events and Roles (SemBEaR). pp. 64–69. Valencia, Spain (2017)

13. Fancellu, F., Lopez, A., Webber, B.: Neural Networks For Negation Scope Detection. In: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 495–504. Association for Computational Linguistics, Berlin, Germany (Aug 2016). <https://doi.org/10.18653/v1/P16-1047>, <https://www.aclweb.org/anthology/P16-1047>
14. Fancellu, F., Lopez, A., Webber, B., He, H.: Detecting Negation Scope is Easy, except when it isn't. In: Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers. pp. 58–63. Association for Computational Linguistics, Valencia, Spain (Apr 2017), <https://www.aclweb.org/anthology/E17-2010>
15. Felbo, B., Mislove, A., Sogaard, A., Rahwan, I., Lehmann, S.: Using Millions of Emoji Occurrences to Learn Any-domain Representations for Detecting Sentiment, Emotion and Sarcasm. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing. pp. 1615–1625. Association for Computational Linguistics, Copenhagen, Denmark (Sep 2017). <https://doi.org/10.18653/v1/D17-1169>, <https://www.aclweb.org/anthology/D17-1169>
16. Hu, M., Liu, B.: Mining Opinion Features in Customer Reviews. In: Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2004). pp. 168–177 (2004)
17. Jiménez-Zafra, S.M., Cruz Díaz, N.P., Morante, R., Martín-Valdivia, M.T.: NEGES 2019 Task: Negation in Spanish. In: Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2019). CEUR Workshop Proceedings, CEUR-WS, Bilbao, Spain (2019)
18. Jiménez-Zafra, S.M., Díaz, N.P.C., Morante, R., Martín-Valdivia, M.T.: NEGES 2018: Workshop on Negation in Spanish. *Procesamiento del Lenguaje Natural* **62**, 21–28 (2019)
19. Jiménez-Zafra, S.M., Taulé, M., Martín-Valdivia, M.T., Ureña-López, L.A., Martí, M.A.: SFU ReviewSP-NEG: a Spanish Corpus Annotated with Negation for Sentiment Analysis. A Typology of Negation Patterns. *Language Resources and Evaluation* **52**(2), 533–569 (Jun 2018)
20. Jiménez-Zafra, S.M., Martín-Valdivia, M.T., Molina-González, M.D., Ureña-López, L.A.: Relevance of the SFU ReviewSP-NEG Corpus Annotated with the Scope of Negation for Supervised Polarity Classification in Spanish. *Information Processing & Management* **54**(2), 240 – 251 (2018). <https://doi.org/https://doi.org/10.1016/j.ipm.2017.11.007>, <http://www.sciencedirect.com/science/article/pii/S0306457317306258>
21. Kim, Y.: Convolutional Neural Networks for Sentence Classification. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP) (2014)
22. Lapponi, E., Read, J., Øvrelid, L.: Representing and Resolving Negation for Sentiment Analysis. In: Proceedings of the 2012 IEEE 12th International Conference on Data Mining Workshops. pp. 687–692. ICDMW '12, IEEE Computer Society, Washington, DC, USA (2012). <https://doi.org/10.1109/ICDMW.2012.23>, <https://doi.org/10.1109/ICDMW.2012.23>
23. Linzen, T., Dupoux, E., Goldberg, Y.: Assessing the Ability of LSTMs to Learn Syntax-Sensitive Dependencies. *Transactions of the Association for Computational Linguistics* **4**, 521–535 (2016). https://doi.org/10.1162/tacl_a00115, <https://www.aclweb.org/anthology/Q16-1037>
24. Martínez Alonso, H., Plank, B.: When is Multitask Learning Effective? Semantic Sequence Prediction under Varying Data Conditions. In: Proceedings of the 15th

- Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers. pp. 44–53. Association for Computational Linguistics (2017), <http://aclweb.org/anthology/E17-1005>
25. Morante, R., Liekens, A., Daelemans, W.: Learning the Scope of Negation in Biomedical Texts. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP) (2008)
 26. Morante, R., Blanco, E.: *SEM 2012 Shared Task: Resolving the Scope and Focus of Negation. In: *SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012). pp. 265–274. Association for Computational Linguistics, Montréal, Canada (7-8 Jun 2012), <https://www.aclweb.org/anthology/S12-1035>
 27. Morante, R., Daelemans, W.: A Metalearning Approach to Processing the Scope of Negation. In: Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL) (2009)
 28. Packard, W., Bender, E.M., Read, J., Oepen, S., Drīdan, R.: Simple Negation Scope Resolution through Deep Parsing: A Semantic Solution to a Semantic Problem. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (2014)
 29. Pang, B., Lee, L., Vaithyanathan, S.: Thumbs up? Sentiment Classification using Machine Learning Techniques. In: Proceedings of the ACL-02 Conference on Empirical methods in natural language processing-Volume 10. pp. 79–86. Association for Computational Linguistics (2002)
 30. Peng, N., Dredze, M.: Multi-task Domain Adaptation for Sequence Tagging. In: Proceedings of the 2nd Workshop on Representation Learning for NLP. pp. 91–100. Association for Computational Linguistics, Vancouver, Canada (Aug 2017). <https://doi.org/10.18653/v1/W17-2612>, <https://www.aclweb.org/anthology/W17-2612>
 31. Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., Zettlemoyer, L.: Deep Contextualized Word Representations. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers). pp. 2227–2237. Association for Computational Linguistics (2018). <https://doi.org/10.18653/v1/N18-1202>, <http://aclweb.org/anthology/N18-1202>
 32. Qian, Z., Li, P., Zhu, Q., Zhou, G., Luo, Z., Luo, W.: Speculation and Negation Scope Detection via Convolutional Neural Networks. In: The 2016 Conference on Empirical Methods in Natural Language Processing (2016)
 33. Read, J., Velldal, E., Øvreliid, L., Oepen, S.: UiO1: Constituent-Based Discriminative Ranking for Negation Resolution. In: Proceedings of the First Joint Conference on Lexical and Computational Semantics (*SEM). Montreal (2012)
 34. Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C., Ng, A., Potts, C.: Recursive Deep Models for Semantic Compositionality over a Sentiment Treebank. Proceedings of the EMNLP 2013 pp. 1631–1642 (2013)
 35. Søggaard, A., Goldberg, Y.: Deep Multi-task Learning with Low Level Tasks Supervised at Lower Layers. In: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers). pp. 231–235. Association for Computational Linguistics (2016). <https://doi.org/10.18653/v1/P16-2038>, <http://aclweb.org/anthology/P16-2038>
 36. Taboada, M., Brooke, J., Tofiloski, M., Voll, K., Stede, M.: Lexicon-based Methods for Sentiment Analysis. Computational Linguistics **37**(2), 267–307 (Jun

- 2011). https://doi.org/10.1162/COLI_a00049, http://dx.doi.org/10.1162/COLI_a_00049
37. Tai, K.S., Socher, R., Manning, C.D.: Improved Semantic Representations from Tree-structured Long Short-term Memory Networks. Association for Computational Linguistics 2015 Conference pp. 1556–1566 (2015)
 38. Tang, D., Wei, F., Yang, N., Zhou, M., Liu, T., Qin, B.: Learning Sentiment-Specific Word Embedding for Twitter Sentiment Classification. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 1555–1565 (2014)
 39. Velldal, E., Øvrelid, L., Read, J., Oepen, S.: Speculation and Negation: Rules, Rankers, and the Role of Syntax. Computational Linguistics **38**(2), 369–410 (2012)
 40. Vincze, V., Szarvas, G., Farkas, R., Móra, G., Csirik, J.: The BioScope Corpus: Biomedical Texts Annotated for Uncertainty, Negation and their Scopes. BMC bioinformatics **Suppl 11**(Suppl 11) (2008)
 41. White, J.P.: UWashington: Negation Resolution using Machine Learning Methods. In: Proceedings of the First Joint Conference on Lexical and Computational Semantics (*SEM). Montreal (2012)
 42. Wiegand, M., Balahur, A., Roth, B., Klakow, D., Montoyo, A.: A Survey on the Role of Negation in Sentiment Analysis. In: Proceedings of the Workshop on Negation and Speculation in Natural Language Processing. pp. 60–68 (2010)