





AE-CharCNN: Char-Level Convolutional Neural Networks for Aspect-Based Sentiment Analysis

Ulisses Brisolara Corrêa^{1,2}(✉)  and Ricardo Matsumura Araújo² 

¹ Computer Science Program (PPGC), Center for Technological Advancement (CDTec), Federal University of Pelotas (UFPEL), Pelotas, RS 96010-610, Brazil
`{ub.correa,ricardo}@inf.ufpel.edu.br`

² Sul-rio-grandense Federal Institute of Education, Science, and Technology (IFSul), Charqueadas, RS 96745-000, Brazil
`ulissescorrea@charqueadas.ifsul.edu.br`

Abstract. Sentiment Analysis was developed to support individuals in the harsh task of obtaining significant information from large amounts of non-structured opinionated data sources, such as social networks and specialized reviews websites. A yet more challenging task is to point out which part of the target entity is addressed in the opinion. This task is called Aspect-Based Sentiment Analysis. The majority of work focuses on coping with English text in the literature, but other languages lack resources, tools, and techniques. This paper focuses on Aspect-Based Sentiment Analysis for Accommodation Services Reviews written in Brazilian Portuguese. Our proposed approach uses Convolution Neural Networks with inputs in Character-level. Results suggest that our approach outperforms lexicon-based and LSTM-based approaches, displaying state-of-the-art performance for binary Aspect-Based Sentiment Analysis.

Keywords: Aspect-Based Sentiment Analysis · Convolutional Neural Networks · Char-level embedding

1 Introduction

We are living in the social media era. People's opinions about a plethora of topics can be abundantly found on the Internet. This kind of information is extremely valuable to other individuals, organizations, and government entities.

According to Liu [12], whenever people have to make decisions they seek out others' opinions, usually friends and family, to support their choices. Social networks allows for opinions to be more easily accessible. Now one can easily get opinions directly from users that already have bought a product or from some experts in this kind of product. Thus individuals can take advantage of the experiences of others. But the amount of opinionated content generated by users

on the Internet day by day makes it too difficult to manually distill significant information, even for big organizations.

Sentiment Analysis is the Computer Science response to this problem. This research area, also known as Opinion Mining, was developed to support individuals in the harsh task of obtain significant information from large amounts of non-structured opinionated data sources, like social networks and specialized reviews websites (such as: IMDB, Rotten Tomatoes, Metacritic, etc.).

Extract significant information about opinions from raw text is a hard task. Opinionated text can present opinions towards different entities, or conflicting opinions about same entity. Correctly address the target of an opinion is crucial to ensure the perfect understanding of the expressed sentiment. Aspect-based Sentiment Analysis (ABSA) is a sentiment analysis process that is done paying particular attention to the aspect to which the sentiment is addressed.

Opinion Mining research focuses on tasks related to: sentiment classification, subjectivity detection, sentiment orientation detection, etc. Works in this research area usually apply techniques based on classical Natural Language Processing (NLP), such as: hand-crafted linguistic feature extraction, structural pattern detection, and shallow models (e.g. Support Vector Machines and Logistic Regression) [21].

Deep Learning architectures and algorithms are getting significant results in several Computer Science problems. Computer vision is clearly the field where Deep Learning (DL) application has displayed more significant gains, mostly with Convolutional Neural Networks (CNN), starting with the work of Krizhevsky et al. [11]. Based on these good results in computer vision and pattern recognition, NLP researchers have been trying to apply Deep Learning on textual input. Recently, NLP research has been gradually adopting more complex and flexible models, based on Deep Learning, as can be seen in recent top conferences [21].

Sentiment Analysis, as an NLP task, has also been taking advantage of these techniques [17]. Works based on character-level convolutional networks have achieved promising results in text classification and document-level sentiment analysis [19,20,23]. This kind of approach deals naturally with typos, encoding errors, and rare words. Moreover, they are less susceptible to external resources influence on performance, when comparing with lexicon-based sentiment classifiers, which are highly dependent of those resources [8].

Works focused in dealing with English text are plenty in the literature, however other languages lack good tools and techniques to NLP and Sentiment Analysis [6,7]. Because of this disparity some authors tried approaches based on language translation with reasonable results [1,15], but making the approach dependant of the translation tool.

This work aims to apply Aspect-Embedding to make Character-level Convolution Neural Networks compliant with sentiment classification in ABSA. Previous works used Char-level CNNs to classify sentiment in document-level. But, following Chaves et al. [3], up to 86% of hotel's reviews present both positive and negative opinions about different hotel's aspects (parts or attributes of the

hotel). Then document-level sentiment analysis will have limited performance in this domain. In our approach, the network receives the review and the aspect to be assessed in Brazilian Portuguese. We evaluate our approach against two state-of-art approaches [7, 18].

The remaining of this text is organized as follow. A succinct overview about Sentiment Analysis is shown in Sect. 2. Section 3 discusses relevant related work. Next, Sect. 4 presents our proposed approach. Experimental methodology and results are shown and analyzed in Sect. 5. Finally Sect. 6 summarizes our contributions and future work.

2 Sentiment Analysis

2.1 Opinion Definition

Liu [12] proposes a taxonomy for opinions. Two main classes are highlighted: regular and comparative opinions. *Regular Opinions* refers a entity directly or indirectly, while *Comparative Opinions* expresses relations between two or more entities, focusing on their similarities or differences.

Most works in Sentiment Analysis deal with Regular Opinions, frequently referenced simply as *Opinions*. Liu [12] formally defines an opinion as a quadruple,

$$Opinion : (g, s, h, t), \quad (1)$$

where:

- g – entity towards sentiment is addressed;
- s – sentiment of the opinion;
- h – opinion holder, the person or organization who holds the opinion;
- t – time when the opinion was expressed.

The full understanding of all these elements is crucial to ensure a Sentiment Analysis system with the best results. In a more fine grained way we can decompose g to discover the specific *aspect* towards which the opinion is targeted. Aspects can be parts or attributes of entities, e.g. room and comfort are aspects of hotel.

Liu [12] states that an entity (e) can be a product, service, topic, person, organization, issue, or event. Entities can be described as pairs:

$$e : (T, W), \quad (2)$$

where:

- T – hierarchy of parts, subparts, and so on;
- W – set of attributes of e .

Sentiment Analysis can be done at different levels of granularity. Some works assign a single sentiment orientation to a whole document. Other works break documents in finer grained pieces (such as: blocks, sentences or paragraphs), then

extract their sentiment orientation. Both approaches suppose that the text being analyzed has an opinion about an entity chosen previously. This assumption can lead the classifier to high error rates by assigning sentiment orientation to the wrong target entity.

A more precise Sentiment Analysis approach is aspect-based. Aspect-Based Sentiment Analysis (ABSA) summarizes the sentiment orientation towards each entity's aspects. These fine-grained results are more meaningful, enabling a better understanding of the entity pros and cons.

According to Pang et al. [13], ABSA should be split in two subtasks:

1. the identification of product aspects (features and subparts), and
2. extraction of the sentiment orientation (opinions) associated with these aspects.

This work proposes a novel convolutional network architecture to tackle with the last subtask.

3 Related Works

Aspect-Based Sentiment Analysis is crucial to ensure that companies, governments, or even individuals get adequate information from large data sets of opinionated text. Liu [12] states that traditional ML methods, such as Support Vector Machines or Naïve Bayes, are inadequate to deal with ABSA. Currently, just a few works in literature deals with Portuguese texts in aspect-level [7].

Freitas [7] proposes a novel lexicon-based approach to sentiment analysis in aspect-level. Proposed approach comprises four steps: pre-processing, feature identification (aspect extraction), polarity identification, and polarity summarizing. Work was evaluated over a data set of accommodation sector. Author made a fine-grain analysis, at aspect-level, where explicit and implicit references to features were recognized using a domain ontology. The best results were obtained using POS-tagging extracted with TreeTagger, sentiment lexicon synsets with polarities from Onto.PT, and syntactic pattern detection, adjective position rule for negative polarity identification and presence or absence of polarity words in surrounding window for positive polarity identification.

Research from Blaz and Becker [1] proposes a hybrid approach, where the sentiment classification is made using both Lexicon-based and Machine Learning (ML) based approaches. They propose the automatically creation of a Domain Dictionary with opinionated terms in the Information Technology (IT) domain. Then they use these terms to filter ticket for sentiment analysis. For evaluate the work, the authors propose 3 approaches (one based on dictionary and two based on position of tokens in tickets) to calculating the polarity of terms in tickets. The study was developed using 34,895 tickets from 5 organizations, from which 2,333 tickets were randomly selected to compose a Gold Standard. The best results display an average precision and recall of 82.83% and 88.42%, which outperforms the compared sentiment analysis solutions.

Carvalho et al. [2] investigate the expression of opinions about human entities in UGC. A set of 2,800 online comments about news was manually annotated, following a rich annotation scheme designed for this purpose. The authors conclude that it is difficult to correctly identify positive opinions because they are much less frequent than negative ones and particularly exposed to verbal irony. The human entities recognition is also a challenge because they are frequently mentioned by pronouns, definite descriptions, and nicknames.

When analyzing works focused on other languages, we can find approaches based on newer techniques, such as Convolutional Neural Networks, Recurrent Networks, and Attention Mechanisms.

Convolutional Neural Networks already have been used to infer sentiment present on microblogs' posts [5, 14, 16, 17, 19, 20, 23]. Commonly using word-embeddings as input for Document-Level Sentiment Analysis, mainly over microblogs' posts. This kind of approach lacks support for words unseen on training, demanding extra pre-processing steps. For instance, input texts can be pre-processed with stemming or lemmatization to group words, reducing the effect of absent words in the *corpus* training set.

FastText word embedding method proposes to look at words as sets of n-grams, what can reduce this kind of problem, since the unknown word can be created by a composition of trained n-grams [10]. One can argue that word embeddings pre-trained over large data sets, like Wikipedia, are available, notwithstanding Sentiment Analysis is highly dependent on the text domain. Thus, this kind of approach tends to be hampered by the usage of domain bias and jargon.

In the other hand, Character-Level models cope with text in a smaller granularity, intrinsically dealing with typos and enabling a behaviour similar to stemming [23]. CharCNN [23], Conv-Char-S [19], and NNLS [20] exploit these advantages to text classification, including Document-Level Sentiment Analysis.

Wang et al. [18] propose a novel approach to ABSA using Long Short-Term Memories (LSTM) with Attention Mechanism. Their network has two inputs, one representing the text to gauge sentiment and another to represent which aspect we want the network to pay attention to.

Authors claim that it's the first approach to generate distinct embeddings for each input, one for aspects and another for reviews. Aspect-Embeddings (AE) differs from the embedding used to represent words in the networks' input. Wang et al. [18] evaluate three different approaches: AT-LSTM, AE-LSTM, and ATAE-LSTM. *AT* denotes approaches that employ Attention mechanism, while *AE* denotes approaches that use Aspect Embedding.

We propose approaches that differs from early work by applying char-level CNNs to evaluate sentiment in Aspect-level. In order to do so, we added an extra input to CharCNN [23] and NNLS [20] models. This extra input is an Aspect-Embedding, based on the one proposed by Wang et al. [18]. As far as we know, there's no work using char-level CNNs with an extra input representing aspects in an ABSA approach. A more detailed explanation is given in Sect. 4.

4 Proposed Approach

This work proposes an Aspect-Based Sentiment Analysis approach based on Character-level Convolutional Neural Networks, with an additional Aspect Embedding input.

ABSA is a specific kind of text classification problem. To extract the sentiment orientation of an opinionated text towards a specific entity’s aspect demands an extra input, addressing the aspect of interest or a model that predicts both sentiment orientation and target aspect.

Our approach takes advantage of the concepts of different approaches. We adopted Aspect-Embedding idea from Wang et al. [18] and applied in a Character-Level Convolutional Neural Networks, CharCNN [22]. Besides that, we tried a shallower char-level architecture, NNLS, which is smaller and presents better results than CharCNN on Document-Level Sentiment Analysis [20].

We call our approach Aspect-Embedding Character-Level Convolutional Neural Network (shortened as AE-CharCNN). AE-CharCNN takes advantage of the character-level CNN approach, avoiding complex pre-processing because it copes with text in a smaller granularity, intrinsically dealing with typos, encoding errors, and enabling a behaviour similar to stemming and making it easy to train a method tolerant to typos and encoding errors [19, 20, 23].

Figures 1 and 2 presents the two char-level CNN architectures we propose to solve ABSA. Figure 1 presents the architecture of AE-CharCNN-v1, based on the CharCNN model [23]. Figure 2 presents the AE-CharCNN-v2, based on the NNLS-v2 model [20]. Green blocks represent original model’s flow, while red blocks represent the flow we added to deal with aspects.

Both models are fed with char-level embeddings, starting the green flow. Aspect-embedding also is a char-level embedding, but does not share information with the embedding used to reviews. The features at the end of both flows are concatenated in a single input to the classifying stage, composed of three dense (fully connected) layers in AE-CharCNN-v1 and a single dense layer in AE-CharCNN-v2.

As can be seen in Figs. 1 and 2, we have two flows in ours network architectures, one for aspects and another for the reviews. Both flows receive a character-level input. Features generated by these flows are concatenated before the application of the fully connected (dense) layers.

In our approaches, aspect-embedding inputs cannot exceed 100 characters. Review inputs are set to 2204 characters, length of the longest review in the corpus. In both cases, if bigger inputs are fed they are cut in the limit length. Otherwise, if smaller inputs are given, they are padded to the dimensions aforementioned.

Both architectures weights were adjusted using Adam optimizer, during 20 epochs, 32 samples batch size, and with a learning rate set to $10e^{-3}$. All hidden layers use ReLU activation function. The kernel size of feature extraction layers can be seen in Figs. 1 and 2.

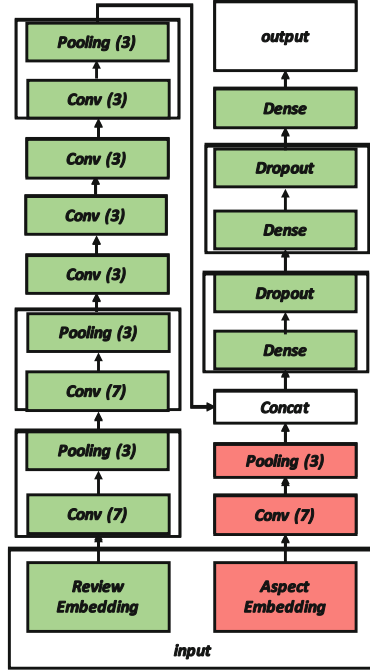


Fig. 1. Aspect-Embedding CharCNN (AE-CharCNN). (Color figure online)

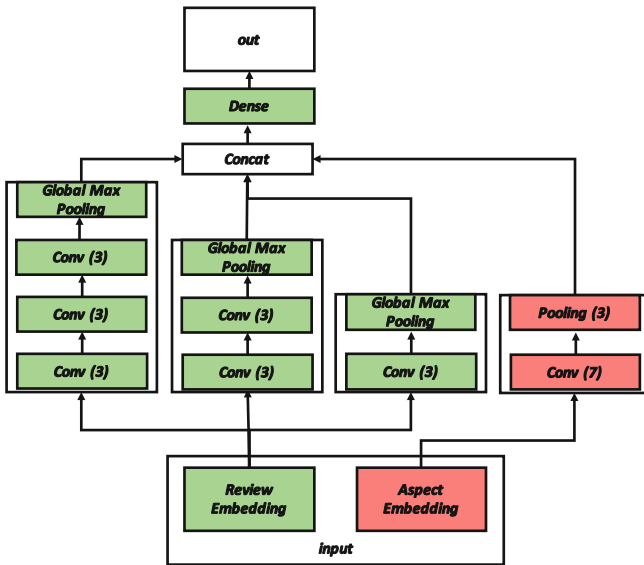


Fig. 2. Aspect-Embedding CharCNN v2 (AE-CharCNN-v2). (Color figure online)

Our model was trained with Cross Entropy Loss, which combines a LogSoftmax with Negative Log Likelihood Loss (NLLoss)¹. Original work from LeCun [22] was developed using torch library and uses NLLoss and an additional LogSoftmax layer at the end of network architecture. Weights were initialized using a Gaussian distribution. The mean and standard deviation used for initializing the large model is (0, 0.02) and small model (0, 0.05).

5 Experiments and Results

This Section presents our experiments over the proposed approach, as well as the comparison with two state-of-the-art approaches: Lexicon-based for Portuguese [7] and LSTM-based with Attention Mechanism adapted from English [18]. All approaches were tested over the same *corpus*. Source code for experiments presented here is available in <https://github.com/ulissesbcorrea/ae-charcnn>.

5.1 *Corpus* – TripAdvisor’s Reviews About Porto Alegre Accommodation Services

Our experiments were conducted over a *corpus* of reviews about Porto Alegre Accommodation Services extracted from TripAdvisor. Freitas [7] created that *corpus* manually annotated with aspect level sentiment orientation.

The *corpus* contains all TripAdvisor’s reviews about hotels located in Porto Alegre until May 2014. There are 987 entries with opinions about hotel’s aspects. A total of 194 hotel reviews were manually annotated, assigning polarity to opinions targeted to every hotel aspect directly or indirectly referenced in the review.

Each data set entry links one entity’s aspect with one concept from a domain ontology (HOntology [4]). The annotation process collected explicit references to hotel’s aspects, the data set contains 987 explicit opinions about HOntology concepts, we use this entries as input to our model. Besides that, annotators also collected implicit and explicit references to TripAdvisor’s categories (Localization, Room, Service, Value, and Cleanliness). There are 269 explicit and 71 implicit mentions to these categories.

5.2 Baselines

ATAE-LSTM [18]. Our first baseline is an implementation of ATAE-LSTM from Wang et al. [18], detailed in Sect. 3.

This approach depends on word embedding. Authors claim that they have used pre-trained Glove word vectors, obtained from Stanford² repository.

¹ This work was developed using PyTorch 1.0.

² Pre-trained English Glove models can be obtained from <http://nlp.stanford.edu/projects/glove/>.

Glove models for Brazilian Portuguese are not available in Stanford’s repository. Then we adapted the code to use Glove word embedding trained to Brazilian Portuguese³ [9].

Results presented in this Section were obtained by executing authors code as is, except by the exchange of Glove pre-trained model.

Lexicon-Based ABSA Approach [7]. We also tried the lexicon-based method proposed by Freitas [7], the only approach for ABSA fully developed focusing only on Brazilian Portuguese, as far as we know.

This approach was proposed in the same work were the TripAdvisor’s Corpus was created. This is a knowledge-based model, fully based on linguistic rules to detect opinions and lexicons to obtain sentiment orientation.

Results presented here were obtained by modifying authors code to calculate overall accuracy. Freitas [7] code calculated other metrics per aspect (ontology’s concepts).

5.3 Results

This Section presents a performance comparison between baselines and CharCNN approaches. Table 1 summarizes accuracies obtained by the approaches tested. Both binary (positive/negative) and three-way (positive/neutral/negative) accuracies were assessed. As one can see, the best results are presented by the CharCNN-v2.

Accuracy values presented for ATAE-LSTM and AE-CharCNN were obtained from the mean of 10 executions. For each execution we randomly split the *corpus* in train and test sets, in a 70/30 proportion. Accuracy values in Table 1 were assessed on test set.

Table 1. Test accuracy

	Binary (+, -)	Three-way (+, Neutral, -)
Large AE-CharCNN-v1 (ours)	0.69 ± 0.02	0.54 ± 0.03
Small AE-CharCNN-v1 (ours)	0.68 ± 0.03	0.55 ± 0.03
Large AE-CharCNN-v2 (ours)	0.71 ± 0.03	0.55 ± 0.04
Small AE-CharCNN-v2 (ours)	0.71 ± 0.02	0.55 ± 0.02
ATAE-LSTM	0.69 ± 0.19	0.49 ± 0.13
Lexicon-Based*	0.66 ± 0.00	0.55 ± 0.00

*Freitas [7] approach is deterministic, resulting in no standard deviation

It is important to notice that the Lexicon-Based approach is a knowledge-based approach developed by the author of the *corpus* used in our experiments.

³ Pre-trained Brazilian Portuguese Glove models can be obtained from <http://nilc.icmc.usp.br/embeddings>.

Besides that, this method is highly dependent of external resources, such domain ontology, sentiment lexicon, and PoS-Tagger. These external resources demands a lot of effort to be created. Moreover, all external resources have an important impact on lexicon-based sentiment classifiers, as can be seen on the work of Freitas and Vieira [8].

6 Conclusion and Further Works

This work presented results about the application of character-level approaches to ABSA focused on Brazilian Portuguese. Results showed that data-driven approaches can perform better than knowledge-based (lexicon-based) approach, without the high dependence of additional external resources [8].

AE-CharCNN takes advantage of the character-level approach, avoiding complex pre-processing and making it easy to train a method tolerant to typos, rare or absent words during training, and encoding errors.

There are several possible paths for further works. For instance, to test the proposed approach with other *corpora* should bring more clarity about the performance difference between the work of Freitas [7] and ML approaches. Since there is a possibility that their results are biased by the rules created with those TripAdvisor's Reviews in mind.

We are also working on expanding TripAdvisor's data set. All reviews written in Portuguese were collected and are being annotated with the same protocol proposed by Freitas [7]. It should improve the performance of data-driven models.

Besides that, data-driven models are sensitive to data set where they were trained. Our experiments were done with hyperparameter proposed in earlier works, applied to other corpora. A grid search of hyperparameters can improve the AE-CharCNN performance over TripAdvisor's data set.

Acknowledgments. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan X GPU used for this research.

References

1. Blaz, C.C.A., Becker, K.: Sentiment analysis in tickets for it support. In: 2016 IEEE/ACM 13th Working Conference on Mining Software Repositories (MSR), pp. 235–246, May 2016. <https://doi.org/10.1109/MSR.2016.032>
2. Carvalho, P., Sarmiento, L., Teixeira, J., Silva, M.J.: Liars and saviors in a sentiment annotated corpus of comments to political debates. In: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers-Volume 2, pp. 564–568. Association for Computational Linguistics (2011)
3. Chaves, M.S., de Freitas, L.A., Souza, M., Vieira, R.: PIRPO: an algorithm to deal with polarity in Portuguese online reviews from the accommodation sector. In: Bouma, G., Ittoo, A., Métails, E., Wortmann, H. (eds.) NLDB 2012. LNCS,

- vol. 7337, pp. 296–301. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31178-9_37
4. Chaves, M.S., Freitas, L.A., Vieira, R.: Hontology: a multilingual ontology for the accommodation sector in the tourism industry. In: 4th International Conference on Knowledge Engineering and Ontology Development, pp. 149–154 (2012)
 5. Dos Santos, C., Gatti, M.: Deep convolutional neural networks for sentiment analysis of short texts. In: Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers, pp. 69–78 (2014)
 6. Fernandes, A.A., de Freitas, L.A., Corrêa, U.B.: Minerando tweets. In: Pesquisas e Perspectivas em Linguística de Corpus, pp. 283–302. Mercado de Letras (2015)
 7. Freitas, L.A.: Feature-level sentiment analysis applied to Brazilian Portuguese reviews. Ph.D. thesis, Programa de Pós-Graduação em Ciência da Computação. Pontifícia Universidade Católica do Rio Grande do Sul, Porto Alegre (2015)
 8. de Freitas, L.A., Vieira, R.: Exploring resources for sentiment analysis in portuguese language. In: 2015 Brazilian Conference on Intelligent Systems (BRACIS), pp. 152–156, November 2015. <https://doi.org/10.1109/BRACIS.2015.52>
 9. Hartmann, N., Fonseca, E., Shulby, C., Treviso, M., Rodrigues, J., Aluisio, S.: Portuguese word embeddings: evaluating on word analogies and natural language tasks. arXiv preprint [arXiv:1708.06025](https://arxiv.org/abs/1708.06025) (2017)
 10. Joulin, A., Grave, E., Bojanowski, P., Mikolov, T.: Bag of tricks for efficient text classification. In: Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers, pp. 427–431. Association for Computational Linguistics, April 2017
 11. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1, pp. 1097–1105, NIPS’12, Curran Associates Inc., USA (2012). <http://dl.acm.org/citation.cfm?id=2999134.2999257>
 12. Liu, B.: Sentiment Analysis: Mining Opinions, Sentiments, and Emotions. Cambridge University Press, Cambridge (2015)
 13. Pang, B., Lee, L., et al.: Opinion mining and sentiment analysis. *Found. Trends®Inf. Retrieval* **2**(1–2), 1–135 (2008)
 14. Severyn, A., Moschitti, A.: Twitter sentiment analysis with deep convolutional neural networks. In: Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 959–962. ACM (2015)
 15. Singhal, P., Bhattacharyya, P.: Borrow a little from your rich cousin: using embeddings and polarities of English words for multilingual sentiment classification. In: Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: technical papers, pp. 3053–3062 (2016)
 16. de Souza, J.G.R., de Paiva Oliveira, A., de Andrade, G.C., Moreira, A.: A deep learning approach for sentiment analysis applied to hotel’s reviews. In: Silberstein, M., Atigui, F., Kornysheva, E., Métais, E., Meziane, F. (eds.) NLDB 2018. LNCS, vol. 10859, pp. 48–56. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-91947-8_5
 17. Tang, D., Zhang, M.: Deep learning in sentiment analysis. In: Deng, L., Liu, Y. (eds.) Deep Learning in Natural Language Processing, pp. 219–253. Springer, Singapore (2018). https://doi.org/10.1007/978-981-10-5209-5_8
 18. Wang, Y., Huang, M., Zhao, L., et al.: Attention-based LSTM for aspect-level sentiment classification. In: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, pp. 606–615 (2016)

19. Wehrmann, J., Becker, W., Cagnini, H.E.L., Barros, R.C.: A character-based convolutional neural network for language-agnostic twitter sentiment analysis. In: 2017 International Joint Conference on Neural Networks (IJCNN), pp. 2384–2391, May 2017. <https://doi.org/10.1109/IJCNN.2017.7966145>
20. Wehrmann, J., Becker, W.E., Barros, R.C.: A multi-task neural network for multilingual sentiment classification and language detection on twitter. *Mach. Translation* **2**(32), 37 (2018)
21. Young, T., Hazarika, D., Poria, S., Cambria, E.: Recent trends in deep learning based natural language processing. CoRR abs/1708.02709 (2017). <http://arxiv.org/abs/1708.02709>
22. Zhang, X., LeCun, Y.: Text understanding from scratch. arXiv preprint [arXiv:1502.01710](https://arxiv.org/abs/1502.01710) (2015)
23. Zhang, X., Zhao, J., LeCun, Y.: Character-level convolutional networks for text classification. In: *Advances in Neural Information Processing Systems*, pp. 649–657 (2015)