

Clustering Word Embeddings with Self-Organizing Maps. Application on LaRoSeDa – A Large Romanian Sentiment Data Set

Anca Maria Tache^{1,*}, Mihaela Găman^{1,‡}, Radu Tudor Ionescu^{1,2,◊}

¹Department of Computer Science, ²Romanian Young Academy
University of Bucharest

14 Academiei, Bucharest, Romania

*anca.taches@gmail.com ‡mp.gaman@gmail.com

◊raducu.ionescu@gmail.com

Abstract

Romanian is one of the understudied languages in computational linguistics, with few resources available for the development of natural language processing tools. In this paper, we introduce LaRoSeDa, a **Large Romanian Sentiment Data Set**¹, which is composed of 15,000 positive and negative reviews collected from one of the largest Romanian e-commerce platforms. We employ two sentiment classification methods as baselines for our new data set, one based on low-level features (character n-grams) and one based on high-level features (bag-of-word-embeddings generated by clustering word embeddings with k-means). As an additional contribution, we replace the k-means clustering algorithm with self-organizing maps (SOMs), obtaining better results because the generated clusters of word embeddings are closer to the Zipf’s law distribution, which is known to govern natural language. We also demonstrate the generalization capacity of using SOMs for the clustering of word embeddings on another recently-introduced Romanian data set, for text categorization by topic.

1 Introduction

Perhaps one of the most studied tasks in computational linguistics is sentiment classification, a.k.a. opinion mining or polarity classification. The task has been studied across several languages, the most popular being English (Blitzer et al., 2007; Dos Santos and Gatti, 2014; Fu et al., 2018; Giménez-Pérez et al., 2017; Huang et al., 2017; Ionescu and Butnaru, 2019; Kim, 2014; Maas et al., 2011; Pang and Lee, 2005; Shen et al., 2018; Socher et al., 2013), Chinese (Peng et al., 2017; Wan, 2008; Zagibalov and Carroll, 2008; Zhai et al., 2011; Zhang and He, 2013; Zhang

et al., 2008), Arabic (Al-Ayyoub et al., 2018; El-sahar and El-Beltagy, 2015; Dahou et al., 2016; Nabil et al., 2013, 2015) or Spanish (Brooke et al., 2009; Molina-González et al., 2015; Navas-Loro and Rodríguez-Doncel, 2019; Vilares et al., 2015; Zafra et al., 2017). However, studying this task for under-studied languages, e.g. Romanian, is difficult without access to large data sets. We hereby introduce LaRoSeDa, a **Large Romanian Sentiment Data Set**, which is freely available for download at <https://github.com/ancatache/LaRoSeDa> for non-commercial use. With a total of 15,000 positive and negative reviews collected from one of the most popular Romanian e-commerce websites, to our knowledge, LaRoSeDa is the largest data set for Romanian polarity classification.

We experiment with two baseline methods on our novel data set. The first baseline employs string kernels, an approach based on low-level features (character n-grams), that was found to work well for sentiment analysis across multiple languages, e.g. English (Giménez-Pérez et al., 2017; Ionescu and Butnaru, 2018), Chinese (Zhang et al., 2008) and Arabic (Popescu et al., 2017), requiring no linguistic resources besides a labeled training set of samples. The second baseline employs bag-of-word-embeddings (Ionescu and Butnaru, 2019; Fu et al., 2018), an approach based on high-level features (clusters of word embeddings generated by k-means), that attains good results in various text classification tasks (Butnaru and Ionescu, 2017; Cozma et al., 2018; Fu et al., 2018; Ionescu and Butnaru, 2019), including sentiment analysis. As an additional contribution, we replace the k-means clustering algorithm in the second baseline method with self-organizing maps (SOMs) (Kohonen, 2001), obtaining better results because the generated clusters of word embeddings are closer to the Zipf’s law distribution (see Figure 2), which is known to govern natural language (Powers, 1998). To our knowl-

¹<https://github.com/ancatache/LaRoSeDa>

edge, we are the first to apply SOMs to cluster word embeddings, showing performance gains for both word2vec (Mikolov et al., 2013) and Romanian BERT (Dumitrescu et al., 2020) embeddings. We also demonstrate the generalization capacity of using SOMs in the bag-of-word-embeddings on another recently-introduced Romanian data set (Butnaru and Ionescu, 2019), for the task of text categorization by topic.

In summary, our contribution is twofold:

- We introduce LaRoSeDa, one of the largest corpora for Romanian sentiment analysis, along with a set of strong baselines to be used as reference in future research.
- To our knowledge, we are the first to employ SOMs as a technique to cluster word embeddings. We provide empirical evidence showing that SOMs produce better results than the popular k-means.

2 Related Work

To date, a small number of works targeting sentiment classification in the Romanian language have been published. Preceding the sentiment analysis efforts on Romanian texts, there are a few studies on subjectivity, that have introduced two corpora built through cross-lingual projections from English to Romanian (Mihalcea et al., 2007) or through machine translation (Banea et al., 2008). An extensive study conducted by Banea et al. (2011) looks at sentiment and subjectivity from a computational linguistics perspective, in a multilingual setup in which Romanian is also included. However, in these initial works, Romanian is studied only from a subjectivity perspective, which does not go down to the level of polarity.

On our topic (i.e. sentiment analysis in Romanian), the first study that we have found describes two word sets tagged with polarity for Romanian and Russian (Sokolova and Bobicev, 2009). Gînscă et al. (2011) introduced a sentiment analysis service intended for multiple languages, that also supports Romanian. They perform sentiment identification using a list of manually-built triggers which, to our knowledge, is not publicly available. Another effort (Colhon et al., 2016) in creating an opinion lexicon with polarity annotations introduced a collection of 2,521 Romanian tourist reviews and an extensive linguistic analysis of the corpus. The data set is not released for public use. Similarly, we did not find any public link to RoEmoLex, a lexicon with




Figure 1: The rating distribution of Romanian product reviews. Negative reviews are those rated with one or two stars, while positive reviews are those rated with four or five stars. Neutral reviews are not included in our data set. Best viewed in color.

approximately 11,000 Romanian words tagged for emotion and sentiment (Briciu and Lupea, 2017; Lupea and Briciu, 2017). The only Romanian data set annotated for sentiment that we have found freely available is rather small, with 1,000 movie reviews manually extracted from several blogs and sites (Russu et al., 2014). With 15,000 reviews, our corpus is much larger.

3 Data Set

In order to build LaRoSeDa, we collected product reviews from one of the largest e-commerce websites in Romania. Along with the textual content of each review, we collected the associated star ratings in order to automatically assign labels to the collected text samples. Following the same approach used for data sets containing English reviews (Blitzer et al., 2007; Maas et al., 2011; Pang and Lee, 2005), we assigned positive labels to the reviews rated with four or five stars and negative labels to the reviews rated with one or two stars. However, the star rating might not always reflect the polarity of the text. We thus acknowledge that the automatic labeling process is not optimal, i.e. some labels might be noisy. Since automatic labeling based on star ratings is a commonly accepted practice for opinion mining data sets of product reviews, we leave the analysis of noisy labels and manual labeling for future work.

We also imitate the data collection approach used for English review data sets (Blitzer et al., 2007; Maas et al., 2011; Pang and Lee, 2005), selecting a balanced set of Romanian reviews. More precisely, LaRoSeDa is formed of a total of 15,000 reviews that are perfectly balanced, i.e. half of them (7,500)

Set	Positive		Negative	
	#samples	#words	#samples	#words
Training	6,000	187,813	6,000	244,499
Test	1,500	47,661	1,500	60,314
Total	7,500	235,474	7,500	304,813

Table 1: The number of positive and negative samples (#samples) and the corresponding number of words (#words) contained in the training set and the test set of LaRoSeDa.

are positive reviews and the other half (7,500) are negative reviews. In Figure 1, we show the distribution of reviews with respect to the star ratings. We note that most of the negative reviews (5,561) are rated with one star. Similarly, most of the positive reviews (6,238) are rated with five stars. Hence, the corpus is highly polarized. We divide LaRoSeDa into a training set containing 80% of the data samples and a test set containing the remaining 20%. In Table 1, we present the number of positive and negative reviews inside each subset, along with the number of words. Our data set contains a total of 540,287 words, with an average of 36 words per review. We observe that positive reviews contain 235,474 words (44.6%) and negative reviews contain 304,813 words (56.4%). We note that, in negative reviews, people are likely to complain about several points or to explain what is wrong with the reviewed products. This could provide a natural explanation for the fact that the negative reviews contain more words than the positive reviews.

4 Methods

String kernels. A simple language-independent and linguistic-theory-neutral approach is to interpret text samples as sequences of characters (strings) and to use character n-grams as features. The number of character n-grams is usually much higher than the number of samples, so representing the text samples as feature vectors may require a lot of space. String kernels provide an efficient way to avoid storing and using the feature vectors (primal form), by representing the data through a kernel matrix (dual form). Each component K_{ij} in a kernel matrix represents the similarity between data samples x_i and x_j . In our experiments, we use the histogram intersection string kernel (HISK) (Ionescu et al., 2014, 2016) as the similarity function. For two strings x_i and x_j over a set of characters S , HISK is defined as follows:

$$k^\cap(x_i, x_j) = \sum_{g \in S^n} \min\{\#(x_i, g), \#(x_j, g)\}, \quad (1)$$

where $\#(x, g)$ is a function that returns the number of occurrences of n-gram g in x , and n is the length of n-grams. While being a rather shallow approach, string kernels attained strong results in some specific tasks. For instance, string kernels ranked first in the Arabic Dialect Identification tasks of VarDial 2017 (Ionescu and Butnaru, 2017) and VarDial 2018 (Butnaru and Ionescu, 2018).

Bag-of-word-embeddings. Following the seminal paper of Mikolov et al. (2013) introducing *word2vec*, word embeddings became one of the mainstream approaches in various computational linguistics tasks (Cheng et al., 2018; Conneau et al., 2017; Cozma et al., 2018; Fu et al., 2018; Kim, 2014; Kiros et al., 2015; Shen et al., 2018; Torki, 2018; Zhou et al., 2018). In order to build the bag-of-word-embeddings (BOWE), we first trained *word2vec* on the collected Romanian reviews using the continuous bag-of-words (CBOW) model. Before training, we transformed all letters to lowercase and removed punctuation. In addition to *word2vec*, we consider a recently introduced Romanian BERT model (Dumitrescu et al., 2020) as an alternative way to produce word embeddings, which is likely to produce much better results, considering the success of BERT (Devlin et al., 2019) in English NLP tasks. Instead of averaging the word embeddings to obtain document-level representations (Shen et al., 2018), we follow a different and more effective path suggested by some recent works (Butnaru and Ionescu, 2017; Cozma et al., 2018; Fu et al., 2018; Ionescu and Butnaru, 2019). More specifically, we cluster the word embeddings collected from the entire training set using k-means. For a document D of n words, $D = (w_1, w_2, \dots, w_n)$, a word embedding model, be it *word2vec* or BERT, outputs a matrix of $n \times m$ components (or a set of n m -dimensional vectors), the m -dimensional vector at index i corresponding to word w_i . We apply clustering on the word vectors extracted from all training documents, thus obtaining a set of k clusters. A document D is then represented as a bag-of-word-embeddings

(histogram) $H = (h_1, h_2, \dots, h_k)$ in which each component h_i retains the number of word embeddings from the document D that fall in cluster i , where $i \in \{1, 2, \dots, k\}$. We note that the size of the bag-of-word-embeddings is equal to the number of clusters k . In the case of BERT, we emphasize that, although the embedding vector of a word depends on the context, it is likely that the embedding vectors corresponding to a specific word will fall in the same cluster. Hence, BOWE is able to cope well with this situation.

Replacing k-means with SOMs. Quantitative linguistics studies (Powers, 1998) have pointed out that, given a corpus of text documents, the frequency of any word is inversely proportional to its rank in the frequency table, giving rise to a Zipf’s law distribution of words in natural language. However, the k-means algorithm tends to ignore the data density, producing equally-sized clusters. We therefore propose to replace the k-means algorithm with an approach that takes into account the density in the word embedding space, producing a set of clusters that follow the Zipf’s law. We propose to perform clustering using self-organizing maps (SOMs) (Kohonen, 2001), since these models are known to preserve the topological properties of the input space. Indeed, Figure 2 shows that SOMs produce clusters of Romanian word embeddings closer to the Zipf’s law distribution than k-means.

It is important to emphasize that k-means can produce clusters of different size, as shown in Figure 2. Our observation refers only to the fact that the data density is not particularly modeled by the k-means optimization process, while SOMs are optimized by shifting the neural units following the density of the data (units tend to migrate where the space is more dense). Our observation with respect to k-means is also confirmed by other studies. For example, Raykov et al. (2016) note that: “*even when all other implicit geometric assumptions of k-means are satisfied, it will fail to learn a correct, or even meaningful, clustering when there are significant differences in cluster density*”. Since natural language involves such significant differences (due to the presence of the Zipf’s law), we believe that k-means is a sub-optimal choice.

5 Experiments

Corpora. First and foremost, we perform experiments on LaRoSeDa with the goal of introducing some benchmark results on our new data set. We




Figure 2: The distribution of words into clusters generated by k-means (in red) or by SOMs (in blue), for LaRoSeDa. The Zipf’s law distribution (in green) is included for reference. Best viewed in color.

also perform experiments on MOROCO (Butnaru and Ionescu, 2019), a data set with Moldavian and Romanian news articles, with the goal of showing the generalization capacity of using SOMs instead of k-means.

Experimental setup. On LaRoSeDa, we present two sets of results, one based on the established train-test split and one based on 10-fold cross-validation. On MOROCO, we choose to present 10-fold cross-validation results for the intra-dialect multi-class categorization by topic task, on the 18,161 samples written in the Romanian dialect.

Parameter and model choices. For HISK, we combined character 3-grams, 4-grams and 5-grams. For BOWE-word2vec and BOWE-BERT, we set the number of clusters to $k = 500$, just as Cozma et al. (2018). We trained *word2vec* to produce 300-dimensional Romanian word embeddings, while the Romanian BERT outputs 768-dimensional embeddings. In the learning stage, we employed the linear Support Vector Machines (SVM) implementation from Scikit-learn (Pedregosa et al., 2011), providing as input pre-computed kernels. For BOWE-word2vec and BOWE-BERT, we opt for the PQ kernel, based on the findings of Butnaru and Ionescu (2017). We set the regularization parameter of SVM to $C = 10^3$ in all the experiments. We also fuse HISK with BOWE-word2vec or BOWE-BERT in the dual form by summing up the corresponding kernel matrices. We employed an open source implementation of SOMs². We used the default choices for most hyperparameters, the modifications being detailed next. We set the learning rate to 0.25 and the number of epochs to 200. Before starting the training, the SOM is configured to randomly choose a number of training samples equal to number of expected outputs. We

²<http://neupy.com/pages/home.html>

Method	10-fold CV	Test
HISK	84.38%	84.73%
BOWE-word2vec with k-means	72.24%	70.73%
BOWE-word2vec with SOMs	75.99%	75.57%
BOWE-BERT with k-means	78.63%	77.36%
BOWE-BERT with SOMs	79.75%	80.73%
HISK+BOWE-word2vec with k-means	85.08%	83.57%
HISK+BOWE-word2vec with SOMs	87.17%	88.93%
HISK+BOWE-BERT with k-means	88.81%	89.42%
HISK+BOWE-BERT with SOMs	89.54%	90.90%

Table 2: Accuracy rates of HISK, BOWE-word2vec and BOWE-BERT with clustering based on k-means or SOMs, as well as ensemble models on LaRoSeDa. Results are reported in two cases: using a 10-fold cross-validation procedure and using the train-test split.

Method	10-fold CV
HISK (Butnaru and Ionescu, 2019)	71.27%
BOWE-BERT with k-means	63.42%
BOWE-BERT with SOMs	68.50%
HISK+BOWE-BERT with k-means	72.21%
HISK+BOWE-BERT with SOMs	73.35%

Table 3: Accuracy rates of HISK, BOWE-BERT with k-means, BOWE-BERT with SOMs and their combinations on MOROCO, for the Romanian intra-dialect multi-class categorization by topic task. Results are reported using a 10-fold cross-validation procedure.

opted for the cosine distance between data samples and SOM’s weights.

Results on LaRoSeDa. In Table 2, we present the results on LaRoSeDa. Among the individual baselines, we observe that HISK attains the best accuracy rates, surpassing all BOWE configurations. We also note that by replacing k-means with SOMs, the accuracy rate of BOWE-word2vec grows by 4 or 5%. The improvements brought by SOMs can be explained by the fact that, unlike k-means, SOMs produce clusters that are closer to the Zipf’s law distribution. This is proven by the word embedding counts per cluster illustrated in Figure 2. When we combine HISK with BOWE-BERT, we notice significant performance gains.

Results on MOROCO. In Table 3, we present the results on MOROCO, for the Romanian intra-dialect multi-class categorization by topic task. We notice that HISK attains better results than BOWE-BERT with k-means and BOWE-BERT with SOMs, although the differences in terms of accuracy seem to be smaller. As for LaRoSeDa, we observe a significant improvement (higher than 5%) when k-means is replaced by SOMs. There is an observable improvement over the plain HISK, when

HISK is combined with BOWE-BERT based on k-means. Nonetheless, we notice a larger improvement when we combine HISK and BOWE-BERT based on SOMs.

6 Conclusion

In this paper, (i) we introduced LaRoSeDa, a large data set for polarity classification of Romanian reviews, and (ii) we employed self-organizing maps, a clustering approach that preserves the density of words in the embedding space, resulting in a more effective bag-of-word-embeddings representation. Our top accuracy rates on LaRoSeDa are 89.54% for the cross-validation procedure and 90.90% on the test set. We note that SOMs had a significant contribution in attaining these high accuracy rates. We conclude that the combination of HISK and BOWE-BERT with SOMs is a strong baseline which should encourage future research in proposing non-trivial models for Romanian polarity classification. Furthermore, the results obtained on MOROCO confirm that SOMs provide better accuracy rates than k-means when it comes to building document-level representations based on clustering word embeddings.

Acknowledgments

The authors thank reviewers for their useful remarks. This work was supported by a grant of the Romanian Ministry of Education and Research, CNCS - UEFISCDI, project number PN-III-P1-1.1-TE-2019-0235, within PNCDI III. This article has also benefited from the support of the Romanian Young Academy, which is funded by Stiftung Mercator and the Alexander von Humboldt Foundation for the period 2020-2022.

References

- Mahmoud Al-Ayyoub, Aya Nuseir, Kholoud Alsmearat, Yaser Jararweh, and Brij Gupta. 2018. [Deep learning for Arabic NLP: A survey](#). *Journal of Computational Science*, 26:522–531.
- Carmen Banea, Rada Mihalcea, and Janyce Wiebe. 2011. [Multilingual Sentiment and Subjectivity Analysis](#). In *Multilingual Natural Language Processing Applications: From Theory to Practice*, chapter 7. IBM Press.
- Carmen Banea, Rada Mihalcea, Janyce Wiebe, and Samer Hassan. 2008. [Multilingual Subjectivity Analysis Using Machine Translation](#). In *Proceedings of EMNLP*, page 127–135.
- John Blitzer, Mark Dredze, and Fernando Pereira. 2007. [Biographies, Bollywood, Boom-boxes and Blenders: Domain Adaptation for Sentiment Classification](#). In *Proceedings of ACL*, pages 187–205.
- Anamaria Briciu and Mihaela Lupea. 2017. [RoEmoLex – A Romanian Emotion Lexicon](#). *Studia Universitatis Babeş-Bolyai Informatica*, 62(2):45–56.
- Julian Brooke, Milan Tofiloski, and Maite Taboada. 2009. [Cross-linguistic sentiment analysis: From English to Spanish](#). In *Proceedings of RANLP*, pages 50–54.
- Andrei Butnaru and Radu Tudor Ionescu. 2017. [From Image to Text Classification: A Novel Approach based on Clustering Word Embeddings](#). In *Proceedings of KES*, pages 1784–1793.
- Andrei Butnaru and Radu Tudor Ionescu. 2019. [MO-ROCO: The Moldavian and Romanian Dialectal Corpus](#). In *Proceedings of ACL*, pages 688–698.
- Andrei M. Butnaru and Radu Tudor Ionescu. 2018. [UnibucKernel Reloaded: First Place in Arabic Dialect Identification for the Second Year in a Row](#). In *Proceedings of VarDial*, pages 77–87.
- Zhou Cheng, Chun Yuan, Jiancheng Li, and Haiqin Yang. 2018. [TreeNet: Learning Sentence Representations with Unconstrained Tree Structure](#). In *Proceedings of IJCAI*, pages 4005–4011.
- Mihaela Colhon, Mădălina Cerban, Alex Becheru, and Mirela Teodorescu. 2016. [Polarity shifting for Romanian sentiment classification](#). In *Proceedings of INISTA*, pages 1–6.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised Learning of Universal Sentence Representations from Natural Language Inference Data](#). In *Proceedings of EMNLP*, pages 670–680.
- Mădălina Cozma, Andrei Butnaru, and Radu Tudor Ionescu. 2018. [Automated essay scoring with string kernels and word embeddings](#). In *Proceedings of ACL*, pages 503–509.
- Abdelghani Dahou, Shengwu Xiong, Junwei Zhou, Mohamed Houcine Haddoud, and Pengfei Duan. 2016. [Word embeddings and convolutional neural network for Arabic sentiment classification](#). In *Proceedings of COLING*, pages 2418–2427.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). In *Proceedings of NAACL*, pages 4171–4186.
- Cícero Nogueira Dos Santos and Maira Gatti. 2014. [Deep Convolutional Neural Networks for Sentiment Analysis of Short Texts](#). In *Proceedings of COLING*, pages 69–78.
- Ştefan Daniel Dumitrescu, Andrei-Marius Avram, and Sampo Pyysalo. 2020. [The birth of Romanian BERT](#). In *Findings of EMNLP*, pages 4324–4328.
- Hady Elsahar and Samhaa R. El-Beltagy. 2015. [Building large Arabic multi-domain resources for sentiment analysis](#). In *Proceedings of CICLing*, pages 23–34.
- Mingsheng Fu, Hong Qu, Li Huang, and Li Lu. 2018. [Bag of meta-words: A novel method to represent document for the sentiment classification](#). *Expert Systems with Applications*, 113:33–43.
- Rosa M. Giménez-Pérez, Marc Franco-Salvador, and Paolo Rosso. 2017. [Single and Cross-domain Polarity Classification using String Kernels](#). In *Proceedings of EACL*, pages 558–563.
- Alexandru-Lucian Gînscă, Emanuela Boros, Adrian Iftene, Diana Trandabăt, Mihai Toader, Marius Corîci, Cene-Augusto Perez, and Dan Cristea. 2011. [Sentimatrix – Multilingual Sentiment Analysis Service](#). In *Proceedings of WASSA*, pages 189–195.
- Xingchang Huang, Yanghui Rao, Haoran Xie, Tak-Lam Wong, and Fu Lee Wang. 2017. [Cross-Domain Sentiment Classification via Topic-Related TrAdaBoost](#). In *Proceedings of AAAI*, pages 4939–4940.
- Radu Tudor Ionescu and Andrei Butnaru. 2019. [Vector of Locally-Aggregated Word Embeddings \(VLAWE\): A Novel Document-level Representation](#). In *Proceedings of NAACL*, pages 363–369.
- Radu Tudor Ionescu and Andrei M. Butnaru. 2017. [Learning to Identify Arabic and German Dialects using Multiple Kernels](#). In *Proceedings of VarDial*, pages 200–209.
- Radu Tudor Ionescu and Andrei M. Butnaru. 2018. [Improving the results of string kernels in sentiment analysis and Arabic dialect identification by adapting them to your test set](#). In *Proceedings of EMNLP*, pages 1084–1090.
- Radu Tudor Ionescu, Marius Popescu, and Aoife Cahill. 2014. [Can characters reveal your native language?](#)

- A language-independent approach to native language identification. In *Proceedings of EMNLP*, pages 1363–1373.
- Radu Tudor Ionescu, Marius Popescu, and Aoife Cahill. 2016. *String Kernels for Native Language Identification: Insights from Behind the Curtains*. *Computational Linguistics*, 42(3):491–525.
- Yoon Kim. 2014. *Convolutional Neural Networks for Sentence Classification*. In *Proceedings of EMNLP*, pages 1746–1751.
- Ryan Kiros, Yukun Zhu, Ruslan R. Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. *Skip-Thought Vectors*. In *Proceedings of NIPS*, pages 3294–3302.
- Teuvo Kohonen. 2001. *Self-Organizing Maps*, 3rd edition. Springer-Verlag.
- Mihaiela Lupea and Anamaria Briciu. 2017. *Formal concept analysis of a Romanian emotion lexicon*. In *Proceedings of ICCP*, pages 111–118.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. *Learning Word Vectors for Sentiment Analysis*. In *Proceedings of ACL*, pages 142–150.
- Rada Mihalcea, Carmen Banea, and Janyce Wiebe. 2007. *Learning Multilingual Subjective Language via Cross-Lingual Projections*. In *Proceedings of ACL*, pages 976–983.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. *Distributed Representations of Words and Phrases and their Compositionality*. In *Proceedings of NIPS*, pages 3111–3119.
- Dolores M. Molina-González, Eugenio Martínez-Cámara, Teresa M. Martín-Valdivia, and Alfonso L. Ureña-López. 2015. *A Spanish semantic orientation approach to domain adaptation for polarity classification*. *Information Processing & Management*, 51(4):520–531.
- Mahmoud Nabil, Mohamed Aly, and Amir Atiya. 2015. *ASTD: Arabic Sentiment Tweets Dataset*. In *Proceedings of EMNLP*, pages 2515–2519.
- Mahmoud Nabil, Mohamed A. Aly, and Amir F. Atiya. 2013. *LABR: A Large Scale Arabic Book Reviews Dataset*. In *Proceedings of ACL*, pages 494–498.
- María Navas-Loro and Víctor Rodríguez-Doncel. 2019. *Spanish corpora for sentiment analysis: a survey*. *Language Resources and Evaluation*, pages 1–38.
- Bo Pang and Lillian Lee. 2005. *Seeing Stars: Exploiting Class Relationships For Sentiment Categorization With Respect To Rating Scales*. In *Proceedings of ACL*, pages 115–124.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. *Scikit-learn: Machine Learning in Python*. *Journal of Machine Learning Research*, 12(85):2825–2830.
- Haiyun Peng, Erik Cambria, and Amir Hussain. 2017. *A Review of Sentiment Analysis Research in Chinese Language*. *Cognitive Computation*, 9(4):423–435.
- Marius Popescu, Cristian Grozea, and Radu Tudor Ionescu. 2017. *HASKER: An efficient algorithm for string kernels. Application to polarity classification in various languages*. In *Proceedings of KES*, pages 1755–1763.
- David M.W. Powers. 1998. *Applications and Explanations of Zipf’s Law*. In *Proceedings of CoNLL*, pages 151–160.
- Yordan P. Raykov, Alexis Boukouvalas, Fahd Baig, and Max A. Little. 2016. *What to Do When K-Means Clustering Fails: A Simple yet Principled Alternative Algorithm*. *PLoS One*, 11(9):1–28.
- Roxana Monica Russu, Mihaela Dinsoreanu, Oana Luminița Vlad, and Rodica Potolea. 2014. *An opinion mining approach for Romanian language*. In *Proceedings of ICCP*, pages 43–46.
- Dinghan Shen, Guoyin Wang, Wenlin Wang, Martin Renqiang Min, Qinliang Su, Yizhe Zhang, Chunyuan Li, Ricardo Henao, and Lawrence Carin. 2018. *Baseline Needs More Love: On Simple Word-Embedding-Based Models and Associated Pooling Mechanisms*. In *Proceedings of ACL*, pages 440–450.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, D. Christopher Manning, Andrew Ng, and Christopher Potts. 2013. *Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank*. In *Proceedings of EMNLP*, pages 1631–1642.
- Marina Sokolova and Victoria Bobicev. 2009. *Classification of Emotion Words in Russian and Romanian Languages*. In *Proceedings of RANLP*, pages 416–420.
- Marwan Torki. 2018. *A Document Descriptor using Covariance of Word Vectors*. In *Proceedings of ACL*, pages 527–532.
- David Vilares, Miguel A. Alonso, and Carlos Gómez-Rodríguez. 2015. *A syntactic approach for opinion mining on Spanish reviews*. *Natural Language Engineering*, 21(1):139–163.
- Xiaojun Wan. 2008. *Using Bilingual Knowledge and Ensemble Techniques for Unsupervised Chinese Sentiment Analysis*. In *Proceedings of EMNLP*, pages 553–561.

- Salud Maria Jimenez Zafra, Teresa M. Martín-Valdivia, Eugenio Martinez Camara, and Alfonso L. Ureña-López. 2017. [Studying the scope of negation for Spanish sentiment analysis on Twitter](#). *IEEE Transactions on Affective Computing*, 10(1):129–141.
- Taras Zagibalov and John Carroll. 2008. [Unsupervised classification of sentiment and objectivity in Chinese text](#). In *Proceedings of IJCNLP*, pages 304–311.
- Zhongwu Zhai, Hua Xu, Bada Kang, and Peifa Jia. 2011. [Exploiting effective features for Chinese sentiment classification](#). *Expert Systems with Applications*, 38(8):9139–9146.
- Changli Zhang, Wanli Zuo, Tao Peng, and Fengling He. 2008. [Sentiment Classification for Chinese Reviews Using Machine Learning Methods Based on String Kernel](#). In *Proceedings of ICCIT*, volume 2, pages 909–914.
- Pu Zhang and Zhongshi He. 2013. [A weakly supervised approach to Chinese sentiment classification using partitioned self-training](#). *Journal of Information Science*, 39(6):815–831.
- Qianrong Zhou, Xiaojie Wang, and Xuan Dong. 2018. [Differentiated Attentive Representation Learning for Sentence Classification](#). In *Proceedings of IJCAI*, pages 4630–4636.