Tweet sentiment visualization and classification using manifold dimensionality reduction

Francisco GRIMALDO [a], Emilia LÓPEZ-IÑESTA [b], Juan José GARCÉS-INIESTA [c], Juan GÓMEZ-SANCHIS [c], Daniel GARCÍA-COSTA [a], Eusebio MARQUÉS-BENÍTEZ [a], and Emilio SORIA-OLIVAS [c]

[a] *Departament d'Informàtica, Universitat de València*
Av. de la Universitat s/n, 46100-Burjassot (Spain)
[b] *Facultat de Ciències de l'Activitat i l'Esport, Universitat Catòlica de València*
Av. Virgen de la Soledad s/n, 46900-Torrent (Spain)
[c] *Departament d'Enginyeria Electrónica, Universitat de València*
Av. de la Universitat s/n, 46100-Burjassot (Spain)
francisco.grimaldo@uv.es, emilia.lopez@ucv.es
juangarces@uv.es, juan.gomez-sanchis@uv.es
daniel.garcia@uv.es, eumarbe@alumni.uv.es
emilio.soria@uv.es

Abstract. The growth of online content generated by users on the Web and social media has turned Sentiment Analysis into one of the most active research areas of Natural Language Processing aimed to computationally identify the underlying opinions, attitudes or feelings of a given text. In particular, a variety of methods have emerged that focus on Twitter due to its challenging traits, being a central issue how to provide a proper feature-based representation for its short and incomplete messages, i.e. tweets. The contribution of this paper is twofold: on the one hand, we present an approach to perform sentiment tagging of tweets based on text and emoji polarity; on the other hand, a set of manifold dimensionality reductions are carried out that allow a convenient 3D visualization and a rapid prototyping of sentiment patterns. We then compare the performance obtained with such reduced feature spaces when applied to the classification of sentiments of a collection of tweets from a real-life case study of people experiencing the celebration of a massive festivity.

Keywords. Sentiment Analysis, Manifold, Visualization, Classification, Natural Language Processing, Social media, Twitter

1. Introduction

In the last decades technological advances have allowed the transmission and collection of data in an increasingly simple and efficient way. As a result, there has been a paradigm shift in the role of people, who have changed from being data consumers to become data producers of the ever-growing content available on the Web and Social Media (e.g. Twitter, Facebook, Instagram, StackOverflow, TripAdvisor, etc.). In particular, user-generated content contributes every day to the creation of large textual databases full of opinions that provide an excellent context to perform Sentiment Analysis (SA), a subcategory of Natural Language Processing (NLP) that makes it possible to detect emotions and feelings [24,23].

Sentiment Analysis techniques are widely used in a broad range of applications such as marketing [21], politics [17] or tourism [12]. For instance, SA has been succesfully applied in Customer Management Relationship services to deal with user suggestions and complaints in social networks or *chatbots* [7,5]. Other examples can be found in the measurement of the impact of marketing campaigns [8], the detection of social trends in large events like music festivals [15] or the design of recommendation systems [3], to name a few.

Processing the vast amount of textual data coming from Social Media is a big deal for NLP methods. In this regard, the popular micro-blogging platform Twitter poses a special challenge for SA techniques due to the characteristics of its 280 character limited messages, known as tweets. Some of these traits are related with the frequent use of abbreviations, low-quality grammar, multilingual, slang or mispelling language [2]. Besides that, the fact that messages can include enriched information (e.g. *hashtags*, mentions to user profiles, *retweeted* content and *emojis*) increases the difficulty of carrying out such a linguistic analysis.

One of the goals when analyzing tweets is that of automatically identifying the sentiment behind the text and, more specifically, the ability to classify them visually. Graphical representations facilitate the early detection of (un)desired emotions and open up the possibility of a fast reaction by stakeholders. A number of SA methods have emerged that focus on how to provide a proper feature-based representation for tweets [34,11], though, finding a low-dimensional viewable space of representation may go against its expressive power.

Thus, the contribution of this paper is twofold: on the one hand, we present an approach to perform sentiment tagging of tweets based on text and emoji polarity; on the other hand, a set of manifold dimensionality reductions are carried out that allow a convenient 3D visualization and a rapid prototyping of sentiment patterns. We then compare the performance obtained with such reduced feature spaces when applied to the classification of sentiments of a collection of tweets from a real-life case study of people experiencing the celebration of a massive festivity.

The rest of the paper is organized as follows: Section 2 describes related work on Sentiment Analysis and manifold dimensionality reduction. Section 3 presents our approach to tag tweets as positive or negative and to plot them in a 3D space. The experimental setting and the obtained results are analyzed in Section 4. Finally, Section 5 states the conclusions and discusses future work.

2. Related work

Natural Language Processing has a long-standing tradition in artificial intelligence research but it was not till the year 2000 that digging into people's opinion and sentiments attracted most attention [17]. Indeed, seminal work on sentiment classification can be dated as of 2002, when two approaches were introduced that used alternative machine learning methods to discover the semantic orientation of a text. On the one hand, Pang, Lee and Vaithyanathan [24] proposed the use of supervised models such as Naive Bayes, Maximum Entropy and Support Vector Machines to determine the positive or negative sentiment of movie reviews. On the other hand, Turney [31] presented a lexicon-based method to assess polarity, i.e. a real-number measuring the positive, negative or neutral sentiment behind a text. This approach involved the use of an unsupervised algorithm that employed a dictionary of words and a corpus of phrases with associated polarity scores. A couple of examples of prominent lexicons for opinion mining are SentiWordNet and SentiStrength[1].

The different methods in the Sentiment Analysis literature are then characterised by the learning approach followed (supervised, unsupervised or hybrid), the level of granularity at which SA is done (mainly document, sentence or entity-aspect level) and the great amount of related tasks involved such as polarity classification, sentiment extraction, and opinion summarization, among others [17]. Although polarities of most opinion words can be used across domains, the performance of SA is usually context-dependent and a universal sentiment lexicon able to capture contextual polarities of words, especially when dealing with short texts, is still to arrive.

Meanwhile, applying SA to Twitter data has become a hot topic of research in the last years. The publication of a number of complete surveys [1,22] together with the organization of competitions about sentiment analysis, classification or prediction of tweets in international conferences[2] give an idea of this increasing interest. Twitter data poses a difficult challenge to NLP techniques due to the special characteristics of tweets and make formal mathematical representation of tweets particularly important to detect sentiments and opinions.

To face this challenge, a series of pre-processing methods have been developed to build proper feature-based representations of tweets [18,33]. Traditional representations have used syntactic and linguistic hand-crafted features such as Bag-Of-Words (BOW), Part-Of-Speech (POS) tags, n-grams or word counts[17]. An alternative representation consists in modelling tweets using a word embedding function that maps words to a feature vector space (e.g. a Neural Network [20]), learned from a large collection of texts. Word embeddings have shown promising results in NLP tasks [26,13,19], capturing the semantic and syntactic relationships of similar words (as measured by cosine vector similarity).

Rich feature-based representation methods may end up in high-dimensional spaces that can harden sentiment classification and visualization. Hence, dimension reduction methods are commonly used to project tweet representations in a low-dimensional space where the structure of the original data is preserved. There are multiple dimension reduction methods, linear and non-linear, that can be applied such as: Principal Component Analysis (PCA) [25], Multidimensional Scaling (MDS) [10] or Laplacian Eigenmaps (LE) [4]. Recent research has focused on manifold techniques [29] to enhance the understanding of data structures through data visualization and to facilitate classification and clustering [30,28].

[1] Available, respectively, at http://sentiwordnet.isti.cnr.it/ and http://sentistrength.wlv.ac.uk/

[2] See, for instance, the International Workshop on Semantic Evaluation (*SemEval*, http://alt.qcri.org/semeval2018) or the International Conference of the Spanish Society for Natural Language Processing (*Sepln*, http://www.sepln.org)

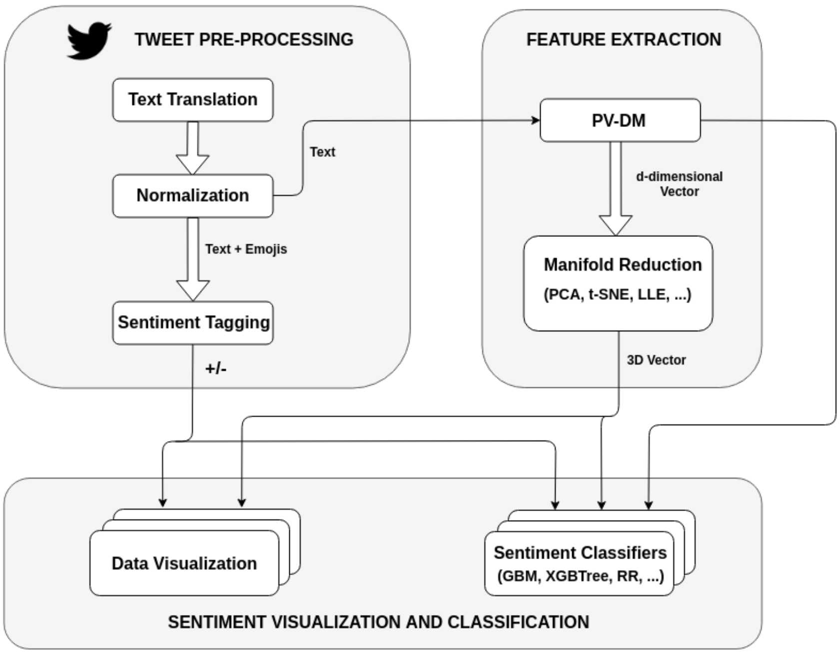

Figure 1. System architecture for sentiment visualization and classification of tweets.

In this paper, we propose a system that combines word embedding and dimensionality reduction to perform sentiment visualization and classification of tweets tagged by aggregating text and emoji polarity. A set of manifold dimensionality reductions is carried out that allows a convenient 3D visualization that can be useful in real applications to conduct a rapid prototyping of sentiment patterns and fasten decision-making.

3. System description

In this section we propose a system able to visualize and classify tweets based on polarity. The architecture of the system is depicted in Figure 1, where data goes through three stages, namely: tweet pre-processing, feature extraction and, sentiment visualization and classification.

3.1. Tweet pre-processing

The first task carried out aims to generate a clean dataset of tweets by following some commonly used pre-processing steps and a sentiment tagging:

- **Translation**: all tweets are translated into English to take advantage of more complete English lexicons and off-the-shelf NLP libraries. This is done in Python using the Google Translate REST API[3].

[3] Available at https://cloud.google.com/translate/

- **Normalization**: we remove duplicates (e.g. retweets), URLs, hashtags and user mentions (i.e. words starting with @). We keep unicode strings (i.e. emojis) separately and the remaining text is converted to lowercase.
- **Sentiment tagging**: we perform a tweet-level sentiment classification based on text and emoji polarity. On the one hand, we calculate the polarity score of texts by summing the sentence-level values obtained from the R library `sentimentr`. This library performs a quick augmented dictionary lookup that also takes into account valence shifters (i.e., negators, amplifiers (intensifiers), de-amplifiers (downtoners), and adversative conjunctions). On the other hand, we add the sentiment score of each emoji in the tweet by using the Emoji Sentiment Ranking [14]. The combined polarity score is used to tag each tweet as negative $(-\infty, threshold_{neg}]$ or positive $[threshold_{pos}, +\infty)$, where $threshold_{neg}$ and $threshold_{pos}$ can be set accordingly depending on the domain or the data at hand.

3.2. Feature extraction: word embedding and manifold reduction

After the pre-processing stage, we extract a fixed-length feature representation for each tweet. We use as word embedding function the Paragraph Vector with Distributed Memory model (PV-DM) [16], which is an extension the the word2vec algorithm [20]. PV-DM is an unsupervised framework able to capture the semantics of words while also taking into consideration the word order in small contexts that are learned from variable-length pieces of texts. We apply the PV-DM model in two steps: 1) we train the model using as corpus the clean tweets; and 2) we infer a mapping of each tweet to a d dimension real-valued vector, where d usually ranges from tens to hundreds.

Thus, we perform a set of dimentionality reductions of the feature space resulting from the PV-DM model to a three-dimensional space that will be later used for sentiment visualization. In this paper, we apply a pool of manifold learning methods that may provide interesting graphic capabilities. The following methods have been included from the `scikit-learn` Python library:

- Principal Component Analysis (PCA) [25] is a statistical procedure where the original data is rotated using an orthogonal set of axes that maximally preserves variance in the dataset.
- Multi-dimensional scaling (MDS) [10] seeks a low-dimensional representation of the data in which the distances in the original high-dimensional space are retained.
- Isometric Mapping (ISOMAP) [29] can be viewed as an extension of MDS where the geodesic distances between all points are preserved.
- t-Distributed Stochastic Neighbour Embedding (t-SNE) [32] represents the similarities between data points in the original high-dimensional space by Gaussian joint probabilities and the similarities in the embedded space by Student's t-distributions allowing the technique to be particularly sensitive to local structure.
- Locally linear embedding (LLE) [27] searchs for a lower-dimensional data projection that maintais the local geometry i.e., distances within local neighborhoods.
- The modified LLE (MLLE) [35] uses multiple weight vectors for each point in reconstruction of lower-dimensional embedding when solving the regularization problem of LLE.
- Totally Random Trees (TRT) [9] implement an unsupervised transformation using a forest of completely random trees. A datapoint is coded according to which leaf

of each tree it is sorted into and we use PCA to further reduce the space to three components.
- Spectral Embedding [6] calcutes a non-linear embedding using a spectral decomposition of a generated graph Laplacian, which can be considered as a discrete approximation of the low-dimensional manifold in the high-dimensional space.

3.3. Sentiment visualization and classification

In this final stage, we plot the 3D points describing tweets in each reduced space taking as labels the sentiment tags obtained in the pre-processing stage. We claim that having several points of view of a dataset of tweets can be a robust way to conduct a rapid prototyping of sentiment patterns when analyzing different collections of tweets. Each manifold technique projects tweets in a different space, where the sentiment of a message can be visually compared to that of its spatial neighbours. So while exploring these plots, outliers can eventually be scrutinized by an interested observer.

Moreover, the reduced feature vectors can be used to classify the sentiment of new incoming tweets both graphically and by training a classification technique. In particular, we evaluate the predictive power of this reduced data input format by measuring the performance of six well-known classification methods from the `caret` and `e1071` R libraries: Gradient Boosting Method (GBM), eXtreme Gradient Boosting (XGBTree), Random Forest (RF), Averaged Neural Network (AVNNet), Naive Bayes (NB), and Support Vector Machine (SVM).

4. Experimental evaluation

To evaluate our approach, we have applied the proposed sentiment visualization and classification to a collection of tweets from a real-life case study of people experiencing the Fallas Festivity, a massive celebration that takes place in València each March. We retrieved from Twitter the set of tweets including one of the 31 hashtags commonly used during the Fallas Festivity 2018 (e.g. `#fallas`, `#fallas18`, `#fallas2018`, `#fallasunesco`, etc.).

Tweets were originally written in 37 different languages and, after the pre-processing stage, we chose randomly a balanced corpus of 12970 text messages tagged as positive and negative. We set $threshold_{pos} = 0.5$ and $threshold_{neg} = -0.2$ as parameters for the sentiment tagging and we extracted $d = 20$ features from the PV-DM model. PV-DM feature extraction took 181 seconds on an Intel i5 processor (6600K, 3.50 GHz) with 32Gb of RAM and a 240Gb SSD. We set the number of neighbours to 15 for the manifold methods LLE, MLLE, ISOMAP and Spectral Embedding.

Figure 2 shows the sentiment visualization that is obtained for each dimensionality reduction method. The numbers in parentheses indicate the seconds that are necessari to run the transformation. We can observe how visual classification of positive and negative tweets is mainly achieved when using PCA, ISOMAP, Spectral Embedding and t-SNE. Other manifold learning methods in the pool such as TRT, MDS, LLE and MLLE do not provide a clear spatial class separation for this dataset. Therefore, we just selected the 3D vector spaces of PCA, ISOMAP, Spectral Embedding and t-SNE to compare the predictive power of these viewable vector representations with that of the higher-dimensional feature representation initially obtained from the PV-DM method.

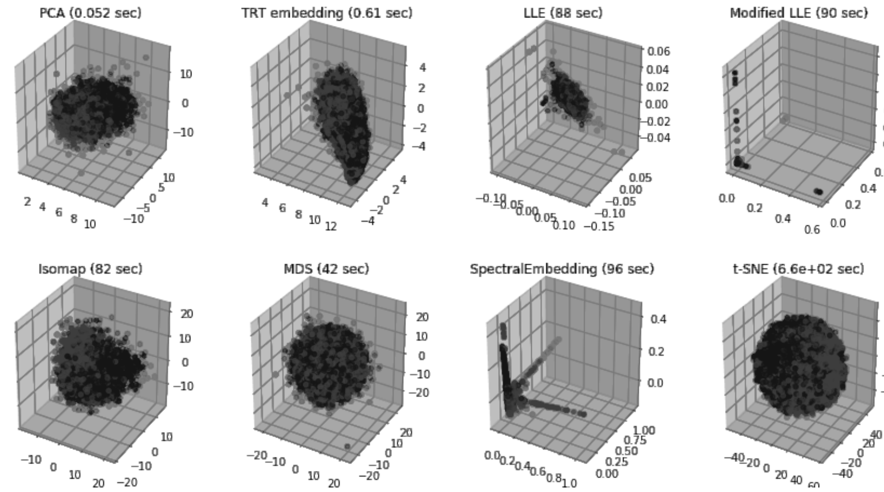

Figure 2. Visual classification of positive (light gray) and negative (dark gray) tweets.

We used 9079 random tweets labeled as positive and negative (the 70% of the dataset) as training instances. The performance of each classification algorithm was assessed on the remaining 30% of tweets by looking at the following measures: Accuracy (i.e. percentage of correctly classified tweets), Sensitivity (i.e. the proportion of positive tweets correctly classified) and Specificity (i.e. the proportion negative tweets correctly classified) and the F1-measure (i.e. the harmonic mean of Sensitivity and Specificity).

Table 1 shows the performance of the sentiment classification for the different input data formats. The Random Forest classifier achieves the best accuracy both when using the 20 dimensional PV-DM feature vectors (0.7658) and the 3 dimensional t-SNE vectors (0.7476). Classifiers perform slightly worse when low-dimensional features are used (e.g. the best accuracies differ in less than a 2%) but we point out that the output of the manifold methods has the advantage of providing a visual representation of polarity patterns in tweets (as shown in Figure 2). The temporal cost depends on the classification algorithm. Training times are worse for XGBTree or RF while they are better for GBM, AVNNet, NB and SVM. Regardless the classification method, though, test times are mostly faster when classifying tweets represented in reduced input data formats.

The low-dimensional representations can also be good at generating higher specificity values. See, for instance, the effect of training RF with the output of t-SNE (0.7728) and the combination of PCA or Spectral Embedding with other classifiers such as GBM and XGBTree. This is particularly interesting for this case study since negative opinions are more informative for decision-makers to identify improvements, being positive attitudes the default behaviour in a festivity. Finally, F1-measurements are comparable although better results are still obtained when using all features.

5. Conclusion

We have developed a system for visualizing and classifying tweets based on polarity. Our approach manages tweet pre-processing, feature extraction and, finally, sentiment

Table 1. Performance of the sentiment classification of tweets.

Method	Classifier	Accuracy	Sensitivity	Specificity	F1	t_{train}(s)	t_{test}(s)
PV-DM	GBM	0.7504	0.8252	0.6756	0.7678	30.6638	0.0156
PV-DM	XGBTree	0.7542	0.8082	0.7003	0.7668	1.6067	0.0350
PV-DM	RF	**0.7658**	0.7871	**0.7445**	0.7707	4.4359	0.2490
PV-DM	AVNNet	0.7566	0.7969	0.7162	0.7660	7.7932	0.0312
PV-DM	NB	0.7355	0.7357	0.7352	0.7355	35.0044	1.3456
PV-DM	SVM	0.7355	0.7357	0.7352	0.7355	11.1012	1.5116
ISOMAP	GBM	0.7242	0.7491	0.6992	0.7309	13.1520	0.0156
ISOMAP	XGBTree	0.7237	0.7589	0.6884	0.7331	57.6016	0.0000
ISOMAP	RF	0.7087	0.7378	0.6797	0.7170	54.6254	0.2328
ISOMAP	AVNNet	0.7252	0.7342	0.7162	0.7276	5.3752	0.0166
ISOMAP	NB	0.7201	0.7491	0.6910	0.728	8.5366	1.1726
ISOMAP	SVM	0.7201	0.7491	0.6910	0.728	8.9060	1.1204
t-SNE	GBM	0.7180	0.7640	0.6720	0.7304	13.3458	0.0120
t-SNE	XGBTree	0.7152	0.8144	0.6159	0.7409	58.6518	0.0000
t-SNE	RF	**0.7476**	0.7224	**0.7728**	0.7410	54.4308	0.2360
t-SNE	AVNNet	0.7247	0.7527	0.6967	0.7322	4.8772	0.0260
t-SNE	NB	0.7159	0.7342	0.6977	0.7210	8.7712	1.1992
t-SNE	SVM	0.7159	0.7342	0.6977	0.7210	9.0740	1.0696
PCA	GBM	0.6982	0.6416	0.7548	0.6801	13.3324	0.0050
PCA	XGBTree	0.6967	0.6391	0.7542	0.6781	58.4236	0.0000
PCA	RF	0.6799	0.5871	0.7727	0.6472	59.4890	0.2450
PCA	AVNNet	0.6997	0.5979	0.8015	0.6657	5.1324	0.0312
PCA	NB	0.6933	0.6617	0.7249	0.6833	8.5628	1.1388
PCA	SVM	0.6933	0.6617	0.7249	0.6833	9.5590	1.1262
Spectral	GBM	0.7177	0.6834	0.7520	0.7096	13.5312	0.0040
Spectral	XGBTree	0.7193	0.6926	0.7460	0.7136	57.9376	0.0000
Spectral	RF	0.6885	0.6715	0.7055	0.6853	55.2912	0.2340
Spectral	AVNNet	0.7173	0.7216	0.7129	0.7207	5.3140	0.0280
Spectral	NB	0.7007	0.7289	0.6725	0.7113	8.6428	1.1362
Spectral	SVM	0.7007	0.7289	0.6725	0.7113	12.4702	1.5912

visualization and classification. One of the most important pre-processing steps consists in tagging tweets as positive or negative through the computation of a polarity score that takes into account sentiment from both text and emojis. Another important aspect is how to obtain a feature representation for each tweet. In this paper we use a PV-DM word embedding to transform tweets into real-valued vectors and, then, a series of manifold methods that project tweets into a 3D space and allow visual classification. The system is open to alternative algorithms and it is by no means restricted to the use of any particular word embedding, dimension reduction or classification methods.

Future work will incorporate the content of hashtags in the pre-processing stage, as they can provide useful information for the sentiment tagging. We are also in the process of characterizing the effect of higher-dimensional word embeddings on the sentiment visualization and on the performance obtained by 3D and 2D feature vectors. Lastly, we plan to develop an interactive interface that allows any user to load Twitter data, to parametrize feature extraction and to visualize the sentiments of tweets using different manifold methods.

Acknowledgements

This work was partially supported by the Spanish Ministry of Science and Innovation Project TIN2015-66972-C5-5-R. The authors would also like to thank the Scientific Culture and Innovation Unit Chair for Scientific Dissemination at the University of Valencia for the financial support to the citizen science project La Falla Immaterial.

References

[1] A. Agarwal, B. Xie, I. Vovsha, O. Rambow, and R. Passonneau. Sentiment analysis of twitter data. In *Proceedings of the Workshop on Languages in Social Media*, LSM'11, pages 30–38, Stroudsburg, PA, USA, 2011. Association for Computational Linguistics.

[2] F. Aisopos, G. Papadakis, K. Tserpes, and T. Varvarigou. Content vs. context for sentiment analysis: A comparative analysis over microblogs. In *Proceedings of the 23rd ACM Conference on Hypertext and Social Media*, HT '12, pages 187–196, New York, NY, USA, 2012. ACM.

[3] K. Bauman, B. Liu, and A. Tuzhilin. Aspect based recommendations: Recommending items with the most valuable aspects based on user reviews. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 717–725. ACM, 2017.

[4] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15:1373–1396, 2002.

[5] A. Burgess. Ai in action. In *The Executive Guide to Artificial Intelligence*, pages 73–89. Springer, 2018.

[6] F. R. K. Chung. *Spectral Graph Theory*. Number 92 in CBMS Regional Conference Series in Mathematics. American Mathematical Society, 1997.

[7] C. Dijkmans, P. Kerkhof, and C. Beukeboom. A stage to engage: Social media use and corporate reputation. *Tourism Management*, 47:58–67, 2015.

[8] C. Ding, H. K. Cheng, Y. Duan, and Y. Jin. The power of the "like button": The impact of social media on box office. *Decision Support Systems*, 94:77 – 84, 2017.

[9] P. Geurts, D. Ernst, and L. Wehenkel. Extremely randomized trees. *Machine Learning*, 63(1):3–42, Apr 2006.

[10] P. E. Green, F. J. Camone, Jr., and S. M. Smith. *Multidimensional scaling: concepts and applications*. Allyn and Bacon, Needham Heights, Massachusetts, 1989.

[11] M. Jabreel and A. Moreno. Sitaka at semeval-2017 task 4: Sentiment analysis in twitter based on a rich set of features. In *Proceedings of the 11th International Workshop on Semantic Evaluation, SemEval@ACL 2017, Vancouver, Canada, August 3-4, 2017*, pages 694–699, 2017.

[12] M. Jabreel, A. Moreno, and A. Huertas. Do local residents and visitors express the same sentiments on destinations through social media? In R. Schegg and B. Stangl, editors, *Information and Communication Technologies in Tourism 2017*, pages 655–668, Cham, 2017. Springer International Publishing.

[13] Y. Kim. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1746–1751, 2014.

[14] P. Kralj Novak, J. Smailović, B. Sluban, and I. Mozetič. Sentiment of emojis. *PLoS ONE*, 10(12):e0144296, 2015.

[15] A. Laplante, T. D. Bowman, and N. Aamar. "i'm at #osheaga!": Listening to the backchannel of a music festival on twitter. In *Proceedings of the 18th International Society for Music Information Retrieval Conference, ISMIR 2017, Suzhou, China, October 23-27, 2017*, pages 585–591, 2017.

[16] Q. V. Le and T. Mikolov. Distributed representations of sentences and documents. In *ICML*, volume 32 of *JMLR Workshop and Conference Proceedings*, pages 1188–1196. JMLR.org, 2014.

[17] B. Liu. *Sentiment analysis: Mining opinions, sentiments, and emotions*. Cambridge University Press, 2015.

[18] E. Martínez-Cámara, M. Martín-Valdivia, L. López, and A. Montejo-Ráez. Sentiment analysis in twitter. *Natural Language Engineering*, 20(1):1–28, 2014.

[19] L. Medrouk and A. Pappa. Deep learning model for sentiment analysis in multi-lingual corpus. In *International Conference on Neural Information Processing*, pages 205–212. Springer, 2017.

[20] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 3111–3119, 2013.

[21] M. M. Mostafa. More than words: Social networks' text mining for consumer brand sentiments. *Expert Systems with Applications*, 40(10):4241 – 4251, 2013.

[22] P. Nakov. Semantic sentiment analysis of twitter data. *CoRR*, abs/1710.01492, 2017.

[23] B. Pang and L. Lee. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1–2):1–135, 2008.

[24] B. Pang, L. Lee, and S. Vaithyanathan. Thumbs up?: Sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing - Volume 10*, EMNLP '02, pages 79–86, Stroudsburg, PA, USA, 2002. Association for Computational Linguistics.

[25] K. Pearson. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2:559–572, 1901.

[26] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pages 1532–1543, 2014.

[27] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *SCIENCE*, 290:2323–2326, 2000.

[28] C. Silva, M. Antunes, J. Costa, and B. Ribeiro. Active manifold learning with twitter big data. *Procedia Computer Science*, 53:208 – 215, 2015. INNS Conference on Big Data 2015 Program San Francisco, CA, USA 8-10 August 2015.

[29] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319, 2000.

[30] F. S. Tsai. Dimensionality reduction techniques for blog visualization. *Expert Systems with Applications*, 38(3):2766 – 2773, 2011.

[31] P. D. Turney. 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*, ACL '02, pages 417–424, Stroudsburg, PA, USA, 2002. Association for Computational Linguistics.

[32] L. van der Maaten and G. Hinton. Visualizing high-dimensional data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, 2008.

[33] D. Vilares, M. A. Alonso, and C. Gómez-Rodríguez. Supervised sentiment analysis in multilingual environments. *Inf. Process. Manage.*, 53(3):595–607, 2017.

[34] D.-T. Vo and Y. Zhang. Target-dependent twitter sentiment classification with rich automatic features. In *Proceedings of the 24th International Conference on Artificial Intelligence*, IJCAI'15, pages 1347–1353. AAAI Press, 2015.

[35] Z. Zhang and J. Wang. Mlle: Modified locally linear embedding using multiple weights. In B. Schölkopf, J. C. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 1593–1600. MIT Press, 2007.