

Improving Sentiment Analysis over non-English Tweets using Multilingual Transformers and Automatic Translation for Data-Augmentation

Valentin Barriere

European Commission – DG-JRC

Via Enrico Fermi, 2749

21027 Ispra (VA), Italy

name.surname@ec.europa.eu

Alexandra Balahur

European Commission – DG-JRC

Via Enrico Fermi, 2749

21027 Ispra (VA), Italy

name.surname@ec.europa.eu

Abstract

Tweets are specific text data when compared to general text. Although sentiment analysis over tweets has become very popular in the last decade for English, it is still difficult to find huge annotated corpora for non-English languages. The recent rise of the transformer models in Natural Language Processing allows to achieve unparalleled performances in many tasks, but these models need a consequent quantity of text to adapt to the tweet domain. We propose the use of a multilingual transformer model, that we pre-train over English tweets on which we apply data-augmentation using automatic translation to adapt the model to non-English languages. Our experiments in French, Spanish, German and Italian suggest that the proposed technique is an efficient way to improve the results of the transformers over small corpora of tweets in a non-English language.

1 Introduction

Monitoring social media at the scale of a continent, like Europe, requires to process multiple languages. Data from Twitter is noisy and can drastically change in term of words distribution in one language when compared with general texts. It becomes even more challenging when tackling a multilingual task.

In accordance with the RGPD, it is impossible to make available tweets that have been deleted by their authors. This makes it more difficult to find tweet corpora annotated in sentiment. The SemEval challenges assure a big database of tweets in English, with a total of more than 62k examples annotated at the level of the tweet (Rosenthal et al., 2018). For other languages, this is more complicated.

Bidirectional transformers like BERT (Devlin et al., 2018) revolutionized the world of Natural Language Processing (NLP). Even when pre-trained over text from general domain, these models need a substantial amount of data to adapt to a domain where the syntax is different, like Twitter (Nguyen et al., 2020).

This paper presents the experiments carried out on several datasets of tweets in five different languages: English, French, Spanish, German and Italian. The general idea is pretty simple: instead of using a monolingual model, we chose to use a multilingual model that we can train over a large dataset of English tweets, over the original non-English tweets and over their automatic translations. We chose the multilingual transformer model XLM-RoBERTa from (Lample and Conneau, 2019) with a data-augmentation technique using machine translation. We investigated the effects of pre-training with English data and data-augmentation. We also compared performances of multilingual models against their monolingual French (Martin et al., 2020) and English (Liu et al., 2019) counterparts and found interesting improvements.

2 State of the Art

The related works section of this paper is shared between multilingual sentiment analysis, data-augmentation and sentiment analysis over tweets.

We can count several challenges tackling sentiment analysis over tweets like SemEval (Nakov et al., 2013) in English, TASS (Villena-Román et al., 2013) in Spanish or DEFT (Hamon et al., 2015) in French.

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

Lang.	Tweet
English	I'd rather dump gasoline all over myself and run into a burning building than use Excel.
French	Je préférerais jeter de l'essence partout et tomber dans un immeuble en feu plutôt que d'utiliser Excel.
German	Ich würde lieber Benzin auf mich werfen und in ein brennendes Gebäude laufen, als Excel zu benutzen.
Spanish	Prefiero tirar gasolina sobre mí mismo y correr hacia un edificio en llamas que usar Excel.
Italian	PREFERIREI buttarmi la benzina addosso e correre in un edificio in fiamme piuttosto che usare Excel.

Table 1: Examples of automatically translated tweets (original language in bold)

For the last years, the neural networks are ruling the sentiment analysis over tweets. (Cliche, 2017) won the sentiment polarity subtask of the SemEval-2017 challenge (Rosenthal et al., 2018) using a neural network approach. (Singh et al., 2019) improve the state-of-the-art with a different incorporation of the emojis: they use the descriptions instead of their unicodes before a BLSTM with word embeddings. Finally, (Nguyen et al., 2020) proposed a BERT model with the RoBERTa pre-training procedure (Liu et al., 2019) over 850M English tweets. This model gives state-of-the-art results on the SemEval-2017 dataset, but like all the other models, it is not adapted to multilingual data.

Although data-augmentation is not as developed for textual data as it is for images, there are ways to apply it to text (Wei and Zou, 2020). (Sennrich et al., 2016) augment their datasets using back-translation for pairing sentences for a MT task. (Kobayashi, 2018) uses language models in context in order to create new plausible sentences. (Fadaee et al., 2017) use data-augmentation for MT, by automatically translating at the level of words using plausible substitutions in order to create a new sentence.

The work that is closest to the one we are presenting in this article remains that of (Balahur and Turchi, 2013; Balahur et al., 2014) who tackle a sentiment analysis task over multilingual tweets. They show that use of multilingual, machine-translated data can help to better distinguish relevant features for sentiment classification, using SVM models with Bag-of-N-Grams. We distinguish from this work by using real datasets for testing instead of artificially created test sets made of translated tweets re-edited by humans.

3 Proposed Method

The method we propose is very simple. It basically consists in using a multilingual model instead of a monolingual model, pre-trained it over available annotated English tweets, combined with a data-augmentation technique that is based on MT. With this augmentation, we have each tweet in five examples, in five different languages. The languages that we use for the tweets are the same languages that the datasets we used to test the tweets: French, English, German, Spanish and Italian.

3.1 Pre-training over External Datasets

As we said earlier, we found it was more difficult to find tweet datasets in languages that were not English. We then investigated the potential of using multilingual model pre-trained over English tweets only, and over English tweets automatically translated in other languages.

We investigated the effect of using other available English datasets with multilingual model. To that end, we pre-trained the neural network with tweets annotated for the SemEval-2013 to SemEval-2016 challenges. We used the original tweets in English, but also their automatic translations in the 4 other languages we studied, leading to a 5 times bigger training dataset.

3.2 Data Augmentation and Multilingual Training

Translating the tweets into other languages allows our model to see 5 times the amount of data that it should have originally seen. The translations from the source language to the 4 other languages were made by the automatic translation tool of the European Commission, which is comparable to Google Translate. You can find examples of tweets and their translation in Table 1. It is important to note that the quality of the translation is not optimal since the translator has been learned over general text, whereas tweets can be noisy data containing abbreviations, and modernisms.

Finally, we always fine-tune the model over the non-English target datasets. The results of the models that are not fine-tuned on the target languages are poor and not even reported in this article.

Dataset	Language	Train	Dev	Test	All
SB-10k (Cieliebak et al., 2017)	German	4925	330	1315	6570
TASS-2019 (Díaz-Galiano et al., 2019)	Spanish	2133	506	581	3220
TASS-2018 (Martínez-Cámara et al., 2018)					
DEFT-2015 (Hamon et al., 2015)	French	6489	407	2938	9427
Sentipolc-16 (Basile et al., 2016)	Italian	6534	436	1964	8934
SemEval-2017 (Rosenthal et al., 2018)	English	47762	2000	12284	62046
SemEval-2013 (Nakov et al., 2013)					
SemEval-2014 (Rosenthal et al., 2015b)					
SemEval-2015 (Rosenthal et al., 2015a)					
SemEval-2016 (Nakov et al., 2016)					

Table 2: Datasets used in our experiments

4 Experiments and Results

4.1 Methodology

The pre-trained models that we used were made available online using the `transformers` library (Wolf et al., 2019). The same library has been employed for the training of the models. We used the Adam algorithm (Kingma and Ba, 2014) with early stopping for the optimization of the training loss, using a learning rate of $2e^{-6}$ for the pre-training of the model over English tweets, and $5e^{-7}$ for the fine-tuning over non-English tweets. We computed the performance on the development set after each training epoch, and kept the model obtaining the best performance. We used a batch size of 32.

We trained our models over 10 datasets and tested them over five different test sets in five languages. A summary of the datasets is shown in Table 2. It was impossible to obtain all of the original datasets because of the nature of the data: if a tweet has been deleted, it should not be available online. Nevertheless, we think that our results are competitive since we are using state-of-the-art models and obtain better results than what are reported in the articles using the original datasets. For the English test set, which is the exact one used for SemEval-2017, our results are higher than the winner of the challenge.

For the pre-training, we used the SemEval-2013 to SemEval-2016 challenges for a total of 47,762 tweets, and 238,810 using data-augmentation with automatic translation. The 2000 tweets from the devtest of SemEval-2016 were used as development set and the test set from SemEval-2017 was used as test set. For the French, German and Italian datasets, we used the same partition than the one used in the original challenges, with the tweets that were available online. For the Spanish dataset, the test sets were not available, hence we used the development set of TASS-2019 as test set, and the development set of TASS-2018 as development set. For the Italian dataset, we discarded the tweets with mixed sentiments.

We computed metrics that are broadly employed for this kind of tasks in order to compare our models: the Average-Recall, the average of the F1 score between positive and negative example, as well as the macro F1 score.

4.2 Results

The results of the experiments are shown in Table 3. We can compare the results of vanilla models with models pre-trained over English datasets for sentiment analysis and with models pre-trained over English and automatically translated datasets for sentiment analysis. MT was used to translate all the training sets into 4 other languages. One important thing to note is that we do not compare our system to other state-of-the-art systems on those datasets. This is due to the non availability of the complete datasets. Here we are focusing on the impact of data-augmentation using automatic translation combined with a multilingual model pre-trained over English tweets.

Nevertheless, we believe that the results of the first configuration, without any pre-training neither data-augmentation are very competitive. For example, the best result reported by the authors of the SB10k has a $F1_{PN}$ of 65.09, which is below the performance of 67.1 we obtained with the Vanilla configuration.

The best results overall non-English languages are obtained using the pre-training as well as the data-augmentation technique.

Language	Model	Using English	D-A	Rec _{avg}	F1 _{mac}	F1 _{PN}
English	(Cliche, 2017) (winner SemEval-2017)	✓	✗	68.1	∅	68.5
	(Nguyen et al., 2020) (SOTA)	✓	✗	73.2	∅	72.8
	Monolingual	✓	✗	72.8	71.7	72.3
	Multilingual	✓	✗	71.9	70.0	70.3
German	Multilingual	✓	✓	71.6	69.3	70.2
		✗	✗	72.6	73.9	67.1
		✓	✗	74.1	74.8	68.7
Spanish	Multilingual	✓	✓	74.2	74.7	68.5
		✗	✗	63.5	63.2	72.7
		✓	✗	68.3	68.1	76.0
French	Multilingual	✓	✓	69.8	69.6	78.2
		✗	✗	72.9	72.8	71.6
		✗	✗	72.5	72.4	71.0
Italian	Multilingual	✓	✓	73.8	73.7	72.2
		✓	✓	74.4	74.5	72.8
		✗	✗	63.0	60.7	55.3
All (non English)	Multilingual	✓	✓	67.1	64.4	60.2
		✓	✓	68.1	66.1	62.0
		✗	✗	68.0	67.6	66.6
All (non English)	Multilingual	✓	✓	70.8	70.3	69.3
		✓	✓	71.6	71.2	70.4

Table 3: Results of the different configurations. All the models were originally pre-trained over general text data.

4.3 Analysis

Because the original tweets datasets are not available online, it is difficult for us to compare with the results in the literature for the datasets other than English. Nevertheless, the focus of our paper is not on beating the state-of-the-art but propose an method to use multilingual data to enhance the performance of a model using non-english data. Interestingly, we found better results than (Nguyen et al., 2020) for both the RoBERTa and XLM-RoBERTa over SemEval-2017. We think that this may be the result of adjusting class weights in our loss function to manage imbalanced classes.

Monolingual versus multilingual As it is pointed out by (Nguyen et al., 2020), the best results over English are obtained using a monolingual model, when compared to the same multilingual model. Hence, RoBERTa reaches higher performances than its multilingual counterpart the XLM-RoBERTa. This behavior is reproduced on the French datasets using CamemBERT, the French version of RoBERTa (Martin et al., 2020). Nevertheless, the pre-training of the multilingual model allow to obtain an increase in the performance of the French model. We think that this may be due to the lack of available examples in the target language. This confirms the hypothesis that pre-training a multilingual model with available data to use it on another language can be a good strategy to improve the results on a target language having less available examples for training.

Effect of data-augmentation Finally, the data-augmentation technique improves slightly the results for almost every language in different proportions. The biggest amelioration is obtained for Spanish, with an improvement of 1.5 points of the average recall compared with the model only pre-trained over English. The improvement over German is questionable. This may be due to the size of the dataset. The German train set is more than twice the size of the Spanish, which is the language were pre-training gives the better boost to the performances.

5 Conclusion

We presented a technique that helps to improve the results of a sentiment analysis system over non-english tweets. We use multilingual model that is able to process external English data available in big quantities to pre-train the model, and machine translation to augment the dataset. This technique is simple and yet allows to take advantage of the multilingual models for non-English tweet datasets of limited size.

References

- Alexandra Balahur and Marco Turchi. 2013. Improving sentiment analysis in twitter using multilingual machine translated data. *International Conference Recent Advances in Natural Language Processing, RANLP*, (September):49–55.
- Alexandra Balahur, Marco Turchi, Ralf Steinberger, Jose Manuel Perea-Ortega, Guillaume Jacquet, Dilek Küçük, Vanni Zavarella, and Adil El Ghali. 2014. Resource creation and evaluation for multilingual sentiment analysis in social media texts. *Proceedings of the 9th International Conference on Language Resources and Evaluation, LREC 2014*, pages 4265–4269.
- Valerio Basile, Francesco Barbieri, Danilo Croce, Malvina Nissim, Nicole Novielli, and Viviana Patti. 2016. Evalita 2016 Sentipolc Task - Task Guidelines. pages 1–12.
- Mark Cieliebak, Jan Milan Deriu, Dominic Egger, and Fatih Uzdilli. 2017. A Twitter Corpus and Benchmark Resources for German Sentiment Analysis. pages 45–51.
- Mathieu Cliche. 2017. BB.twtr at SemEval-2017 Task 4: Twitter Sentiment Analysis with CNNs and LSTMs. *SemEval-2017*, (2014):573–580.
- Jacob Devlin, Ming-wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding.
- Manuel Carlos Díaz-Galiano, Manuel García-Vega, Edgar Casasola, Luis Chiruzzo, Miguel García-Cumbreras, Eugenio Martínez Cámara, Daniela Moctezuma, Arturo Montejo Ráez, Marco Antonio Sobrevilla Cabezudo, Eric Tellez, Mario Graff, and Sabino Miranda. 2019. Overview of TASS 2019: One more further for the global Spanish sentiment analysis corpus. *CEUR Workshop Proceedings*, 2421:550–560.
- Marzieh Fadaee, Arianna Bisazza, and Christof Monz. 2017. Data augmentation for low-Resource neural machine translation. *ACL 2017 - 55th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers)*, 2:567–573.
- Thierry Hamon, Amel Fraise, Patrick Paroubek, Pierre Zweigenbaum, and Cyril Grouin. 2015. Analyse des émotions , sentiments et opinions exprimés dans les tweets : présentation et résultats de l’ édition 2015 du défi fouille de texte (DEFT). *Actes de la 22e conférence sur le Traitement Automatique des Langues Naturelles*.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations*, pages 1–13.
- Sosuke Kobayashi. 2018. Contextual Augmentation : Data Augmentation by Words with Paradigmatic Relations. pages 452–457.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual Language Model Pretraining.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. (1).
- Louis Martin, Benjamin Muller, Ortiz Suarez Pedro Javier, Yoann Dupont, Laurent Romary, Eric Villemonte de la Clergerie, Djamé Seddah, and Benoît Sagot. 2020. CamemBERT: a Tasty French Language Model. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Eugenio Martínez-Cámara, Yudián Almeida-Cruz, Manuel Carlos Díaz-Galiano, Suilan Estévez-Velarde, Migueí A García-Cumbreras, Manuel García-Vega, Yoan Gutiérrez, Arturo Montejo-Ráez, Andrés Montoyo, Rafael Muñoz, Alejandro Piad-Morffis, and Julio Villena-Román. 2018. Overview of TASS 2018: Opinions, Health and Emotions Resumen de TASS 2018: Opiniones, Salud y Emociones. *CEUR Workshop Proceedings*, 2172:13–27.
- Preslav Nakov, Zornitsa Kozareva, Alan Ritter, Sara Rosenthal, Veselin Stoyanov, and Theresa Wilson. 2013. SemEval-2013 task 2: Sentiment analysis in Twitter. **SEM 2013 - 2nd Joint Conference on Lexical and Computational Semantics*, 2(SemEval):312–320.
- Preslav Nakov, Alan Ritter, Sara Rosenthal, Fabrizio Sebastiani, and Veselin Stoyanov. 2016. SemEval-2016 task 4: Sentiment analysis in twitter. *SemEval 2016 - 10th International Workshop on Semantic Evaluation, Proceedings*, pages 1–18.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. BERTweet: A pre-trained language model for English Tweets.

- Sara Rosenthal, Preslav Nakov, Svetlana Kiritchenko, Saif Mohammad, Alan Ritter, and Veselin Stoyanov. 2015a. SemEval-2015 Task 10: Sentiment Analysis in Twitter. (SemEval):451–463.
- Sara Rosenthal, Alan Ritter, Preslav Nakov, and Veselin Stoyanov. 2015b. SemEval-2014 Task 9: Sentiment Analysis in Twitter. (SemEval):73–80.
- Sara Rosenthal, Noura Farra, and Preslav Nakov. 2018. SemEval-2017 Task 4: Sentiment Analysis in Twitter. pages 502–518.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving neural machine translation models with monolingual data. *54th Annual Meeting of the Association for Computational Linguistics, ACL 2016 - Long Papers*, 1:86–96.
- Abhishek Singh, Eduardo Blanco, and Wei Jin. 2019. Incorporating emoji descriptions improves tweet classification. *NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, 1:2096–2101.
- Julio Villena-Román, Sara Lana-Serrano, Eugenio Martínez-Cámara, and José Carlos González-Cristóbal. 2013. TASS - Workshop on sentiment analysis at SEPLN. *Procesamiento de Lenguaje Natural*, 50:37–44.
- Jason Wei and Kai Zou. 2020. EDA: Easy data augmentation techniques for boosting performance on text classification tasks. *EMNLP-IJCNLP 2019 - 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing, Proceedings of the Conference*, pages 6382–6388.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace’s Transformers: State-of-the-art Natural Language Processing.