Comparative study of neural models for the COSET shared task at IberEval 2017

Luca Ambrosini[1] and Giancarlo Nicolò[2]

[1] Scuola Universitaria Professionale della Svizzera Italiana
[2] Univesitat Politècnica De València
luca.ambrosini@supsi.ch
giani1@inf.upv.es

Abstract. This paper describes our participation in the *Classification Of Spanish Election Tweets (COSET)* task at IberEval 2017. During the searching process for the best classification system, we developed a comparative study over possible combinations of corpus preprocessing, text representations and classification models. After an initial models exploration, we focus our attention over specific neural models. Interesting insight can be drawn from the comparative study helping future practitioners tackling tweets classification problems to create system baseline for their work.

Key words. Neural Networks, Natural Language Processing, Text Classification.

1 Introduction

Nowadays the pervasive use of social media as a mean of communication helps researchers to found useful insight over open problems in the field of Natural Language Processing. In this context, the *Twitter* social network has a huge role in text classification problems, because thanks to its *API* is possible to retrieve specific formatted text (i.e., a sentence of maximum 140 characters called tweet) from a huge real-time text database, where different users publish their daily statements.

This huge availability of data gives raise to the investigation of new text classification problems, with special interest in prediction problems related to temporal events that can influence statements published by social network users. An example of this problem category is the text classification related to general election, where the *Classification Of Spanish Election Tweets (COSET)* task at IberEval 2017 is a concrete example.

In COSET, the aim is to classify a corpus of political tweets in five categories related to specific political topics. This task can be analysed as a domain-dependent (i.e., political domain) constrained-text (i.e., tweet sentence) classification problem.

To tackle the above problem we built a classification system that can be decomposed in three main modules, each representing a specific approach widely used in the Natural Language Processing literature: text pre-processing, text representation and classification model. During the modules design, we explore different design combinations leading the system development to a comparative study over the possible modules interactions. Analysing the produced study interesting insight can be drawn to create a system baseline for the tweet classification problem.

In the following sections we firstly describe the COSET task (Section 2), then we report the development process of the classification system and its module design (Section 3), after that, the evaluation of deployed systems over the provided corpus is analysed (Section 4), finally, conclusion over the whole work are outlined (Section 5).

2 Task definition

The COSET shared task's [1] aim was to classify Spanish written tweets talking about the 2015 Spanish General Election, where each tweet had to be classified into one of five different categories: (i) political issues, related to the most abstract electoral confrontation; (ii) policy issues, about sectorial policies; (iii) personal issues, on the life and activities of the candidates; (iv) campaign issues, related with the evolution of the campaign; (v) and other issues.

Participants had access to a labelled corpus composed of training set (2242 tweets) and development set (250 tweets) for system benchmarking. We analysed it and find the following statistical information presented in table 1.

Table 1. Statistical analysis of given corpus' tweets.

	Average length	Maximum length
Chars	135	140
Words	140	49

3 Systems description

In this section we describe the tweet classification systems we built. From a module perspective we can describe our systems as composed of three main blocks: text pre-preprocessing (Section 3.2), text representation (Section 3.3) and classification model (Section 3.4).

3.1 Initial investigation

To address the tweets classification problem we began our investigation analysing some of the most widely used text representations and classifiers. In the analysing for possible text representations we began focusing our attention on lexical features based on: *Bag Of Words* [6],*Bag Of N-Grams* (bigrams and trigrams), both with and without *term frequency-inverse document frequency* normalization (i.e., TF-IDF norm). In relation to the classification models that can exploit the above representations, we analysed *Random Forest, Decision Trees, Support Vector Machines* and *Multi Layer Perceptron*. Since the results obtained with the combination of those *model + representation* were outperformed by neural network based models, due to space limitations their analysis is not reported in this paper, but rather we will focus on the module description of the neural models.

3.2 Text pre-processing

Regarding the text pre-preprocessing, it has to be mentioned that the corpus under observation can not be treated as proper written language, because computer-mediated communication (CMC) is highly informal, affecting diamesic[3] variation with creation of new items supposed to pertain lexicon and graphematic domains [7,8]. Therefore, in addition to well know pre-processing

---

[3] The variation in a language across medium of communication (e.g. Spanish over the phone versus Spanish over email)

approach, as stemming (i.e., ST), removal of stopwords (i.e., SW) , removal of punctuation (i.e., PR), specific tweets pre-processing techniques have to be taken in consideration.

From the previous consideration, we define a set of specific tweet pre-processing approach that take into consideration the following items: (i) mentions (i.e., MT), (ii) smiley (i.e., SM), (iii) emoji (i.e., EM), (iv) hashtags (i.e., HT), (v) numbers (i.e., NUM), (vi) URL (i.e., URL) (vii) and Tweeter reserve-word as RT and FAV (i.e., RW).

For each of these items we left the possibility to be removed or substituted by constant string (e.g. (i) *@ierrejon preparados para salir #PodemosRemontada* $\xrightarrow{substitution}$ *$MENTION preparados para salir $HASHTAG*, (ii) *@ierrejon preparados para salir #PodemosRemontada* $\xrightarrow{removing}$ *preparados para salir*).

To implement above pre-processing technique we took advantage of the following tools: (i) NLTK [4] and (ii) Preprocessor[4].

3.3 Text representation

The use of neural modelS suggest us to exploit recent trend over text representation, in particular we decided to use embedding vectors as representation following the approach described by [5], where tweet elements like *words* and *word n-grams* are represented as vectors of real number with fixed dimension $|v|$. In this way a whole sentence s, with length $|s|$ its number of word, is represented as a *sentence-matrix M* of dimension $|M| = |s| \times |v|$. $|M|$ has to be fixed a priori, therefore $|s|$ and $|v|$ have to be estimated. $|v|$ was fixed to 300 following [5]. $|s|$ was left as a system parameter that after optimization (via grid search) was fixed to $|s| = 30$, with this choice input sentences longer than $|s|$ are truncated, while shorter ones are padded with null vectors (i.e., a vector of all zeros). Depending of chosen tweets elements a different embedding function has to be estimated (i.e., learnt), following we are going to analyse the possible choices.

Word embedding. Choosing words as elements to be mapped by the embedding function, raise some challenge over the function estimation related to data availability. In our case the available corpus is very small and estimated embeddings could lead to low performance. To solve this problem, we decided to use a pre-trained embeddings estimated over Wikipedia using a particular approach called *fastText* [5], this choice was made after previous tries over other embeddings estimated from other corpus that lead to poor performance.

Using this approach, after the sentence-matrix embeddings are calculated, following [5] terminologies, matrix's weights can be set to *static* or *non-static*. In the latter case, backward propagation will be able to adjust its values otherwise they will stay fixed as initially calculated by the embedding function.

In this way four possible combinations of sentence-matrix embeddings can be formulated: (i) the use of a pre-trained embedding function (i.e., FastText from Wikipedia) and (ii) static or non-static weights. From this combination the one composed of static weight without pre-trained embeddings won't be take in consideration for obvious reasons, meaning that the cases in consideration will be three: (i) ES static, (ii) ES non-static, (iii) (no pre-trained embeddings) non-static.

N-gram embedding. Choosing n-grams as elements to be mapped by the embedding function, raises more challenges respect simple words, because no pre-trained embeddings are available and

---

[4] Preprocessor is a preprocessing library for tweet data written in Python, https://github.com/s/preprocessor

in this case the corpus has to be significantly big, otherwise n-gram frequencies will be really low and the estimation algorithm is not able to learn a valid embedding. Our insight was empirically validated by a very low performance. Nevertheless, as explained in the following, this embedding will be used in a particular model that won't rely its performance just over n-gram embeddings.

3.4 Classification models

Following, we describe the neural models used for the classification module, where for each of them the input layer uses text representations described in Section 3.3 (i.e., sentence-matrix).

Fast text. This model was introduced in [3], where its main difference from our neural model is the use of a particular input layer. In details, rather than using only words or only n-gram as element for the embedding, both elements are embedded with the aim of capturing partial information about words order. The architecture's idea is illustrated in Figure 1. Here the input layer is directly fed into a *Global Average Pooling* layer, that transforms the sentence-matrix in a single vector, which is projected into two dense layers. Regarding the architectural references in [3], they used a number of hidden layers fixed to ten, but we measured better performance using just two layers, moreover we integrate both dropout, gaussian noise and batch normalization.


Fig. 1. Model architecture of *fastText* for a sentence with N n-gram features $x_1, \ldots, x_N$. The features are embedded and averaged to form the hidden variable [3].

Convolutional Neural Network. Convolutional Neural Networks (CNN) are considered state of the art in many text classification problem. Therefore, we decide to use them in a simple architecture composed by a convolutional layer, followed by a *Global Max Pooling* layer and two dense layers. Between the two dense layers we used dropout(0.2) to avoid overfitting.

KIM. This model was introduced in [2]. It can be seen as a particular CNN where the convolutional layer has multiple kernels' size and feature maps. The complete architecture is illustrated in Figure 2, here the input layer (i.e., sentence-matrix) is processed in a convolutional layer of multiple filters with different sizes, each of these results are fed into *Max Pooling* layers and finally the concatenation of them (previously flatten to be dimensional coherent) is projected into a dense layer. The intuition behind this model is that smaller filter should be able to capture short sentence patterns similar to n-grams, while bigger ones should capture sentence level features. Regarding the architectural references in [2], the number filter $|f|$ and their size was optimized (via grid search) leading to the following results: $|f| = 4$, $f_1 = 2 \times 2$, $f_2 = 3 \times 3$, $f_3 = 5 \times 5$, $f_4 = 7 \times 7$.
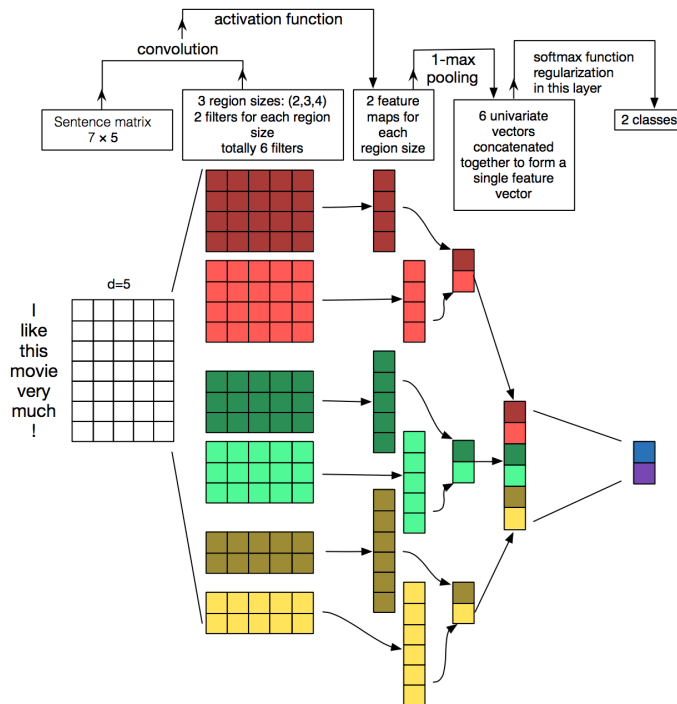
Fig. 2. Illustration of a Convolutional Neural Network architecture for sentence classification [9].

Long short-term memory. LSTM is a type of Recurrent Neural Network (RNN) that is relatively insensitive to gap length. Thanks to this behavior, they are considered state of the art in some NLP problems. Our architecture was made of an embedded input layer followed by an LSTM layer of 128 units, terminated by a dense layer. Moreover, to avoid overfitting we used dropout (0.25) and recurrent dropout (0.25).

Bidirectional LSTM. Similar to the previous model, bidirectional LSTM is a variation of LSTM where the two RNN receive different inputs, the original and its reverse order, and their results are connected through the recurrent layers. Our architecture follows the previous one with an LSTM layer of 128 units terminating with two dense layers, where all layers used dropout (0.25) and recurrent dropout(0.25).

4 Evaluation

In this section we are going to illustrate results from the comparative study elaborated during the system development. First we illustrate the metric used to evaluate the system (Section 4.1) and then we report results produced by a 10-fold cross validation over the given data set (Section 4.2), finally we report our performance at the shared task (Section 4.3).

4.1 Metrics

System evaluation metrics were given by the organizers and reported here in the following equations (1) to (4). Their choice was to use an $F_{1-macro}$ measure due to class unbalance in the corpus.

$$F_{1-macro} = \frac{1}{|L|} \sum_{l \in L} F_1(y_l, \hat{y}_l) \qquad (1) \qquad\qquad precision = \frac{1}{|L|} \sum_{l \in L} Pr(y_l, \hat{y}_l) \qquad (3)$$

$$F_1 = 2 \cdot \frac{precision \cdot recall}{precision + recall} \qquad (2) \qquad\qquad recall = \frac{1}{|L|} \sum_{l \in L} R(y_l, \hat{y}_l) \qquad (4)$$

where L is the set of classes, y_l is the set of correct label and $\hat{y}_l$ is the set of predicted labels.

4.2 Comparative study

Following, we present a comparative study over possible combinations of pre-processing (Table 2) and word embeddings (Table 3), in both cases results are calculated from averaging three runs of a 10-fold cross validation over the complete data set. Notations used in Table 2 refer to the one introduced in Section 3.2, where the listing of a notation means its use for the reported result. Regarding the tweet specific pre-processing, all the items have been substituted, with the exception for URL and RW that have been removed. We report the contribution of each analysed pre-processing alone. To not overwhelm reader with verbose data, reported results are focused only over the two best performing model (Kim's model and FastText).

Table 2. Pre-processing study comparing 10-fold cross validation results over the development set in terms of percentage of $F_{1-macro}$ score. For each model processing technique that brought an improvement has its result in bold.

Models	Pre-processing									
	Nothing	ST	SW	URL	RW	MT	HT	NUM	EM	SM
Kim	0.543	0.528	**0.557**	**0.571**	0.533	**0.558**	0.540	**0.554**	0.537	0.539
FastText	0.546	0.533	**0.550**	0.534	**0.553**	0.519	0.538	**0.558**	**0.552**	**0.566**

From the analysis of Table 2 no absolute conclusion can be drawn, meaning that it wasn't possible to find a combination of pre-processing that gives the best performance for all the models, meaning that each model is highly sensible to the performed combination. Nevertheless, some relative observation can be made:

- SW (i.e., removing Spanish stopwords) and NUM (i.e., substitute numbers with a constant string) lead to performance improvement to all the model respect to no pre-processing at all,
- ST (i.e., stemming) and HT (i.e., substitute hashtags with a constant string) decrease the performance of both models respect to non-preprocessing at all.

From above observations, to continue our study and fix system's pre-processing pipeline, for each classification model we decided to use the combination of pre-processings that singularly applied give improvements to the system.

Analysing results in Table 3, here the used notation refers to the one introduced in Section 3.3, where the listing of a notation means its use as embedded input layer for the reported result. From its analysis the following interpretation can be drawn:

Table 3. Word embeddings study comparing 10-fold cross validation results over the development set in terms of percentage of $F_{1-macro}$ score. For each model the best performing word embeddings configuration has its result in bold.

Models	Text representation		
	Non-static	ES static	ES non-static
Kim	0.541 ($\pm$ 0.019)	0.550 ($\pm$ 0.015)	**0.579 ($\pm$ 0.018)**
FastText	0.556 ($\pm$ 0.013)	0.450 ($\pm$ 0.011)	**0.589 ($\pm$ 0.010)**

- Setting as *static* the sentence matrix weights has the worst performance
- Setting as *non-static* leads to a better performance, where this insight can be deduced by corpus characteristic (i.e., a good example of Computer Mediated Communication)
- The use of pre-trained embedding is useful in combination with *non-static* weights (i.e., best performances with ES non-static)

In table 4 we report a complete overview of the evaluated models in respect to their best configurations of text pre-processing and word embedding. As can be seen, best performances are obtained by FastText and Kim's models, while recurrent models have the worst performance.

Table 4. Best configurations study comparing 10-fold cross validation results over the development set in terms of percentage of $F_{1-macro}$ score.

System	$F_{1-macro}$
LSTM	0.556 ($\pm$ 0.012)
Bi-LSTM	0.555 ($\pm$ 0.035)
CNN	0.571 ($\pm$ 0.030)
FastText	**0.589 ($\pm$ 0.018)**
Kim	0.579 ($\pm$ 0.009)

4.3 Competition results

For the system's submission, participants were allowed to send more than a model till a maximum of 5 possible runs, table 5 reports our best performing systems at the COSET shared task.

Table 5. Resulted obtained in the shared task participation. The absolute and team column represent the ranking over the whole participants.

System	$F_{1-macro}$	Absolute	Team
FastText	**0.6157**	7/39	4/17
Kim	0.6065	8/39	4/17

5 Conclusions

In this paper we have presented our participation in the IberEval2017 Classification Of Spanish Election Tweets (COSET) shared task. Five distinct neural models were explored, in combination

with different types of preprocessing and text representation. From the systems evaluation it wasn't possible to find a combination of pre-processing that gives the best performance for all the models, meaning that each model is highly sensible to the pipeline combination. Regarding the analysed text representation, the setting of sentence matrix to non-static always leads to good performance as a result of the specific text under observation (i.e., a CMC corpus) and the use of pre-trained word embedding is always suggested. Moreover, we outline a not so promising performance of the recurrent model, meaning that for this task the word order (a feature well captured by LSTM family model) seems not so prominent as in other tasks.

References

1. Giménez M., Baviera T., Llorca G., Gámir J., Calvo D., Rosso P., Rangel F. Overview of the 1st Classification of Spanish Election Tweets Task at IberEval 2017. In: Notebook Papers of 2nd SEPLN Workshop on Evaluation of Human Language Technologies for Iberian Languages (IBEREVAL), Murcia, Spain, September 19, CEUR Workshop Proceedings. CEUR-WS.org, 2017.
2. Kim, Yoon. "Convolutional neural networks for sentence classification." arXiv preprint arXiv:1408.5882 (2014).
3. Joulin, Armand, et al. "Bag of tricks for efficient text classification." arXiv preprint arXiv:1607.01759 (2016).
4. Edward Loper and Steven Bird. 2002. NLTK: the Natural Language Toolkit. In Proceedings of the ACL-02 Workshop on Effective tools and methodologies for teaching natural language processing and computational linguistics - Volume 1 (ETMTNLP '02), Vol. 1. Association for Computational Linguistics, Stroudsburg, PA, USA, 63-70. DOI=http://dx.doi.org/10.3115/1118108.1118117
5. Bojanowski, Piotr and Grave, Edouard and Joulin, Armand and Mikolov, Tomas. "Enriching Word Vectors with Subword Information" arXiv preprint arXiv:1607.04606 (2016).
6. Harris, Zellig S. "Distributional structure." Word 10.2-3 (1954): 146-162.
7. Bazzanella, Carla. "Oscillazioni di informalità e formalità: scritto, parlato e rete." Formale e informale. La variazione di registro nella comunicazione elettronica. Roma: Carocci (2011): 68-83.
8. Cerruti, Massimo, and Cristina Onesti. "Netspeak: a language variety? Some remarks from an Italian sociolinguistic perspective." Languages go web: Standard and non-standard languages on the Internet (2013): 23-39.
9. Zhang, Ye, and Byron Wallace. "A sensitivity analysis of (and practitioners' guide to) convolutional neural networks for sentence classification." arXiv preprint arXiv:1510.03820 (2015).