

Using Word Embeddings in Twitter Election Classification

Xiao Yang

Craig Macdonald

Iadh Ounis

{firstname.lastname}@glasgow.ac.uk
University of Glasgow, UK

ABSTRACT

Word embeddings and convolutional neural networks (CNN) have attracted extensive attention in various classification tasks for Twitter, e.g. sentiment classification. However, the effect of the configuration used to train and generate the word embeddings on the classification performance has not been studied in the existing literature. In this paper, using a Twitter election classification task that aims to detect election-related tweets, we investigate the impact of the background dataset used to train the embedding models, the context window size and the dimensionality of word embeddings on the classification performance. By comparing the classification results of two word embedding models, which are trained using different background corpora (e.g. Wikipedia articles and Twitter microposts), we show that the background data type should align with the Twitter classification dataset to achieve a better performance. Moreover, by evaluating the results of word embeddings models trained using various context window sizes and dimensionalities, we found that large context window and dimension sizes are preferable to improve the performance. Our experimental results also show that using word embeddings and CNN leads to statistically significant improvements over various baselines such as random, SVM with TF-IDF and SVM with word embeddings.

1. INTRODUCTION

Word embeddings have been proposed to produce more effective word representations. For example, in the *Word2Vec* model [14], by maximising the probability of seeing a word within a fixed context window, it is possible to learn for each word in the vocabulary a dense real valued vector from a shallow neural network. As a consequence, similar words are close to each other in the embedding space [4, 7, 14]. The use of word embeddings together with convolutional neural networks (CNN) has been shown to be effective for various classification tasks such as sentiment classification on Twitter [9, 17]. However, the effect of the configuration used to

generate the word embeddings on the classification performance has not been studied in the literature. Indeed, while different background corpora (e.g. Wikipedia, GoogleNews and Twitter) and parameters (e.g. context window and dimensionality) could lead to different word embeddings, there has been little exploration of how such background corpora and parameters affect the classification performance.

In this paper, using a dataset of tweets collected during the Venezuela parliamentary election in 2015, we investigate the use of word embeddings with CNN in a new classification task, which aims to identify those tweets that are related to the election. Such a classification task is challenging because election-related tweets are usually ambiguous and it is often difficult for human assessors to reach an agreement on their relevance to the election [5]. For example, such tweets may refer to the election implicitly without mentioning any political party or politician. In order to tackle these challenges, we propose to use word embeddings to build richer vector representations of tweets for training the CNN classifier on our election dataset.

We thoroughly investigate the effect of the background corpus, the context window and the dimensionality of word embeddings on our election classification task. Our results show that when the type of background corpus aligns with the classification dataset, the CNN classifier achieves statistically significant improvements over the most effective classification baseline of SVM with TF-IDF on our task. We also show that word embeddings trained using a large context window size and dimension size can help CNN to achieve a better classification performance. Thus, our results suggest indeed that the background corpus and parameters of word embeddings have an impact on the classification performance. Moreover, our results contradict the findings of different tasks such as dependency parsing [3] and named entity recognition (NER) [10] where a smaller context window is suggested. Such a contradiction suggests that the best setup of parameters such as the context window and dimensionality might differ from a task to another.

In the remainder of this paper, we briefly explain the related work in Section 2. We describe and illustrate the CNN architecture used for our classification task in Section 3. In Section 4, we describe our dataset and the experimental setup. In Section 5, we discuss the impact of two background corpora (Wikipedia articles and Twitter microposts) on the effectiveness of the learned classifier. In Section 6, we investigate the impact of the context window size and dimensionality of word embeddings on the classification performance. We provide concluding remarks in Section 7.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Neu-IR '16 SIGIR Workshop on Neural Information Retrieval, July 21, 2016, Pisa, Italy.

©2016 Copyright held by the owner/author(s).

2. RELATED WORK

A number of studies have already shown that the context window and dimensionality of the used word embedding vectors could affect performance in tasks such as dependency parsing [3] and named entity tagging [10]. For instance, using publicly available corpora such as Wall Street Journals and Wikipedia, Bansal et al. [3] investigated *Word2Vec* word embeddings in the dependency parsing task, which aims to provide a representation of grammatical relations between words in a sentence. By only varying the context window size from 1 to 10, their results on the accuracy of part-of-speech (POS) tagging showed that the context window size of *Word2Vec* could affect the type of the generated word embedding. In particular, they observed that a smaller context window gives a better performance on accuracy. In the named entity recognition (NER) task, Godin et al. [10] investigated three context window sizes w of $w = \{1, 3, 5\}$ based on the accuracy of NER tagging. Their results also reached the same conclusion, namely that a smaller context window gives a better performance using the *Word2Vec* word embeddings when the model is trained from a large Twitter corpus containing 400 million tweets.

Using a subset of the semantic-syntactic word relationship test set, Mikolov et al. [14] investigated the dimensionality of the *Word2Vec* word embeddings and the size of background data. In the test set, word pairs are grouped by the type of relationship. For example “brother-sister” and “grandson-granddaughter” are in the same relationship of “man-woman”. The accuracy is measured such that given a word pair, another word pair with the correct relationship should be retrieved. Using this accuracy measure, they noted that at some point increasing the dimensionality or the size of background data only provides minor improvements. Thus, they concluded the dimensionality and background data size should be increased together [14]. However, Mikolov et al. [14] only investigated the *Word2Vec* parameters using the GoogleNews background corpus.

The aforementioned studies provide a useful guide about the effect of the word embeddings configuration on performance in the specific applications they tackled, but their findings were obtained on tasks different from Twitter classification tasks. Hence, the question arises as whether such findings will generalise to classification tasks on Twitter, which is the object of our study in this paper.

In fact, there is little work in the literature tackling the task of election classification on Twitter. However, similar classification tasks such as Twitter sentiment classification have been well studied [9, 17, 19]. In particular, word embeddings were recently used to build effective tweet-level representations for Twitter sentiment classification [17, 19]. For instance, in the Semeval-2015 Twitter Sentiment Analysis challenge, Severyn et al. [17] proposed to use word embeddings learned from two Twitter corpora to build the vector representations of tweets. Using the *Word2Vec* model, default parameter values such as context window size 5 and dimensionality 100 were applied to train the word embedding. In their approach, one Twitter background corpus (50 million tweets) was used to train the word embedding, while another one (10 million tweets) containing positive and negative emoticons was used to refine the learned word embeddings using the proposed CNN classifier. The CNN classifier was then trained on the Semeval-2015 Twitter sentiment analysis dataset, which contains two subsets: phrase-level

dataset and message-level dataset. Each subset contains 5K+ and 9K+ training samples, respectively. The official ranking in Semeval-2015 showed that this system ranked 1st and 2nd on the phrase-level dataset and the message-level dataset, respectively. However, Severyn et al. [17] focused on refining the word embeddings by using another Twitter corpus with emoticons to learn sentiment information, but did not study the impact of the background corpus and the chosen parameters on the classification performance.

In another approach based on the word embeddings model proposed by Collobert et al. [6], Tang et al. [19] proposed a variation to learn sentiment-specific word embeddings (SSWE) from a large Twitter corpus containing positive and negative emoticons. Tang et al. [19] empirically set the context window size to 3 and the embedding dimensionality to 50. The Semeval-2013 Twitter sentiment analysis dataset, which contains 7K+ tweets was used to evaluate the effectiveness of their proposed approach. Compared to the top system of the Semeval-2013 Twitter Sentiment Analysis challenge, their approach of using an SVM classifier with SSWE outperformed the top system on the F1 measure. However, only the Twitter background corpus was used by Tang et al. [19], which contains 10 million tweets with positive and negative emoticons. On the other hand, the parameters of word embeddings such as the context window and dimensionality were not studied by Tang et al. [19], nor in the existing literature for Twitter classification tasks. As such, in this paper, we conduct a thorough investigation of word embeddings together with CNN on a Twitter classification task and explore the impact of both the background corpus, the context window and the dimensionality of word embeddings on the classification performance.

3. THE CNN MODEL

For our Twitter election classification task, we use a simple CNN architecture described by Kim [11] as well as the one proposed by Severyn et al. [18] and highlighted in Fig. 1. It consists of a convolutional layer, a max pooling layer, a dropout layer and a fully connected output layer. Each of these layers is explained in turn.

Tweet-level representation. The inputs of the CNN classifier are preprocessed tweets that consist of a sequence of words. Using word embeddings, tweets are converted into vector representations in the following way. Assuming $w_i \in \mathbb{R}^n$ to be the n -dimensional word embeddings vector of the i th word in a tweet, a tweet-level representation is obtained by looking up the word embeddings and concatenating the corresponding word embeddings vectors of the total k words:

$$TCNN = w_1 \oplus w_2 \oplus \dots \oplus w_k \quad (1)$$

where \oplus denotes the concatenation operation [11]. For training purposes, short tweets in our dataset are padded to the length of the longest tweet using a special token. Hence the total dimension of the vector representation $TCNN$ is always $k \times n$. Afterwards, the tweet-level representation will feed to the convolutional layer.

Convolutional layer. The convolution operation helps the network to learn the important words no matter where they appear in a tweet [17]. In this layer, the filter $F_i \in \mathbb{R}^{m \times n}$ with different sizes of m are applied to the tweet-level representation $TCNN$. By varying the stride s [12], we can shift the filters across s word embeddings vectors at each step. By sliding the filters over m word vectors in $TCNN$

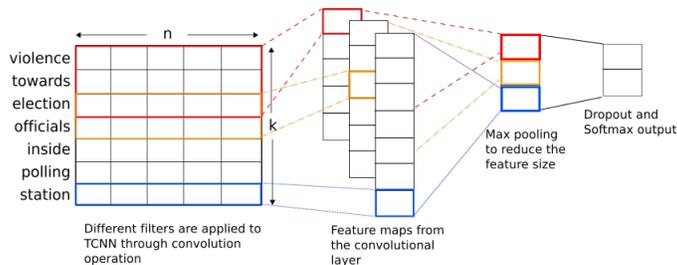


Figure 1: Convolutional neural network architecture for tweet classification. Adapted from [11].

using stride s , the convolution operation produces a new feature map c_i for all the possible words in a tweet:

$$c_i = f(F_i \cdot TCNN_{i:i+m-1} + b_i) \quad (2)$$

where $i : i + m - 1$ denotes the word vectors of word i to word $i + m - 1$ in $TCNN$. b_i is the corresponding bias term that is initialised to zero and learned for each filter F_i during training. In Eq. (2), f is the activation function. In this CNN architecture, we used a rectified linear function (ReLU) as f . No matter whether the input x is positive or negative, the ReLU unit ensures its output (i.e. c_i) is always positive as defined by $f = \max(0, x)$.

Max pooling layer. All the feature maps c_i from the convolutional layer are then applied to the max pooling layer where the maximum value c_i^{max} is extracted from the corresponding feature map. Afterwards, the maximum values of all the feature maps c_i are concatenated as the feature vector of a tweet.

Dropout layer. Dropout is a regularization technique that only keeps a neuron active with some probability p during training [11]. After training, $p = 1$ is used to keep all the neurons active for predicting unseen tweets. Together with the $L2$ regularization, it constraints the learning process of the neural networks by reducing the number of active neurons.

Softmax Layer. The outputs from the dropout layer are fed into the fully connected softmax layer, which transforms the output scores into normalised class probabilities [11]. Using a cross-entropy cost function, the ground truth labels from human assessors are used to train the CNN classifier for our Twitter election classification task.

During training, the weights of each layer are updated according to the loss between the prediction and the target. Once a CNN classifier is trained from a training set, all of its parameters and learned weights are saved into binary files that can be loaded to classify unseen tweets using the same procedures explained in this section.

4. EXPERIMENTAL SETUP

In this paper, we argue that the types of background corpora as well as the parameters of *Word2Vec* model could lead to different word embeddings and could affect the performance on Twitter classification tasks. In the following sections, experiments are tailored to conduct a thorough investigation of word embeddings together with CNN on a Twitter classification task and to explore the impact of the background corpora (Section 5), the context window and the dimensionality of word embeddings (Section 6) on the classification performance. The remainder of this section details our dataset (Section 4.1), our experimental setup and

	Election	Non-Election	Total	# Words
Dataset	2,274	3,474	5,747	9,904

Table 1: Statistics of the dataset used in the experiments. Negative class is the majority in the dataset.

used word embedding models (Section 4.2), baselines (Section 4.3) and measures (Section 4.4).

4.1 Dataset

Our manually labelled election dataset is sampled from tweets collected about the 2015 Venezuela parliamentary election using the well-known pooling method [16]. It covers the period of one month before and after the election date (06/12/2015) in Venezuela. We use the Terrier information retrieval (IR) platform [13] and the DFReeKLIM [2] weighting model designed for microblog search to retrieve tweets related to 21 query terms (e.g. “violencia”, “eleccion” and “votar”). Only the top 7 retrieved tweets are selected per query term per day, making the size of the collection realistic for human assessors to examine and label the tweets. Sampled tweets are merged into one pool and judged by 5 experts who label a tweet as: “Election-related” or “Not Election-related”. To determine the judging reliability, an agreement study was conducted using 482 random tweets that were judged by all 5 assessors. Using Cohen’s *kappa*, we found a moderate agreement of 52% between all assessors. For tweets without a majority agreement, an additional expert of Venezuela politics was used to further clarify their categories. In total, our election dataset consists of 5,747 Spanish tweets, which contains 9,904 unique words after preprocessing (stop-word removal & Spanish Snowball stemmer). Overall, our labelled election dataset covers significant events (e.g. Killing of opposition politician Luis Diaz [1]) in the 2015 Venezuela parliamentary election. From the general statistics shown in Table 1, we observe that the dataset is unbalanced; the majority class (Non-Election) has 1,000 more tweets than the minority class (Election).

4.2 Word embeddings

The word embeddings used in this paper are trained from two different background corpora: a Spanish Wikipedia dump dated 02/10/2015 (denoted *es-Wiki*) and a Spanish Twitter data (denoted *es-Twitter*) collected from the period of 05/01/2015 to 30/06/2015. Over 1 million Spanish articles are observed in *es-Wiki*. In *es-Twitter*, over 20 million Spanish tweets are collected by removing tweets with less than 10 words, hence the short and less informative tweets are not considered. For consistency, we apply the same preprocessing namely stop-word removal and stemmer (see Section 4.1) to both of the background corpora. Af-

	es-Wiki	es-Twitter
# Documents	1M+	20M+
Vocabulary Size	436K	629K
Word Coverage Count	5,111	6,612
Word Coverage Rate	51%	66%

Table 2: Statistics of the background corpora and words coverage on the election dataset.

ter the preprocessing, **es-Wiki** contains 436K unique words while **es-Twitter** has 629K unique words. Salient statistics are provided in Table 2. Indeed, by comparing the unique words in our election dataset with the words in **es-Wiki** and **es-Twitter**, we observe that 5,111 words in our dataset appear in **es-Wiki** while 6,612 words appear in **es-Twitter**. This shows that **es-Twitter** has a better word coverage on our election dataset.

We use the *Word2Vec* implementation in *deeplearning4j* to generate a set of word embeddings by varying the context window size W and the dimensionality D . We use the same context window sizes $W = \{1, 3, 5\}$ that were used by Godin et al. [10]. For each context window W , we use three different dimension sizes $D = \{200, 500, 800\}$ to cover both of the low and high dimensionalities of the word embedding vectors, which were used by Mikolov et al. [14]. Therefore, 9 word embeddings in total are generated by varying W and D . For other parameters, we use the same values that were set by Mikolov et al. [14]: We set the batch size to 50, negative sampling to 10, minimum word frequency to 5 and iterations to 5. As suggested by Kim [11], for a word not appearing in a word embeddings (also known as *out-of-vocabulary OOV*), we generate its vector by sampling each dimension from the uniform distributions $U_i[m_i - s_i, m_i + s_i]$, where m_i and s_i are the mean and standard deviation of the i th dimension of the word embeddings.

4.3 Baselines

To evaluate the CNN classifiers and word embeddings, we use three baselines, namely:

Random classifier: The random classifier simply makes random predictions to the test instances.

SVM with TF-IDF (SVM+TFIDF): As a traditional weighting scheme, TF-IDF is used in conjunction with an SVM classifier for the Twitter election classification.

SVM with word embeddings (SVM+WE): We use a similar scheme that was used by Wang et al. [20] to build the tweet-level representation for the SVM classifiers. The vector representation (i.e. *TWE*) of a tweet is constructed by averaging the word embedding vectors along each dimension for all the words in the tweet:

$$TWE = \sum_{i=1}^k w_i / k \quad (3)$$

where k is the number of words in a tweet and $w_i \in \mathbb{R}^n$ denotes the word embedding vector of the i th word. The vector representation of each tweet has exactly the same dimension as the word embedding vector w_i , which is the input of an SVM classifier.

4.4 Hyperparameters and measures

For all the experiments, we use 3 filter sizes $m = \{1, 2, 3\}$, stride $s = 1$ and dropout probability $p = 0.5$ for our CNN classifier, following the settings used by Kim [11]. For each

Classifier	es-Wiki			es-Twitter		
	P	R	F1	P	R	F1
SVM+WE	74.9	68.0	71.3	74.3	70.1	72.1
CNN	81.6	70.7	75.8	80.9	72.2	76.3

Table 3: Classification results by using the background corpora es-Wiki and es-Twitter.

filter size, 200 filters are applied to the convolutional layer and therefore 600 feature maps are produced in total. For the SVM classifier, we use the default parameter $c = 1$ for the *LinearSVC* implementation in *scikit-learn*¹ [15].

To train the classifiers and evaluate their performances on our dataset, we use a 5-fold cross validation, such that in each fold, 3 partitions are used for training, 1 partition for validation and 1 partition for test. We stop the training process when the classification accuracy on the validation partition declines. Afterwards, the overall performance on the test instances is assessed by averaging the scores across all folds. We report effectiveness in terms of classification measures, precision (denoted P), recall (denoted R) and F1 score (denoted $F1$).

5. EFFECT OF THE BACKGROUND CORPORA

Due to the noisy nature of Twitter data, Twitter posts can often be poor in grammar and spelling. Meanwhile, Twitter provides more special information such as Twitter handles, HTTP links and hashtags which would not appear in common text corpora. In order to infer whether the type of background corpus could benefit the Twitter classification performance, we compare the two background corpora of **es-Wiki** and **es-Twitter**. By considering the various experimental results in [3, 10, 14], the context window size of 5 is said to give a good performance. Thus, in this experiment we set the context window to 5 and the dimensionality to 500 for both word embeddings.

The classification results are shown in Table 3 where the first column shows the classifiers we used. In other columns, we report three measures for both the background corpora **es-Wiki** and **es-Twitter**. Since the SVM+TFIDF and random classifier do not use the background corpus, they are not listed in Table 3. For each classifier, the best scores are highlighted in bold. From Table 3, we observe that when the type of background corpus aligns with our Twitter election dataset, the performance is better for both the SVM+WE and CNN classifiers on Recall and F1 score. In particular, the improvement on recall suggests that **es-Twitter** represents the characteristics of Twitter posts better than the **es-Wiki** corpus.

As shown in the statistics of the two background corpora (Table 2), 66% of the vocabulary of our election dataset appears in **es-Twitter** while only 51% appears in **es-Wiki**. By removing the words covered by both background corpora, we observe that 1,527 unique words are covered by **es-Twitter** but not covered by **es-Wiki**. However, there are only 26 unique words that are covered by **es-Wiki** only. Table 4 categorises the words only found in **es-Twitter**, which are mostly words unique to Twitter, such as Twitter handles and hashtags. This explains why **es-Twitter** works better with our Twitter election dataset. The other 374 words

¹An open source machine learning library in Python.

Twitter handles	Hashtags	Others	Total
818	335	374	1,527

Table 4: Statistics of the vocabulary only covered by es-Twitter.

are mainly incorrect spellings and elongated words such as “bravoooo”, “yaaaa” and “urgenteeee”, which occur more often in Twitter than in other curated types of data such as Wikipedia and News feeds. Our finding on the vocabulary coverage further validates our results shown in Table 3. Thus, the results may generalise to similar Twitter classification tasks that also deal with Twitter posts. In summary, we find that aligning the type of background corpus with the classification dataset leads to better feature representations, and hence a more effective classification using the CNN classifier.

6. EFFECT OF WORD EMBEDDINGS PARAMETERS

In this section, we attempt to investigate the effect of parameters (e.g. context window and dimensionality) for the Twitter election classification task. Since **es-Twitter** gives a better performance, we only use word embeddings generated from **es-Twitter** only. Table 5(a) shows the results of our three baselines, while Table 5(b) shows the results of classifiers using word embeddings, namely SVM with word embeddings (SVM+WE) and CNN. In Table 5(b), the measurements for SVM+WE and CNN are arranged by the dimensionality and context window size of word embeddings. For each row of $W1$, $W3$ and $W5$, Table 5(b) shows results for context window sizes of $W = \{1, 3, 5\}$ along each dimension sizes of $D = \{200, 500, 800\}$. The best overall scores are highlighted in bold.

We first compare the results of the CNN classifiers to the random baseline and the SVM+WE baseline. Clearly, the CNN classifiers outperform these two baselines across all measures. By comparing CNN classifiers to the best baseline SVM+TFIDF, the CNN classifiers consistently outperform the SVM+TFIDF baseline on precision and F1 score. In particular, when $W = 5$ and $D = 800$, the CNN classifier achieves the best scores on all the metrics, which shows the effectiveness of convolution neural networks with word embeddings in the Twitter election classification task. In order to validate whether the best CNN classifiers significantly outperforms the best baseline SVM+TFIDF, the non-parametric McNemar’s test is used to conduct a statistical test as suggested by Dietterich [8] for a reliable and computational inexpensive comparison. Our statistical test result shows that the two-tailed p -value is 0.0042, which means the difference between CNN and SVM+TFIDF is considered to be statistically significant.

In Table 5(b), where both approaches use word embeddings, we observe that SVM+WE and CNN show different preferences in word embeddings dimensionality. When using SVM+WE, a smaller dimension size and larger context window size (for example $W5$ and $D200$) give a better performance on F1 score and precision. However, the CNN classifier prefers both large context window size and dimension size. Therefore, when using a large context window size, word embeddings with higher dimensionality are likely to have a better performance (for example $W5$ and $D800$). The simple scheme used in SVM+WE is problematic with

high dimensional word embeddings. Simply combining all the word vectors of a Twitter post may excessive ambiguity about the topic of the post, particularly as not all the words are meaningful for classification. Hence, this scheme may hurt the semantic representation [20]. As the dimensionality increases, this could introduce further ambiguities and lead to degraded performance in our Twitter election classification task. Nevertheless, results of both SVM+WE and CNN suggest that a higher context window size is most appropriate for our task.

Compared to the studies on other tasks such as named entity recognition (NER) and dependency parsing (see Section 2), our results differ from their conclusions that “a smaller context window size gives a better performance” [3, 10]. Such a contradiction suggests that the best setup of parameters such as context window and dimensionality might differ from a task to another. In summary, for the Twitter election classification task using CNNs, word embeddings with a large context window and dimension size can achieve statistically significant improvements over the most effective classification baseline of SVM with TF-IDF.

7. CONCLUSION

Since previous investigations on the parameter configuration of word embeddings focus on different tasks such as NER [10] and dependency parsing [3], their findings may not generalise to Twitter classification tasks. Meanwhile, similar work on Twitter classification tasks [9, 17, 19] have not studied the impact of background corpora and *Word2Vec* parameters such as context window and dimensionality. Our finding shows that these two factors could affect the classification performance on Twitter classification tasks. Based on experiments on a Twitter election dataset, this paper studies word embeddings when using convolutional neural networks. Using two different types of background corpora, we observe when the type of background corpus aligns with the classification dataset, the CNN classifier can achieve a better performance. In particular, our investigation shows that choosing the correct type of background corpus can potentially cover more vocabulary of the classification dataset. Thus, the alignment between the type of background corpus and classification dataset provides better tweet-level representations. For inferring the best setup of *Word2Vec* parameters (e.g. context window and dimensionality), we applied word embeddings with various parameter setup to convolutional neural networks. As a practical guide for a Twitter classification task, word embedding with both large context window and dimension is preferable with a CNN classifier for a better performance.

8. ACKNOWLEDGMENTS

This paper was supported by a grant from the Economic and Social Research Council, (ES/L016435/1).

9. REFERENCES

- [1] Venezuela opposition politician luis manuel diaz killed. <http://www.bbc.co.uk/news/world-latin-america-34929332>, November 2015. [Accessed: 2016-05-15].
- [2] G. Amati, G. Amodeo, M. Bianchi, G. Marcone, F. U. Bordoni, C. Gaibisso, G. Gambosi, A. Celi, C. Di Nicola, and M. Flammini. FUB, IASI-CNR,

(a) Results of random classifier, SVM with TF-IDF (SVM+TFIDF) and SVM with word embeddings (SVM+WE)

	Precision	Recall	F1 score
Random	38.6	28.5	38.5
SVM+TFIDF	76.0	73.7	74.8
SVM+WE*	77.4	68.4	72.6

(b) Results of SVM with word embedding (SVM+WE) and CNN

		D200			D500			D800											
		SVM+WE			CNN			SVM+WE			CNN								
		P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1			
W5	W1	75.9	66.8	71.0	81.2	71.6	76.1	74.2	68.8	71.4	80.7	71.9	76.0	71.7	69.8	70.7	80.8	71.6	75.9
	W3	77.5	67.7	72.2	81.2	72.0	76.3	75.4	68.9	72.0	80.9	71.6	75.9	71.8	68.9	70.3	81.3	72.2	76.5
	W5	77.4	68.4	72.6	82.0 [†]	71.5 [†]	76.4 [†]	74.3	70.1	72.1	80.9 [†]	72.2 [†]	76.3 [†]	71.5	69.3	70.4	83.7[†]	74.1[†]	77.1[†]

Table 5: Results of our baselines and CNN models in Twitter election classification task. W1 means context window size 1 and D200 denotes word embeddings dimension size 200. * denotes the setting of SVM+WE that exhibits the highest F1 score (W5 and D200) from (b). † indicates that the result is statistically significant compared to the best baseline SVM+TFIDF.

- UNIVAQ at TREC 2011 microblog track. In *Proc. of TREC*, 2011.
- [3] M. Bansal, K. Gimpel, and K. Livescu. Tailoring continuous word representations for dependency parsing. In *Proc. of the 52nd ACL conference*, volume 2, pages 809–815, 2014.
- [4] Y. Bengio, R. Ducharme, P. Vincent, and C. Janvin. A neural probabilistic language model. *Journal of machine learning research*, 3:1137–1155, 2003.
- [5] A. Birmingham and A. F. Smeaton. On using Twitter to monitor political sentiment and predict election results. In *Proc. of SAAIP workshop at IJCNLP*, 2011.
- [6] R. Collobert and J. Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proc. of the 25th ICML*, pages 160–167, 2008.
- [7] R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuksa. Natural language processing (almost) from scratch. *Journal of machine learning research*, 12:2493–2537, 2011.
- [8] T. G. Dietterich. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural computation*, 10(7):1895–1923, 1998.
- [9] S. Ebert, N. T. Vu, and H. Schütze. CIS-positive: Combining convolutional neural networks and SVMs for sentiment analysis in Twitter. In *Proc. of the SemEval workshop*, page 527, 2015.
- [10] F. Godin, B. Vandersmissen, W. De Neve, and R. Van de Walle. Multimedia Lab@ ACL W-NUT NER shared task: Named entity recognition for Twitter microposts using distributed word representations. In *Proc. of the ACL-IJCNLP conference*, page 146, 2015.
- [11] Y. Kim. Convolutional neural networks for sentence classification. In *Proc. of EMNLP conference*, pages 1746–1751, 2014.
- [12] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [13] C. Macdonald, R. McCreddie, R. L. Santos, and I. Ounis. From puppy to maturity: Experiences in developing terrier. In *Proc. of the OSIR workshop at SIGIR*, volume 60, 2012.
- [14] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *Proc. Advances in neural information processing systems*, pages 3111–3119, 2013.
- [15] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of machine learning research*, 12:2825–2830, 2011.
- [16] M. Sanderson. Test collection based evaluation of information retrieval systems. *Foundations and Trends in Information Retrieval*, 2010.
- [17] A. Severyn and A. Moschitti. UNITN: Training deep convolutional neural network for Twitter sentiment classification. In *Proc. of the 9th SemEval workshop*, pages 464–469, 2015.
- [18] A. Severyn, M. Nicosia, G. Barlacchi, and A. Moschitti. Distributional neural networks for automatic resolution of crossword puzzles. In *Proc. of ACL-IJCNLP conference*, 2015.
- [19] D. Tang, F. Wei, N. Yang, M. Zhou, T. Liu, and B. Qin. Learning sentiment-specific word embedding for Twitter sentiment classification. In *Proc. of the 52nd ACL conference*, volume 1, pages 1555–1565, 2014.
- [20] P. Wang, J. Xu, B. Xu, C.-L. Liu, H. Zhang, F. Wang, and H. Hao. Semantic clustering and convolutional neural network for short text categorization. In *Proc. of the 53rd ACL-IJCNLP conference*, volume 2, pages 352–357, 2015.