

Political discourse classification in social networks using context sensitive convolutional neural networks

Aritz Bilbao-Jayo

DeustoTech - Fundación Deusto
Avda Universidades, 24,
48007, Bilbao.
aritzbilbao@deusto.es

Aitor Almeida

DeustoTech - Fundación Deusto
Avda Universidades, 24,
48007, Bilbao.
aitor.almeida@deusto.es

Abstract

In this study we propose a new approach to analyse the political discourse in on-line social networks such as Twitter. To do so, we have built a discourse classifier using Convolutional Neural Networks. Our model has been trained using election manifestos annotated manually by political scientists following the Regional Manifestos Project (RMP) methodology. In total, it has been trained with more than 88,000 sentences extracted from more than 100 annotated manifestos. Our approach takes into account the context of the phrase in order to classify it, like what was previously said and the political affiliation of the transmitter. To improve the classification results we have used a simplified political message taxonomy developed within the Electronic Regional Manifestos Project (E-RMP). Using this taxonomy, we have validated our approach analysing the Twitter activity of the main Spanish political parties during 2015 and 2016 Spanish general election and providing a study of their discourse.

1 Introduction

OSN-s are a commonplace element in most citizens daily lives. A significant amount of the social engagement (com, 2015) between citizens takes place in the OSN-s. The same trend is taking place in the political sphere. The on-line presence of political parties and public servants has increased dramatically in the last decades. Political campaigns include an on-line component and politicians use the OSN-s as another medium for their political discourse (Almeida and Orduna, 2017). As a result, the content of the OSN-s can be used

to analyse different aspects of the political activity. OSN activity can serve as an input to study the possible results of political campaigns (Kalampokis et al., 2017)(Ortiz-Ángeles et al., 2017), to generate profiles (Grčić et al., 2017) of the politicians according to their OSN usage or to analyse their reactions to certain events or topics (Güneyli et al., 2017).

To take advantage of the political data available in the OSN-s, we present in this paper a deep neural network architecture for political discourse analysis. Our architecture takes advantage of the context of the political discourse (what was previously said and who was the transmitter) to improve the classification process. To do so, we have used the annotated political manifestos database created by the Regional Manifestos Project (RMP) (Alonso et al., 2013). To improve the classification we use the simplified taxonomy that have been developed within the Electronic Regional Manifestos Project (E-RMP), which adapts the initial RMP taxonomy to the political discourse analysis in OSN-s. Using this new taxonomy and the created deep neural network architecture we have analysed the discourse during the electoral campaigns of the 2015 and 2016 Spanish general elections.

This paper is organized as follows. In Section 2 we analyse the previous work done in the area of automatic political discourse analysis in social networks. In Section 3 we describe the classification taxonomy that we have used for the analysis of the political discourse. In Section 4 we present our neural network architecture for political discourse classification. In Section 5 we discuss the evaluation of the system. In Section 6 we offer a real use case of the presented system by analysing the political activity on Twitter during the 2015 and 2016 general elections in Spain. Finally, section 7 draws some conclusions and proposes fur-

ther work.

2 Related Work

2.1 Automated use of political manifestos

The automated use of annotated political manifestos as basis for the analysis of other types of political texts besides political manifestos has not been a remarkable research area until recently.

(Nanni et al., 2016) used annotated political manifestos and speeches to analyse the speeches from the last 3 US presidential campaigns in the 7 main domains defined by the manifestos project. The main difference between Nanni et al.'s work and our research is that first, we only use annotated manifestos as training data (while Nanni et al. used annotated speeches too) to later apply this knowledge to another areas such as social networks, and second, this work is applied to analyse the political discourse on social networks and not on political speeches. Moreover, this is the first time that annotated manifestos are used as basis for a political discourse analysis on Twitter to the best of our knowledge.

2.2 Political analysis on Twitter

Since its inception, Twitter has been seen by researchers of several fields as a new source of information where they can conduct their researches. For instance, political scientists have identified Twitter as a platform where they can analyse what a subset of the population says without performing expensive surveys.

Several researchers have measured the predictive power of social networks such as Twitter. (Tumasjan et al., 2010) claimed after analysing more than 100,000 tweets from the 2009th German federal election, that the mere number of messages mentioning a party reflects the election results. Furthermore, (O'Connor et al., 2010) measured the potential Twitter messages might have as a substitute of traditional polling. After using some basic sentiment analysis techniques, O'Connor et al. concluded that a simple sentimental analysis on top of Twitter data produces similar results to polls.

However, there have been diverse criticisms regarding the predictive power of Twitter. For instance, (Gayo Avello et al., 2011) replicated Tumasjan et al.'s and O'Connor's approaches utilising a set of tweets about the 2010 United States House of Representatives elections, obtaining a

mean average error of 17.1% compared to election's real results.

The analysis of political polarization in social networks has also been an important research field in political social network analysis. To do so, one of the principal approaches is to construct the graph representation of the social network and apply some network theory principles. On one hand, (Conover et al., 2011) used a combination of community detection algorithms and manually annotated data to analyse the polarity of two networks constructed after gathering more than 250,000 tweets about 2010 U.S congressional midterm elections. The first network represented the retweets and the second one the mentions between different users. Conover et al. concluded that users tend to retweet tweets of users they agree with. Therefore, communities are evident in the retweet network. However, in the mentions network there were more interactions between people with different political ideas, suggesting the existence of discussions between different polarities.

On the other hand, (Finn et al., 2014) introduced a new approach for the measurement of the polarity using a co-retweeted network. The approach was tested with the most retweeted 3,000 tweets within their dataset. Authors concluded that by using their co-retweeted network were able to measure the polarity of the most important accounts participating in the discussion and the polarity of the analysed event.

Other researchers have detected the polarity of raw text using natural language processing techniques. (Iyyer et al., 2014) using recursive neural networks and (Rao and Spasojevic, 2016) using word embeddings and Long Short-Term Memory (LSTM) in order to identify the political polarity of a sentence.

3 Regional Manifestos Project Annotation Taxonomy

Political scientists have been manually annotating political parties' manifestos for years in order to apply content analysis methods and perform political analyses later on.

The precursors of this methodology were the Manifesto Project, formerly known as the Manifesto Research Group (MRG) and Comparative Manifestos previously (CMP)(Budge, 2001). In 2001, they created the Manifesto Coding Hand-

book(Volkens, 2002) which has evolved over the years. The handbook provides instructions to the annotators about how political parties' manifestos should be coded for later content analysis and a category scheme that indicates the set of codes available for codification. Nowadays, the category scheme for manifestos annotation consists in 56 categories grouped into seven major policy areas (all the categories are available in ¹): *external relations, freedom and democracy, political system, economy, welfare and quality of life* and *social groups*.

Moreover, other manifestos annotation projects such as the RMP (the project to which the dataset we have used in this research belongs to) extended the original annotation to address some other political preferences. In particular, they extended the *centralization, decentralization* and *nationalism* categories in order to perform a deeper analysis of those political phenomena. To do so, they added some new categories to the Manifestos Project category schema, increasing the number of categories from 56 to 78 (the codebook is available at ²).

However, due to the high number of available categories for annotation, it has been proven that manifestos annotation is not an easy task even for trained political scientists as Mikhaylov et al. demonstrated in (Mikhaylov et al., 2012). The authors concluded after examining diverse annotators' intercoder reliability in two preselected manifestos, that the codification process is highly prone to misclassification due to the large number of categories.

To address the problem that annotating political manifestos is not an easy task even for trained annotators with a codification specifically designed for political manifestos, and to adapt the taxonomy to the political discourse analysis in OSN-s, the E-RMP has developed a simplified taxonomy. This new taxonomy has been created redistributing some of the subdomains of the RMP into new 7 categories: external relations, welfare, economy, democratic regeneration, territorial debate, immigration and boasting. The new distribution of subdomains can be seen in Table 1 and it has been designed in order to analyse European politics. Each of the categories would mean the following:

¹https://manifesto-project.wzb.eu/coding_schemes/mp_v5

²<http://www.regionalmanifestosproject.com>

- External Relations: references regarding the position/status of the country inside the European Union.
- Welfare: references to welfare state, equality, education, public health, etc.
- Economy: references to any economic sphere of the country.
- Democratic Regeneration: references to the state of democracy, political corruption and new mechanisms of democratic participation.
- Territorial Debate: references to the distribution of power between the state and lower level governments, patriotism, nationalism, pro-independence movements, etc.
- Immigration: references to how immigration should be handled in the country.
- Boasting: references to the speaking party's competence to govern or other party's lack of such competence.

4 Neural Network Architecture for Political Classification

In order to accomplish the text classification task we have opted for convolutional neural networks with Word2Vec word embeddings. Recently, CNNs have achieved excellent results in several text classification tasks (Kim, 2014) and it has been proven their great performance with tweets too(Severyn and Moschitti, 2015).

The inputs of our model are the sentences which are fed to the neural network as sequences of words. These sequences have a maximum length of 60 words (the maximum length have been decided after an analysis of our corpus' sentences' length). Then, this words are mapped to indexes (1, ..., |D|) in a dictionary, being D the number of unique words in the corpus and using the 0 index for padding purposes. After, an embedding layer transforms the word indexes to their corresponding Word2Vec word embeddings. We have opted for the non-static or trainable embedding layer since it improves model's performance. The used Word2Vec model embedding's size is 400 and it has been trained with a corpus of Spanish raw text of 3 billion words(Almeida and Bilbao, 2018).

External Relations	Economy
European Integration: Positive	Nationalisation: Positive
European Integration: Negative	Controlled Economy: Positive
Democratic regeneration	Protectionism: Positive
Democracy	Keynesian Demand Management: Positive
Constitutionalism: Positive	Economic Planning: Positive
Representative democracy: Positive	Free-Market Economy: Positive
Participatory democracy: Positive	Economic Orthodoxy: Positive
Political Corruption: Negative	Corporatism: Positive
Immigration	Management of natural resources
Equal treatment of immigrants	Market Regulation: Positive
Welfare expansion for immigrants	Economic Goals
Welfare limitations for immigrants	Incentives: Positive
Education expansion for immigrants	Economic Growth
Education limitation for immigrants	Technology and Infrastructure: Positive
Immigrants' negative impact on law and order	Labour Groups: Positive
Multiculturalism: Positive	Labour Groups: Negative
Territorial Debate	Multiculturalism: Negative
Decentralisation: Positive	Boasting
Centralisation: Positive	Political Authority
Regional finance: Positive	Welfare
Differential treatment among regions: Negative	Welfare state expansion
Differential treatment among regions: Positive	Education expansion
National Way of Life: Positive	Equality: Positive
Promotion and protection of vernacular languages	
Cultural links with diaspora	
Bilingualism positive	
National way of life: Negative	
Immigrants: Positive	

Table 1: Proposed taxonomy

Once the input phrase has been converted into a sequence of word vectors, the phrase can finally be fed into the convolutional neural network, since the sequence of word vectors are in fact a matrix which dimensions are $60 \times d$ where d is the embedding size. Then, the model performs convolution operations with 3 different filter sizes, batch normalization (Ioffe and Szegedy, 2015) and ReLU as the activation function. Batch normalization acts as an extra regularizer and increases the performance of the model.

As it can be seen in Figure 1, the defined filter sizes are $2 \times d$, $3 \times d$ and $4 \times d$. In other words, these filter sizes define the sizes of the n-grams which in this case are 2-grams, 3-grams and 4-grams respectively. For example, a filter size of $2 \times d$ will take the whole width of all the possible bigrams of the sentence.

Moreover, as it is stated in (Zhang and Wallace,

2015), multiple filters should be used in order to learn complementary features. Therefore, the proposed model has 100 filters per different filter size. Once a filter has been applied, a feature map is generated. Thus, a different feature map is generated per applied filter as it can be seen in Figure 1, where there are 3 filters instead of 100 for explanatory purposes.

After the convolutional neural networks a pooling layer reduces the dimensionality of the incoming data. There are several pooling strategies, however we have opted for the 1-max-pooling (Boureau et al., 2010) strategy since it has been proved in (Zhang and Wallace, 2015) that is the best approach for natural language processing tasks. It captures the most important feature (the highest value) from each of the feature maps. Therefore, the pooling operation outputs a feature per filter which is later concatenated into a feature

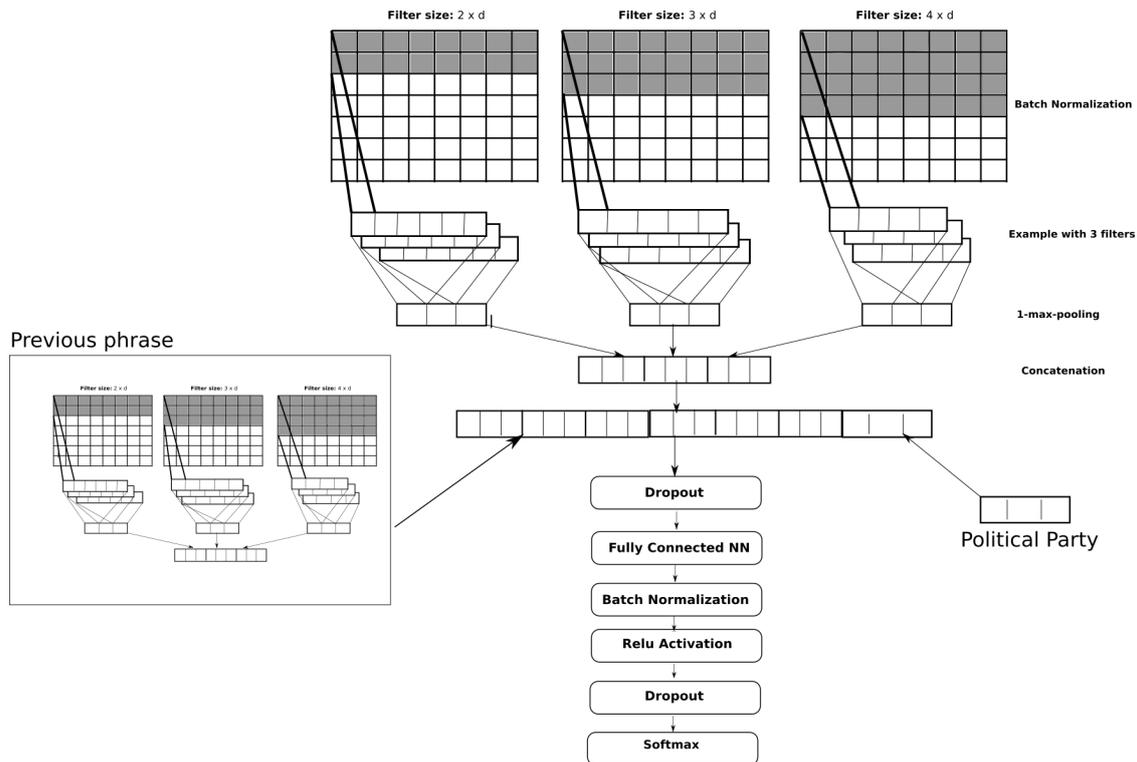


Figure 1: Designed architecture

vector.

Next, a dropout (Srivastava et al., 2014) rate of 0.5 is applied as regularization in order to prevent the network from over-fitting, followed by a fully connected layer with ReLU as the activation function and batch normalization. Then a 0.5 dropout is applied. Finally, the softmax function computes the probability distribution over the labels.

The categorical cross-entropy loss has been used as training objective function since it supports multiclass classifications. Regarding the optimizer, the optimization has been performed using Adam (Kingma and Ba, 2014) with the parameters used in the original manuscript for classification problems.

4.1 Contextual data as new inputs

Two different approaches have been tested in order to insert the previous phrase as an extra input: 1) As a second channel in the convolutional layers. When convolution operations are applied to text only one channel is used. Here we propose the use of an extra channel for the previous context; 2) Replicating for the previous phrase the same convolution-pooling process used for the phrase being classified (see Figure 1).

Regarding the political party, we have decided

to represent each political party with a one-hot-encoding representation and concatenate it to the feature maps obtained after the convolutions (see figure 1).

5 Evaluation

The experimentation performed in this research work has been done with the dataset provided by the Regional Manifestos Project, which has a high annotators' intercoder reliability (Alonso et al., 2013). This dataset has almost two decades of political manifestos in Spain and therefore covers a wider span of political issues with a high language variation. The dataset consists in 88,511 annotated phrases and the distribution of codes is highly imbalanced: External Relations (0.9%), Welfare (35.91%), Economy (47.83%), Democratic Regeneration (4.38%), Immigration (1.77%), Territorial debate (7.81%), Boasting (1.3%). Almost 85% of the dataset belongs to Welfare and Economy categories, leaving around the 15% of the dataset for the remaining 5 categories.

In order to evaluate our approach, we have divided our dataset in 2 different subsets: training and validation sets (85%), and test set (15%). The training and validation set has been used in order to create models with 5-fold cross validation to

Experiment	Accuracy	F1(Macro)
E1	83.79%	69.19
E2	83.8%	69.64
E3	86.36%	72.58
E4	87.63%	75.29
E5	87.55%	74.68

Table 2: Results with political manifestos.

later test their performance with the same test set. The reason why we have split the dataset in 2 subsets and then apply cross-validation to one of them is because we have used early stopping (Prechelt, 1998) in order to stop our model’s training when it started to over-fit. Early stopping compares the training accuracy with the validation accuracy and after some epochs without any improvements in the validation accuracy it stops the training. Nevertheless, the model may have over-fitted with respect to the validation set, therefore, a third set (test set) is needed in order to measure the real performance of the model. Furthermore, since we work with an imbalanced dataset, we have applied stratification in order to preserve the same percentage of samples for each class. Using this approach we are able to evaluate how each class is classified since it ensures that in each of the subsets there will be a representation of each class. Taking into account both the high number of classes and the imbalance between them, we have used the f-measure as the evaluation metric. Additionally we also provide the accuracy of each experiment.

We have performed five different experiments to analyse the importance of the context (both the what was said previously and who is saying it) when classifying the political discourse: 1) Only the sentence to be classified with no additional context (*E1*); 2) the sentence plus the political party who belongs to (*E2*); 3) the sentence plus the previous sentence in an additional channel on the CNNs (*E3*); 4) the sentence plus the previous sentence in another CNNs structure, concatenating the features extracted by both networks (*E4*); and 5) the sentence, the political party who belongs to and the previous sentence in another CNN(*E5*).

As it is shown in table 2, the performance of the classifiers improves when adding the previous sentence and the political party as extra features. On the one hand, the previous sentence provides a remarkable increase in accuracy and F1 when it is inserted as an additional channel on the CNNs

Experiment	Accuracy	F1(Macro)
T1	67.57%	57.07
T2	69.05%	63.06
T3	66.33%	60.26
T4	70.59%	63.17

Table 3: Results with annotated tweets.

(*E3*) and as as a new structure of CNNs (*E4*). However, the improvement in *E4* is greater than in *E3*. On the other hand, adding the political party who says the phrase as an extra feature (*E2*) improves the F1 in 0.45 points compared with the baseline (*E1*). With regard to *E5*, since combining party and previous phrase does not improve the results of *E4*, we can affirm that those two features are not complementary.

Additionally, we have also tested the performance of our model on Twitter. To do so, we have tested the aforementioned models in a dataset of 404 manually annotated tweets. The category distribution of the test set is the following one: external relations (0.74%), welfare(33.66%), economy(30.69%), democratic regeneration(14.35%), immigration(0.49%), territorial debate(16.58%), boasting(3.46%).

It is important to remark that these models have been trained using the annotated manifestos from the Regional Manifestos Project dataset, without using any tweet during the training process.

We have performed four different experiments to analyse the performance of the previously explained architecture when classifying manually annotated tweets: 1) Only the tweet to be classified with no additional context and a Word2Vec model generated with generic Spanish text (*T1*); 2) the tweet to be classified with no additional context and a Word2Vec model generated with generic Spanish text and on-line trained with the tweets of our Spanish elections dataset (*T2*); 3) the tweet to be classified with the tweet it is answering to in another CNNs structure and a Word2Vec model generated with generic Spanish text (*T3*); 4) the tweet to be classified with the tweet it is answering to in another CNNs structure and a Word2Vec model generated with generic Spanish text and on-line trained with the tweets of our Spanish elections dataset (*T4*).

As it can be seen in table 3, retraining the Word2Vec model with tweets of our Spanish elections dataset significantly increases the accuracy

and F-measure of the model. On the one hand, from T1 to T2 there is an improvement of 2.5 points in accuracy and 6 points in F1. On the other hand, from T3 to T4 there is an improvement of 4 points in accuracy and 3 points in F1. With regard to the use of the previous tweet in the thread, it improves the accuracy of the model in 1.5 points.

6 Use Case

To demonstrate the usefulness of our system, we present a possible use case scenario for our classification model: to analyse the political discourse of the Spanish political parties and candidates during the campaign period of the 2015 and 2016 Spanish general elections on Twitter. In Spain, general elections should be held every 4 years. However, after the results of 2015 Spanish general elections neither of the two most voted parties were capable of obtaining the necessary support to form a government. Therefore, after months of unsuccessful negotiations new general elections were called.

The performed analysis consists in classifying the tweets written by the political parties standing for elections in the previously mentioned 7 categories to later analyse how some political parties prioritise some categories over others. To do so, we gathered from 4th to 18th of December (the 2015 general election was held on the 20th of December) (Almeida et al., 2015) and from 10th to 24th of June (the 2016 general election was held on the 26th of June) (Almeida et al., 2016) all the tweets written by the political parties and candidates standing for election. We gathered more than 80,000 tweets (taking into account both elections) from more than 10 different political parties and their respective candidates.

In order to perform the political discourse analysis, we used the previously mentioned classification model to distribute the tweets from 5 political parties (ignoring retweets) in the 7 categories previously defined. The analysed political parties are:

- People's Party (PP): right-wing, conservative political party. PP had been the ruling party between 2011-2015 having an absolute majority in Parliament.
- Spanish Socialist Workers' Party (PSOE): social democratic, centre-left political party. PSOE had been the ruling party between (2004-2011) when due to the financial crisis

and the high unemployment rate in Spain lost the 2011 Spanish general elections.

- Podemos - We Can: left-wing political party. The party was founded in 2014 and their main objectives were to address unemployment, inequality, corruption and austerity problems.
- Citizens: centre, liberal political party. Even though it was founded in 2006 as regional party in Catalonia, the party started to have influence at national level in the end of 2014.
- Basque Nationalist Party (PNV): centre-right, Christian democratic, Basque nationalist party.

6.1 2015 general elections political discourse analysis

In figure 2, the distribution of the tweets of the 5 analysed Spanish political parties over the 7 categories is shown. On the one hand, the first worth mentioning aspect is how *Boasting* is the dominant category on the 4 main political parties running for the 2015 general elections in all regions of Spain (People's Party, Spanish Socialist Workers' Party - PSOE, Podemos - We Can, Citizens). Moreover, it is also remarkable that People's party, the ruling party when the elections were held, is the political party with the highest percentage in *Boasting*. On the other hand, the Basque Nationalist Party (PNV) focuses its discourse on *Territorial Debate* category. This category includes topics such as the distribution of power between state and lower level governments (Basque Nationalists want more autonomy for their region), promotion and protection of vernacular language such as Basque, bilingualism (in Basque Country there are two official languages: Spanish and Basque) or nationalism which in this case would be Basque nationalism.

It is also noteworthy how differently the two main Spanish political parties (PP and PSOE) prioritised *Welfare* category. The low interest shown by the People's Party on *Welfare* may be due to the austerity measures taken and the performed cutbacks in the welfare state and social protection during their period as the ruling party. Therefore, it would make sense to assume that PSOE (the first opposition party) could see this as an opportunity to take advantage to differentiate themselves from the People's Party. However, People's Party is not the political party which has talked less about *Welfare* and *Quality of Life*. As it can be seen in Fig-

ure 2, Citizens talks even less about Welfare and Quality of Life which may be related to their liberal ideology.

With regard to *Democratic regeneration*, it is clearly seen in figure 2 that mainly Citizens, but also Podemos- We can and Spanish Socialist Workers' Party - PSOE, gave a high importance to this category, unlike PP. *Democratic regeneration* encompasses concepts such as calls for constitutional amendments or changes, favourable mentions of the system of direct democracy, the need of involvement of all citizens in political decision-making, division of powers, independence of courts, etc. These concepts were introduced in Spanish politics after 2011 15-M Movement (Hughes, 2011), and continued to gain in importance during the legislature, being one of the main topics the parties on the opposition addressed during their campaign.

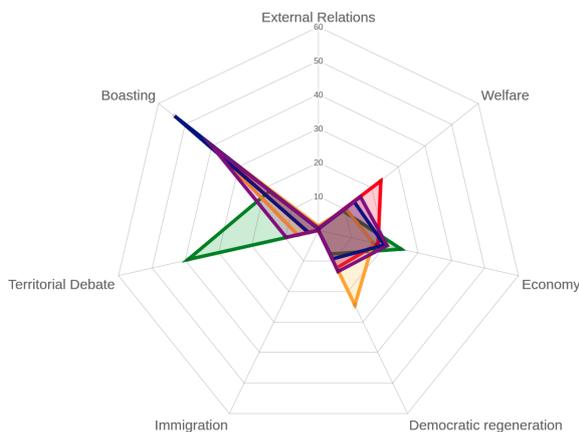


Figure 2: Distribution among 7 categories of the tweets created by the Twitter accounts of PSOE (red), PP (Blue), Podemos - We Can (Purple), Citizens (Orange) and PNV (Green) in 2015 Spanish general elections

6.2 2016 general elections political discourse analysis

One relevant change in the 2016 elections political discourse in Twitter is the use of *External Relations* category. In the previous elections this domain was ignored by all the political parties. However, as it can be seen in figure 3, People's Party and Citizens emphasized more this category than in the previous general elections. This could have happened due to Brexit.

With respect to the rest of categories, it is noteworthy how the 4 main political parties gave

less importance to *Boasting* category in favour of *Democratic regeneration* and *Economy*.

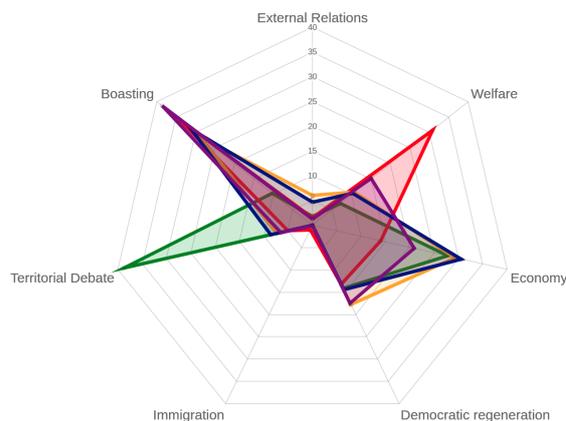


Figure 3: Distribution among 7 categories of the tweets created by the Twitter accounts of PSOE (red), PP (Blue), Podemos - We Can (Purple), Citizens (Orange) and PNV (Green) in 2016 Spanish general elections

7 Conclusions and Future Work

In this paper we present a model, based in a convolutional neural network architecture, which takes advantage of the context to classify the political discourse in OSN-s. The political discourse classification is based in a simplified taxonomy developed within the Electronic Regional Manifestos Project, which has been created to be applied specifically to OSN-s. To demonstrate the utility of our model we have used it to analyse the Twitter activity of the main political parties during the 2015 and 2016 Spanish general elections. The proposed model can be easily retrained to work in other languages, using the for example the dataset of the Manifesto Project³, which provides annotated manifestos in several languages.

As future work, we would like to study how attention mechanisms (Hermann et al., 2015) could be used to improved the classification process, in order to obtain better results. We would also like to take advantage of the inner representation created by the capsule networks(Sabour et al., 2017) to create vectors that represent each one of the target categories, in order to use them for the classification.

³<https://manifesto-project.wzb.eu/>

Acknowledgments

We gratefully acknowledge the support of the Basque Government’s Department of Education for the predoctoral funding; the Ministry of Economy, Industry and Competitiveness of Spain under Grant No. CSO2015-64495-R (Electronic Regional Manifestos Project); and NVIDIA Corporation with the donation of the Titan X used for this research. We thank the Regional Manifestos Project team (Braulio Gómez Fortes and Matthias Scantamburlo) for making available their dataset of annotated political manifestos and tweets.

References

2015. Media metrix cross-platform. Technical report, comScore.
- Aitor Almeida and Aritz Bilbao. 2018. [Spanish 3b words word2vec embeddings](#).
- Aitor Almeida, Pablo Orduña, and Aritz Bilbao. 2015. [Party and candidate tweets for the campaign period of the 2015 spanish general election](#).
- Aitor Almeida, Pablo Orduña, and Aritz Bilbao. 2016. [Party and candidate tweets for the campaign period of the 2016 spanish general election](#).
- Aitor Almeida and Pablo Orduna. 2017. Analyzing political discourse in on-line social networks j. ucs special issue. *Journal of Universal Computer Science*, 23(3):233–235.
- Sonia Alonso, Braulio Gómez, and Laura Cabeza. 2013. Measuring centre–periphery preferences: The regional manifestos project. *Regional & Federal Studies*, 23(2):189–211.
- Y-Lan Boureau, Jean Ponce, and Yann LeCun. 2010. A theoretical analysis of feature pooling in visual recognition. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 111–118.
- Ian Budge. 2001. *Mapping policy preferences: estimates for parties, electors, and governments, 1945-1998*, volume 1. Oxford University Press on Demand.
- Michael Conover, Jacob Ratkiewicz, Matthew R Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. 2011. Political polarization on twitter. *ICWSM*, 133:89–96.
- Samantha Finn, Eni Mustafaraj, and P Takis Metaxas. 2014. The co-retweeted network and its applications for measuring the perceived political polarization.
- Daniel Gayo Avello, Panagiotis T Metaxas, and Eni Mustafaraj. 2011. Limits of electoral predictions using twitter. In *Proceedings of the Fifth International AAI Conference on Weblogs and Social Media*. Association for the Advancement of Artificial Intelligence.
- Klara Grčić, Marina Bagić Babac, and Vedran Podobnik. 2017. Generating politician profiles based on content analysis of social network datasets. *Journal of Universal Computer Science*, 23(3):236–255.
- Ahmet Güneyli, Metin Ersoy, and Sevki Kiralp. 2017. Terrorism in the 2015 election period in turkey: Content analysis of political leaders’ social media activity. *J. UCS*, 23(3):256–279.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems*, pages 1693–1701.
- Neil Hughes. 2011. young people took to the streets and all of a sudden all of the political parties got old: The 15m movement in spain. *Social Movement Studies*, 10(4):407–413.
- Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.
- Mohit Iyyer, Peter Enns, Jordan Boyd-Graber, and Philip Resnik. 2014. Political ideology detection using recursive neural networks. In *Proceedings of the Association for Computational Linguistics*, pages 1113–1122.
- Evangelos Kalampokis, Areti Karamanou, Efthimios Tambouris, and Konstantinos A Tarabanis. 2017. On predicting election results using twitter and linked open data: The case of the uk 2010 election. *J. UCS*, 23(3):280–303.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Slava Mikhaylov, Michael Laver, and Kenneth R Benoit. 2012. Coder reliability and misclassification in the human coding of party manifestos. *Political Analysis*, 20(1):78–91.
- Federico Nanni, Cäcilia Zirn, Goran Glavas, Jason Eichorst, and Simone Paolo Ponzetto. 2016. Topfish: Topic-based analysis of political position in us electoral campaigns.
- Brendan O’Connor, Ramnath Balasubramanyan, Bryan R Routledge, and Noah A Smith. 2010. From tweets to polls: Linking text sentiment to public opinion time series. *ICWSM*, 11(122-129):1–2.

- Sonia Ortiz-Ángeles, Yenny Villuendas-Rey, Itzamá López-Yáñez, Oscar Camacho-Nieto, and Cornelio Yáñez-Márquez. 2017. Electoral preferences prediction of the yougov social network users based on computational intelligence algorithms. *J. UCS*, 23(3):304–326.
- Lutz Prechelt. 1998. Early stopping-but when? In *Neural Networks: Tricks of the trade*, pages 55–69. Springer.
- Adithya Rao and Nemanja Spasojevic. 2016. Actionable and political text classification using word embeddings and lstm. *arXiv preprint arXiv:1607.02501*.
- Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. 2017. Dynamic routing between capsules. In *Advances in Neural Information Processing Systems*, pages 3859–3869.
- Aliaksei Severyn and Alessandro Moschitti. 2015. Twitter sentiment analysis with deep convolutional neural networks. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 959–962. ACM.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958.
- Andranik Tumasjan, Timm Oliver Sprenger, Philipp G Sandner, and Isabell M Welp. 2010. Predicting elections with twitter: What 140 characters reveal about political sentiment.
- Andrea Volkens. 2002. Manifesto coding instructions.
- Ye Zhang and Byron Wallace. 2015. A sensitivity analysis of (and practitioners’ guide to) convolutional neural networks for sentence classification. *arXiv preprint arXiv:1510.03820*.