

A Deep Modular RNN Approach for Ethos Mining

Rory Duthie¹, Katarzyna Budzynska^{2,1}

¹ Centre for Argument Technology, University of Dundee

² Centre for Argument Technology, Institute of Philosophy and Sociology, Polish Academy of Sciences
rwduthie@dundee.ac.uk

Abstract

Automatically recognising and extracting the reasoning expressed in natural language text is extremely demanding and only very recently has there been significant headway. While such argument mining focuses on *logos* (the content of what is said) evidence has demonstrated that using *ethos* (the character of the speaker) can sometimes be an even more powerful tool of influence. We study the UK parliamentary debates which furnish a rich source of ethos with linguistic material signalling the ethotic relationships between politicians. We then develop a novel *deep modular recurrent neural network*, DMRNN, approach and employ proven methods from argument mining and sentiment analysis to create an *ethos mining* pipeline. Annotation of ethotic statements is reliable and its extraction is robust (macro- $F1 = 0.83$), while annotation of polarity is perfect and its extraction is solid (macro- $F1 = 0.84$). By exploring correspondences between ethos in political discourse and major events in the political landscape through ethos analytics, we uncover tantalising evidence that identifying expressions of positive and negative ethotic sentiment is a powerful instrument for understanding the dynamics of governments.

1 Introduction

Arguments play a fundamental role in the workings of society, this is particularly evident in political discourse. Although headway has begun to be made through argument mining (*cf.* [Budzynska and Villata, 2016]), the focus here has been exclusively on *logos*, the content of what is said. Since antiquity [Aristotle, 1991], it has been recognised in rhetoric that ethos, the character of the speaker, can play a more important role in influencing beliefs, attitudes and actions of others. Politicians quickly learn that having sound or popular policies might not be as crucial in winning an election as having the ethos of a successful businessman, for example.

Following Aristotle’s definition, we specify ethos as properties of the individual or the group of agents which can be supported (Example 1, hereafter used as a running example) or attacked (Example 2) in order to influence the audience

through communication. Using transcribed UK parliamentary debates as an exploratory domain, we aim to extract the relationships between politicians or between a politician and a party through ethotic sentiment expressions, **ESEs**: when a politician talks about another politician’s or party’s ethos; and then to classify these relationships as having positive or negative sentiment through **+/-ESEs**: when a politician supports, **+ESE**, or attacks, **-ESE**, ethos of another politician or party. The intuition that we intend to model is that linguistic structure encodes both the target entity of the ethotic statement (in Example 1, “My hon. Friend,” and in Example 2, “the Government”) as well as the polarity of the ethotic statement (in Example 1, positive sentiment is signalled by “assiduously”, “pursuing” and “interests” and in Example 2, negative sentiment is signalled by “sick”). As a result, we are able to identify alliances and enmities between politicians.

Example 1 Mr. John Moore said, *My hon. Friend is assiduously pursuing his constituents’ interests.*

Example 2 Mr. Bruce Grocott said, *Is it not the simple truth that the Government are making the country sick?*

In order to automatically extract and classify ESEs, we develop an ethos mining pipeline featuring a deep modular recurrent neural network, **DMRNN**, approach applied for the first time to the NLP task of text classification. More specifically, we experiment with proven machine learning methods that boost accuracy (e.g. ensemble classification [Larkey and Croft, 1996]) and advances in neural networks through deep learning (e.g. recurrent neural networks, RNNs [Hochreiter and Schmidhuber, 1997]). We then create the DMRNN approach which provides several textual based inputs to RNN models, some featuring long short-term memory (LSTM) layers, and combine them into one model for binary classification of ESEs. The DMRNN allows for an increase in the available data that can be used for classification (due to the allowance of multiple separate modules) and for the application of multiple simple models with unique features from the inputs which, when combined, do not suffer from over-fitting.

The **contribution** of this paper is as follows: (1) the largest publicly available corpus (90,991 tokens) of ethotic relations between politicians or between politicians and political groups; (2) an ethos mining pipeline with two modules employing existing NLP techniques, two modules extending standard text classification methods, and four original

modules; (3) a DMRNN approach which is new to the area of text classification; (4) an automatically generated graph-based ethos analytics which, by processing the output of the ethos mining pipeline, provides an insight into the correspondence between political events and the expressed sentiment towards ethos of politicians involved in these events.

2 Related Work

We position our approach to related work with regard to both the method (deep learning) and the objective of study (automatic processing of discourse and reasoning). In **deep learning** (cf. [LeCun *et al.*, 2015]), RNNs have been used to create state of the art classifiers using word embeddings [Zhou *et al.*, 2015; Lee *et al.*, 2017] and LSTM layers [Sutskever *et al.*, 2014]. Other advances in text classification have been made with convolutional neural networks (CNNs) [Lai *et al.*, 2015] and by combing CNNs and RNNs from one input source in an ensemble [Wang *et al.*, 2017]. In image classification, modular methods are created combing CNNs trained on the same input data [Ju *et al.*, 2017] and on both image and text data [Ma *et al.*, 2016]. We adapt this successful modular method of image classification in our DMRNN approach.

Sentiment analysis and **opinion mining** (cf. [Liu, 2010]) are popular tools for processing discourse. In [Thomas *et al.*, 2006], a corpus of U.S. debates was compiled with added voting records to train a support vector machine (SVM) classifier and determine if a politician supports or opposes a piece of legislation. Although this does not encompass relationships between politicians, it is highlighted there that it would be an intriguing path to explore. In [Hirst *et al.*, 2014], it has been shown that the use of SVM classifiers to group politicians into parties (from the Canadian Hansard) on the basis of what they say is sensitive to whether the speech is aimed to attack or defend. Hirst *et al.* highlight that the majority of the oral question period consists of hostile speech between politicians which would make this an interesting space to extract ethos.

A new rapidly growing area of **argument mining** (cf. [Budzynska and Villata, 2016]) aims to automatically process reasoning. Although the focus here is solely on logos rather than ethos, it is worth noting that the introduction of deep learning techniques has already been explored in argument mining. In [Eger *et al.*, 2017], a dataset of 402 persuasive essays are used to extract arguments. Several deep neural network models are evaluated on their ability to identify different essay components, relations and a combination of the tasks with an LSTM which uses the dependency tree structure and gives *F1*-scores of 0.55 for 100% matches and 0.61 for 50% text span matches. In [Naderi and Hirst, 2017], the reputation argument strategies denial, excuse, justification and concession were classified, from Canadian Hansard, using an SVM and five-fold cross validation which gave *F1*-score 0.57.

Although the methods used for sentiment analysis and opinion and argument mining can be similar to those applied to **ethos mining**, it is important to highlight that the goal of the automatic extraction is different in each case and therefore results are not directly comparable. To the best of our knowledge, the only other work in the area of ethos mining is our previous work [Duthie *et al.*, 2016]. Train-

ing data is used there to manually craft domain specific rules which, along with a pipeline (containing part-of-speech tagging (POS), named entity recognition and reported speech removal), allows for the extraction of positive and negative ethotic statements in the UK Hansard with an *F1*-score of 0.70 (for more detailed comparison of results see Table 4 and 5). The improvements of the approach in this paper are: the size of the dataset (30% more tokens); the annotation scheme (for guidelines, see <http://arg.tech/GuideEthos>); and the NLP method used to extract ethotic relations (a general purpose DMRNN not tailored to specific domain rules).

3 Data

We construct our publicly available corpus, **Ethos_Hansard1** (http://arg.tech/Ethos_Hansard1, see Table 1), from the UK Hansard – UK parliamentary debate record (<http://hansard.millbanksystems.com/>) – by extracting 90 transcripts from Margaret Thatcher’s period as Prime Minister, dating from 1979 to 1990 (a volatile period in the UK). We split our data into 60 training transcripts, of which we use 10% as validation data, and 30 test transcripts to give a wide range of test cases. Each transcript, part of a day parliament sitting, contains an initial question asked by a member of parliament (MP) identified by their name and constituency, to which a reply is given by a government minister and further continued with subsequent turns in the debate.

Corpus	Sessions	Words	Segments (ESEs)	Speakers
Train	60	61,813	395	116
Test	30	29,178	243	82
TOTAL	90	90,991	638	198

Table 1: Summary of the language resources in the Ethos_Hansard1 corpus for mining ethos in the UK Hansard.

We annotate the transcripts for ethotic statements on a sentence level. ESEs contain a parliamentary entity (e.g. “hon. Member”, “Conservative party”) as a target of either an ethotic support or attack. The annotation scheme consists of four types of tag (see Table 2 for occurrences): (1) **speaker**, the author of an ESE; (2) **target**, the referent of an ESE; (3) **ethotic support, +ESE** (Example 1), when: (a) a single target entity, individual or group, can be established; and (b) it puts in a positive frame the target’s character or achievements, or supports their credibility explicitly; and (c) it is not self-referential; (4) **ethotic attack, -ESE** (Example 2), when: (a) a single target entity, individual or group, can be established; and (b) it puts in a negative frame the target’s character, attacking their credibility or associating them with events

Tags	#	κ
Speaker	149	1
Target	188	0.93
Ethotic Support (+ESE)	169	0.67
Ethotic Attack (-ESE)	469	1
TOTAL	975	-

Table 2: Frequency of tags in the Ethos_Hansard1 corpus and inter-annotator agreement (Cohen’s κ).

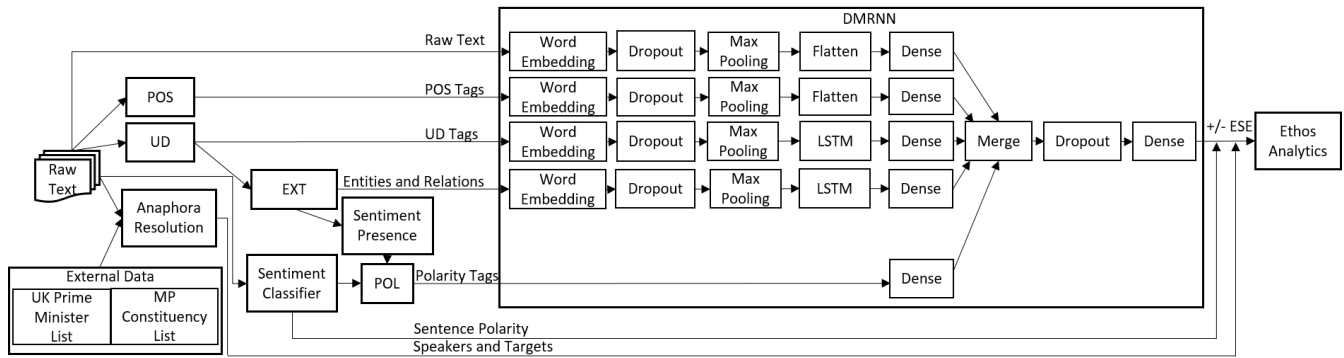


Figure 1: Pipeline for ethos mining featuring raw text; parsing (POS and UD); anaphora resolution with external data from Wikipedia; entity extraction (EXT); sentiment classification; sentiment presence; polarity combination (POL); and ESE/n-ESE classification performed by a DMRNN. The output of the pipeline is processed by ethos analytics.

of a negative connotation explicitly. For instance, Example 1 is tagged, speaker: Mr. Moore, target: Mr. Meyor (referred to as “hon. Friend”), and +ESE. All other sentences (a total of 3,007, 80% of the corpus) not labelled using these tags are considered as non-ethotic sentiment expressions (**n-ESEs**).

For the corpus evaluation, a second annotator analysed a 10% subset of Ethos_Hansard1 (see Table 2). This gave a Cohen’s $\kappa = 0.67$ on ESE and n-ESE (normalised for word count), $\kappa = 1$ for the polarity classification of ESEs (+/- ESEs) when both annotators have already agreed on an ESE, $\kappa = 1$ for speaker tags and $\kappa = 0.93$ for the target tags (errors are mainly due to multiple entities in an ESE).

4 Automatic Extraction

To extract ethos, we create an NLP pipeline (see Figure 1) with an input of raw natural language text and an output of +/-ESEs which are applied in the ethos analytics tool. The pipeline consists of modules which are either employing existing techniques; or modules which extend such methods for the purpose of ethos mining; or modules which contain original techniques developed specifically for this paper.

The raw text is passed to five areas of the pipeline (see arrows coming top-down from raw text in Figure 1): (1) directly to the DMRNN module; (2) the POS tagger; (3) the universal dependency (UD) tagger; (4) the sentiment classifier; and (5) the anaphora resolution module. Modules (3) and (4) are involved in complex processes. The UD tags are passed to the entity extraction (EXT) module (which removes entity references not relevant for ethos mining) and then to the sentiment presence module (which determines whether a sentence contains a sentiment). The output from the sentiment classifier is passed to the polarity (POL) module which combines the output from the sentiment presence module with the sentiment classifications. Next, the raw text, POS tags, UD tags, EXT output and POL tags are passed as separate inputs into the DMRNN, returning ESEs/n-ESEs. The output of sentiment classification determines +/-ESEs with the anaphora resolution module tagging each +/-ESE with a source and target. The resulting +/-ESEs are then applied in our ethos analytics tool (Section 6 for further details). In the following we describe each module of the pipeline in more detail.

Part-of-speech (POS) Tagging is applied to the raw corpus text, using the already existing Stanford parser [Toutanova *et al.*, 2003]. These are passed as an input to the DMRNN.

Universal Dependencies (UD) Tagging is also applied to the raw corpus text, again, using the already existing Stanford parser [Schuster and Manning, 2016]. The subjects of each sentence are tagged (e.g. “nsubj” in Figure 2) allowing the extraction of subject entities later in the EXT module. The UD tags are also passed as an input directly to the DMRNN.

Anaphora Resolution uses manual rules to determine the speaker and target entities of each sentence by extending existing techniques. The speaker is explicitly identified at the beginning of each of their turns. For the target of each sentence, we use manual rules relating to specific entity mentions. For entities where no unique identity details are given (e.g. “he”, “hon. Lady”), the system tracks back over a sentence and a speaker’s turn to determine if there are any entity mentions which do reveal unique values (e.g. “hon. Member for Bath”) or unique roles (e.g. “Prime Minister”), but excluding entities without human features (e.g. “the Government”). In these cases, we use **external data** containing MP details and a UK Prime Minister list scraped from Wikipedia using the month and year of each transcript in order to pinpoint the target. If there are no entity mentions present then the system sets the target as the speaker of the previous turn.




Figure 2: Dependency parse tree structure for Example 1.

Entity Extraction (EXT) is a novel manual rule-based system to remove entity mentions which are not relevant for ethos mining. The dialogue protocol in the Parliament dictates that in order to elicit a response from a Government representative, MPs have to frame their statements explicitly as questions which introduce entities who are addressees of the questions. Thus any initial framing pertaining to entities is ignored (e.g. “Is he aware...”), unless the entity is mentioned again within the sentence. The UD tags are used to extract all

entity mentions in a sentence and any words related to these entities, as long as the tag between them is a nominal subject, clause or modifier (Figure 2 and in Example 1 the following entities and related words are extracted “interests”, “pursuing”, “assiduously”, “constituents”, “Friend” and “his”).

Sentiment Classifier (SVM, NB, LR) extends existing methods considering three machine learning classifiers: a linear SVM classifier, a Naïve Bayes (NB) classifier and a Logistic Regression (LR) classifier, all within scikit-learn [Pedregosa *et al.*, 2011]. They are used to classify the polarity of sentences into positive and negative, performing a binary sentiment analysis (Example 1 gives positive sentiment). We also consider a TF-IDF approach to create vectors from a sentiment word lexicon (LIU) [Liu, 2010] and extend the approach with an ethotic word lexicon (ETH) [Duthie *et al.*, 2016] as well as negative words from a word embedding generated Hansard specific lexicon (HAN) [Rheault *et al.*, 2016]. The addition of ETH and HAN provides a relevant domain classification which is needed for words that play a different sentiment role in day-to-day speech. For example, “rich” has a positive meaning of wealthiness, but much of the time in the UK parliament it is used to describe someone as a hypocrite.

Sentiment Presence extends standard techniques, with an input of entities and related words from EXT. Using sentiment lexicons (LIU, ETH and HAN) sentences are tagged as holding sentiment or not (Example 1 contains sentiment due to the words “interests”, “pursuing” and “assiduously”).

Polarity (POL) is an original module which involves combining any entity relations tagged as containing sentiment with the corresponding sentiment classification. Sentences are grouped into positive, negative and non-sentiment bearing (Example 1 is marked as positive by combining the sentiment presence output and the sentiment classifier output).

Deep Modular Recurrent Neural Network (DMRNN) is developed to classify sentences as ESEs or n-ESEs. We use the keras deep learning library (<https://github.com/fchollet/keras>) to construct a novel DMRNN where the network hyper-parameters are specified using the validation set and early stopping (when the validation *F1*-score no longer increases) to determine the final model for testing. The DMRNN involves an ensemble of models as separate inputs (raw text, POS, UD, EXT and POL) which are then combined through concatenation in a final dense hidden layer with sigmoid activation. All of the models except polarity use the keras embedding layer with the output passed to a dropout layer to reduce over-fitting [Zaremba *et al.*, 2014], and a max pooling layer to reduce the feature set (see Table 3 for parameters). More specifically, the text and POS inputs are flattened to reduce the dimensionality of the vectors and then passed to a dense hidden layer containing 10 units and sigmoid activation. The UD and entity relation inputs are passed to an LSTM layer, ensuring that long distance relations are taken into account. Its output is then passed to a final dense hidden layer containing 10 units and sigmoid activation. The POL input is passed only to a dense hidden layer with 30 units and linear activation, due to the numerical values of the data where context is not needed. A final model is created by concatenating the individual models and it contains a

dropout layer to combat over-fitting; a final output layer with one neuron and sigmoid activation which uses Adam optimization [Kingma and Ba, 2014]; and a binary cross-entropy loss function for the binary classification to ESE and n-ESE. The sentiment classifier is then used to determine the polarity and create +/- ESEs.

Parameter	Value
Embedding dimension	input length: 400, output: 128
Hidden layer dimension	10
POL hidden layer	30
Dropout rate	0.20
LSTM layer	128
Max pooling size	4
Adam	α : 0.001, β_1 : 0.9, β_2 : 0.999

Table 3: Hyper-parameter values for the DMRNN.

5 Results

For the evaluation of the ethos mining pipeline, we provide results for anaphora resolution, the classification of +/-ESEs, the classification of ESE/n-ESE, and the combination of the highest performing classifiers for +/-ESEs and ESE/n-ESE.

The **results for anaphora resolution** are based upon the correct tagging of the target entity in each ethotic statement. The source is not considered as it will always have 100% accuracy due to the structure of Hansard transcripts. Target entities were tagged with an *F1*-score of 0.76 using the anaphora resolution function with errors mainly produced due to multiple different entities appearing in a sentence.

+/-ESE	P	R	m-F1
Baseline	0.63	0.65	0.64
[Duthie <i>et al.</i> , 2016]	0.78	0.78	0.78
SVM + LIU + ETH	0.81	0.81	0.81
NB + LIU + ETH	0.84	0.83	0.83
LR + LIU + ETH	0.80	0.80	0.80
SVM + LIU + ETH + HAN (Full)	0.78	0.79	0.78
NB + LIU + ETH + HAN (Full)	0.80	0.80	0.80
LR + LIU + ETH + HAN (Full)	0.79	0.80	0.79
SVM + LIU + ETH + HAN (Neg)	0.80	0.77	0.78
NB + LIU + ETH + HAN (Neg)	0.82	0.82	0.82
LR + LIU + ETH + HAN (Neg)	0.84	0.84	0.84*

Table 4: Classification of ESEs into positive and negative. Macro-averaged precision, recall and *F1*-score are reported for classification using machine learning classifiers and training lexicons, ethotic (ETH), Liu (LIU) and Hansard (HAN) compared against a baseline classifying on the training set distributions and against our previous work. (*) denotes classifier with highest *F1*-score.

Table 4 provides **+/-ESE results** using machine learning classifiers, TF-IDF vectorization and a combination of lexicons. The results are compared against a baseline classifier classifying on the training set class distributions and against the +/-ESE classification from our previous work on ethos mining [Duthie *et al.*, 2016]. All classifiers outperform the baseline (macro-averaged *F1*-score of 0.64) with the highest performing classifier, 31.3% above the baseline, consisting

ESE/n-ESE	P	R	F1	m-F1
Baseline	0.50	0.78	0.61	0.73
[Duthie <i>et al.</i> , 2016]	0.64	0.76	0.70	-
LR + Text	0.53	0.98	0.69	0.78
LR + Text + UD + POS + EXT	0.51	0.98	0.67	0.77
SVM + Text + UD + POS + EXT	0.52	0.98	0.68	0.77
POL	0.66	0.67	0.66	0.79
RNN + Text	0.52	0.98	0.68	0.77
CNN + Text	0.54	0.87	0.67	0.77
CNN + LSTM + Text	0.57	0.85	0.68	0.79
CNN + Text + UD LSTM	0.54	0.93	0.68	0.79
RNN + Text + UD LSTM	0.55	0.91	0.69	0.79
RNN + Text + UD LSTM + POL	0.58	0.85	0.69	0.80
RNN + Text + UD LSTM + POS + POL	0.58	0.88	0.70	0.80
RNN + Text + UD LSTM + EXT LSTM + POL	0.62	0.85	0.71	0.81
DMRNN	0.60	0.95	0.74*	0.83*

Table 5: Precision (P), recall (R) and $F1$ -score are reported for the classification of ESEs and the macro-averaged $F1$ -score (m- $F1$) is reported for the classification of ESEs and n-ESEs. Results from our previous work, standard machine learning classifiers, experimental classifications using different CNN and RNN modular combinations and our final DMRNN are compared to a baseline which classifies on the training set distributions. (*) denotes classifier with the highest $F1$ -scores.

of LR and the LIU, ETH and HAN (Neg) lexicons (m- $F1$ -score 0.84). The superior performance can be attributed to the domain specific lexicons (ETH and HAN), in particular the HAN (Neg) lexicon extends the set of negative domain specific words. The full HAN lexicon does not perform any better due to the use of word embeddings. Many non-sentiment holding words (e.g. “point”) are considered as positive, because of their context, skewing the classification.

Table 5 shows the **results of the ESE/n-ESE classification** with a combination of standard machine learning algorithms and RNN and CNN models compared against a baseline classifier using the training set distributions and against the ESE/n-ESE classification from [Duthie *et al.*, 2016]. All of the models outperform the baseline $F1$ -score of 0.61 with our DMRNN model giving the highest $F1$ -score 0.74 (21.3% above baseline) and macro-averaged $F1$ -score 0.83 (13.7% above baseline). The DMRNN model involves an ensemble method which does not suffer from the use of voting methods used in standard ensemble classifiers and explains the improvement of performance in the ethos mining pipeline. The addition of POL to any of the models provides an increase in $F1$ -score, highlighting the role sentiment plays in ethotic statements, but alone it is not enough to classify ESEs ($F1$ -score 0.66). CNN models have provided a boost in $F1$ -score in areas like sentiment analysis, but this was not the case for ESEs. Our relatively small dataset, in comparison to datasets in sentiment analysis, can explain this as a complex CNN model may be prone to over-fitting on smaller datasets (for further error analysis of the ESE/n-ESE and +/-ESE classifications see <http://arg.tech/Error>).

Finally, we provide **results for the combination of +/-ESE and ESE/n-ESE**. We take a correct ESE classification and do a +/-ESE classification. If this is correct then the classification is deemed as correct. The baseline trained on the class distributions gives the $F1$ -score 0.39, and our previous work, [Duthie *et al.*, 2016], gives $F1$ -score 0.60. Our highest performing classifier gives $F1$ -score 0.65, outperforming the baseline by 66.7% and our previous work by 8.3%.

6 Ethos Analytics

The availability of online transcripts from Hansard provides the opportunity to run the ethos mining pipeline on large unseen data for qualitative and quantitative analysis. Ethotic aspects of the political landscape are examined by comparing the ethos analytics data with political science analysis and news articles of the time. An automatically generated graph-based tool processes the ethos mining output, determining relationships between entities from the +/-ESEs and creating visualisations and graphs for individual entities (see Figure 3) of monthly attacks and supports of politicians and parties.

We extracted transcripts of the oral answers section of parliament from the period between 1946 (post war) and 1979 (before the Ethos_Hansard1 corpus) giving 314,465 sentences. When processed by the ethos mining pipeline 76,112 sentences were tagged as +/-ESEs. Using information about MPs and parties, we counted ethotic supports and attacks for pairs of political parties, pairs of politicians and individual politicians. The +/-ESE data split into months and years was normalized in two ways: (1) the number of sentences where an entity is the target of an attack on or support of ethos, divided by the sentences where the same entity is a target, and (2) the number of sentences where an entity supports or attacks the ethos of another entity, divided by the sentences where this entity was the source. We finally use the threshold of one standard deviation for each individual entity to investigate significant distributions from the norm.

Qualitative Ethos Analytics. In 1959, Margaret Thatcher became MP for the constituency Finchley representing the Conservative party. We investigate the period from her first appearance in the ethos analytics, 1962, to the year she became Prime Minister, 1979. This provides eight data points outside one standard deviation for attacks (1970:12, 1975:1, 1975:11, 1976:6, 1977:1, 1977:6, 1977:11, 1978:5 and 1978:8) and four data points for supports (1962:7, 1964:11, 1973:5 and 1978:5) of Thatcher’s ethos (see Figure 3). In 1962, when supports increased from 0 to 0.4,




Figure 3: Line chart showing the number of ethotic attacks and supports of Margaret Thatcher split by monthly intervals containing markers for one positive standard deviation (SD) from the mean of supports and attacks on ethos and political milestones.

she was recently promoted to a junior ministerial position. In 1964, she lost her position, due to the Conservative party losing the general election, but became shadow on pensions. This co-occurs with an increase in supports from 0 to 0.4. In 1970, attacks on Thatcher increased from 0.17 to 0.78 corresponding with her appointment as Education Secretary and the cuts she made to funding (<http://bbc.co.uk/timelines/zqp7tyc>). In 1973, supports of Thatcher’s ethos increased from 0 to 0.43. This follows the release of her paper in 1972 “Education: A Framework for Expansion” which reformed education. Thatcher became leader of the Conservative Party in 1975 which corresponds to an increase of attacks on her ethos before the election (0 to 1) and after (0.7 to 0.8). In 1976, attacks increased on Thatcher (0.67 to 0.79) corresponding to her leadership of the Conservative party and through changing dynamics in the parliament which saw the Government lose their majority. Throughout 1977, attacks increased on Thatcher’s ethos (0.69 to 0.9, 0.38 to 1, 0 to 0.8), which correlates with the local council elections in the UK which provided an insight into public attitude. The Conservative party won this resoundingly. Finally, in 1978 attacks increased on Thatcher’s ethos (0.29 to 0.91, 0.75 to 1) as did supports (0.14 to 0.43) which co-occurred with an impending general election where she was prominent and her party later won.

Quantitative Ethos Analytics. General elections took place nine times between 1946 and 1979. We look at the two main parties in the UK (Labour and Conservative parties) using two data points: supports in one quarter prior to an election and supports in one month prior to an election, to determine the slope, $s = (y_m - y_q) / (x_m - x_q)$ (where x_m and x_q are the time points of one month and one quarter earlier and y_m and y_q are the ethotic values), and ascertain the election winner. All supports of ethos were taken into account where the party was tagged except when the Government or the Opposition were targeted as they are heavily mentioned in oral answers to questions sessions. When comparing the slope values of each party in all but one case, the party with the most positive slope won the election. The exception occurred in 1966 when the Conservative party had a more positive slope (-0.003), however, the Labour party (-0.007) held a slim majority.

7 Conclusions and Future Work

This paper advances the state of the art in the new area of ethos mining. To automatically extract information about ethos, we created the largest publicly available corpus of ethotic relations, Ethos.Hansard1, and developed an ethos mining pipeline featuring several original components including a module exploiting a deep learning technique, DMRNN, which has been applied for the first time to the task of text classification. We experimented with various NLP methods and demonstrated that our pipeline outperforms all of them (macro- $F1 = 0.84$ for +/-ESE classification and macro- $F1 = 0.83$ for ESE/n-ESE classification), including the method proposed in our previous work on ethos mining, [Duthie *et al.*, 2016]. We also presented an application of ethos mining technology, ethos analytics, which is able to process large amounts of data, automatically annotated by our pipeline, bridging the gap between interpretation of major events and the raw data allowing for both qualitative analysis (e.g. from our analysis, attacks rose on Thatcher’s ethos whenever she was appointed to governmental position) and quantitative analysis (e.g. the political party with the most positive slope typically wins the general election). Such analyses allow the general public and political scientists to make sense of the constantly changing political landscape directly from the data rather than relying on experts, pollsters or the news.

In future work, we aim to further enrich the resources and improve the technology for ethos mining and analytics by annotating both ethos and logos, and thus the interactions between these; applying tree structures [Zhu *et al.*, 2015] as an input to the DMRNN; exploring the aggregation of the individual modules of the pipeline; developing modules for sarcasm and irony; and introducing ethotic metrics similar to argument metrics [Lawrence *et al.*, 2017] for the automatic identification of the most popular or most divisive speakers. Yet even in its current form, we have demonstrated that ethos mining coupled with ethos analytics can provide unique evidence in understanding the dynamics of governments.

Acknowledgments

The work was supported by EPSRC in the UK under grant EP/M506497/1.

References

- [Aristotle, 1991] Aristotle. *On Rhetoric* (G. A. Kennedy, Trans.). New York: Oxford University Press., 1991.
- [Budzynska and Villata, 2016] Katarzyna Budzynska and Serena Villata. Argument mining. In *The IEEE Intelligent Informatics Bulletin*, volume 17, pages 1–7, 2016.
- [Duthie et al., 2016] Rory Duthie, Katarzyna Budzynska, and Chris Reed. Mining ethos in political debate. In *Proc. of the Sixth Int Conference on Computational Models of Argument (COMMA)*, pages 299–310. IOS Press, 2016.
- [Eger et al., 2017] Steffen Eger, Johannes Daxenberger, and Iryna Gurevych. Neural end-to-end learning for computational argumentation mining. *arXiv preprint arXiv:1704.06104*, 2017.
- [Hirst et al., 2014] Graeme Hirst, Yaroslav Riabinin, Jory Graham, Magali Boizot-Roche, and Colin Morris. From text to political positions: Text analysis across disciplines. In *Text to ideology or text to party status?*, pages 61–79. John Benjamins Publishing Company, Amsterdam, 2014.
- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [Ju et al., 2017] Cheng Ju, Aurélien Bibaut, and Mark J van der Laan. The relative performance of ensemble methods with deep convolutional neural networks for image classification. *arXiv preprint arXiv:1704.01664*, 2017.
- [Kingma and Ba, 2014] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [Lai et al., 2015] Siwei Lai, Liheng Xu, Kang Liu, and Jun Zhao. Recurrent convolutional neural networks for text classification. In *In Proc of the Twenty-Ninth AAAI*, volume 333, pages 2267–2273, 2015.
- [Larkey and Croft, 1996] Leah S. Larkey and W. Bruce Croft. Combining classifiers in text categorization. In *Proc of the 19th Annual Int Conference on Research and Development in Information Retrieval, SIGIR*, pages 289–297, New York, NY, USA, 1996. ACM.
- [Lawrence et al., 2017] John Lawrence, Joonsuk Park, Katarzyna Budzynska, Claire Cardie, Barbara Konat, and Chris Reed. Using argumentative structure to interpret debates in online deliberative democracy and erulemaking. *ACM Trans. Internet Technol.*, 17(3):25:1–25:22, 2017.
- [LeCun et al., 2015] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [Lee et al., 2017] Jinhyuk Lee, Hyunjae Kim, Miyoung Ko, Donghee Choi, Jaehoon Choi, and Jaewoo Kang. Name nationality classification with recurrent neural networks. In *Proc of the Twenty-Sixth Int Joint Conference on Artificial Intelligence, IJCAI-17*, pages 2081–2087, 2017.
- [Liu, 2010] Bing Liu. Sentiment analysis and subjectivity. *Handbook of Natural Language Processing*, 2:627–666, 2010.
- [Ma et al., 2016] Lin Ma, Zhengdong Lu, and Hang Li. Learning to answer questions from image using convolutional neural network. In *Proc of the Thirtieth AAAI*, pages 3567–3573. AAAI Press, 2016.
- [Naderi and Hirst, 2017] Nona Naderi and Graeme Hirst. Recognizing reputation defence strategies in critical political exchanges. In *Proc of the Int Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 527–535, 2017.
- [Pedregosa et al., 2011] 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.
- [Rheault et al., 2016] Ludovic Rheault, Kaspar Beelen, Christopher Cochrane, and Graeme Hirst. Measuring emotion in parliamentary debates with automated textual analysis. *PLOS ONE*, 11:1–18, 12 2016.
- [Schuster and Manning, 2016] Sebastian Schuster and Christopher D. Manning. Enhanced english universal dependencies: An improved representation for natural language understanding tasks. In *LREC 2016*, pages 2371–2378, 2016.
- [Sutskever et al., 2014] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems*, pages 3104–3112, 2014.
- [Thomas et al., 2006] Matt Thomas, Bo Pang, and Lillian Lee. Get out the vote: Determining support or opposition from Congressional floor-debate transcripts. In *Proc of EMNLP*, pages 327–335, 2006.
- [Toutanova et al., 2003] Kristina Toutanova, Dan Klein, Christopher Manning, and Yoram Singer. Feature-rich part-of-speech tagging with a cyclic dependency network. *Proc of HLT-NAACL*, pages 252–259, 2003.
- [Wang et al., 2017] Jin Wang, Zhongyuan Wang, Dawei Zhang, and Jun Yan. Combining knowledge with deep convolutional neural networks for short text classification. In *Proc of the Twenty-Sixth Int Joint Conference on Artificial Intelligence, IJCAI-17*, pages 2915–2921, 2017.
- [Zaremba et al., 2014] Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329*, 2014.
- [Zhou et al., 2015] Chunting Zhou, Chonglin Sun, Zhiyuan Liu, and Francis Lau. A C-LSTM neural network for text classification. *arXiv preprint arXiv:1511.08630*, 2015.
- [Zhu et al., 2015] Xiaodan Zhu, Parinaz Sobhani, and Hongyu Guo. Long short-term memory over recursive structures. In *Proc of the 32nd Int Conference on Machine Learning, ICML 15*, pages 1604–1612, 2015.