

# Target Based Speech Act Classification in Political Campaign Text

Shivashankar Subramanian    Trevor Cohn    Timothy Baldwin

School of Computing and Information Systems  
The University of Melbourne

shivashankar@student.unimelb.edu.au    {t.cohn, tbaldwin}@unimelb.edu.au

## Abstract

We study pragmatics in political campaign text, through analysis of speech acts and the target of each utterance. We propose a new annotation schema incorporating domain-specific speech acts, such as *commissive-action*, and present a novel annotated corpus of media releases and speech transcripts from the 2016 Australian election cycle. We show how speech acts and target referents can be modeled as sequential classification, and evaluate several techniques, exploiting contextualized word representations, semi-supervised learning, task dependencies and speaker meta-data.

## 1 Introduction

Election campaign text is a core artifact in political analysis. Campaign communication can influence a party’s reputation, credibility, and competence, which are primary factors in voter decision making (Fernandez-Vazquez, 2014). Also, modeling the discourse is key to measuring the role of party in *constructive democracy* — to engage in constructive discussion with other parties in a democracy (Gibbons et al., 2017).

Speech act theory (Austin, 1962; Searle, 1976) can be used to study such pragmatics in political campaign text. Traditional speech act classes have been studied to analyze how people engage with elected members (Hemphill and Roback, 2014), and how elected members engage in discussions (Shapiro et al., 2018), with a particular focus on pledges (Artés, 2011; Naurin, 2011, 2014; Gibbons et al., 2017). Also, election manifestos have been analyzed for prospective and retrospective messages (Müller, 2018). In this work, we combine traditional speech acts with those proposed by political scientists to study political discourse, such as *specific pledges*, which can also help to verify the pledges’ fulfilment after an election (Thomson et al., 2010).

In addition to speech acts, it is important to identify the target of each utterance — that is, the political party referred to in the text — in order to determine the discourse structure. Here, we study the effect of jointly modeling the speech act and target referent of each utterance, in order to exploit the task dependencies. That is, this paper is an application of discourse analysis to the pragmatics-rich domain of political science, to determine the intent of every utterance made by politicians, and in part, automatically extract pledges at varying levels of specificity from campaign speeches and press releases.

We assume that each utterance is associated with a unique speech act (similar to Zhao and Kawahara (2017)) and target party,<sup>1</sup> meaning that a sentence with multiple speech acts and/or targets must be segmented into component utterances. Take the following example, from the Labor Party:

- (1) Labor will contribute \$43 million towards the Roe Highway project and we call on the WA Government to contribute funds to get the project underway.

The example is made up of two utterances (with and without an underline), belonging to speech act types *commissive-action-specific* and *directive*, referring to different parties (LABOR and LIBERAL), resp. In our initial experiments, we perform target based speech act classification (i.e. joint speech act classification and determination of the target of the utterance) over gold-standard utterance data (Section 6), but return to perform automatic utterance segmentation along with target based speech act classification (Section 7).

While speech act classification has been applied to a wide range of domains, its application to polit-

<sup>1</sup> Zhao and Kawahara (2017) do not address the target referent classification task in their work.

Utterance	Speech act	Target party	Speaker
Tourism directly and indirectly supports around 38000 jobs in TAS.	<i>assertive</i>	NONE	LABOR
We will invest \$25.4 million to increase forensics and intelligence assets for the Australian Federal Police	<i>commissive-action-specific</i>	LIBERAL	LIBERAL
Labor will prioritise the Metro West project if elected to government.	<i>commissive-action-vague</i>	LABOR	LABOR
A Shorten Labor Government will create 2000 jobs in Adelaide.	<i>commissive-outcome</i>	LABOR	LABOR
Federal Labor today calls on the State Government to commit the final \$75 million to make this project happen.	<i>directive</i>	LIBERAL	LABOR
Good morning everybody.	<i>expressive</i>	NONE	LABOR
The Coalition has already delivered a \$2.5 billion boost to our law enforcement and security agencies.	<i>past-action</i>	LIBERAL	LIBERAL
Malcolm Turnbull’s health cuts will rip up to \$1.4 billion out of Australians’ pockets every year	<i>verdictive</i>	LIBERAL	LABOR

Table 1: Examples with speech act and target party classes. “Speaker” denotes the party making the utterance.

ical text is relatively new. Most speech act analyses in the political domain have relied exclusively on manual annotation, and no labeled data has been made available for training classifiers. As it is expensive to obtain large-scale annotations, in addition to developing a novel annotated dataset, we also experiment with a semi-supervised approach by utilizing unlabeled text, which is easy to obtain.

The contributions of this paper are as follows: (1) we introduce the novel task of target based speech act classification to the analysis of political discourse; (2) we develop and release a dataset (can be found here <https://github.com/shivashankarrs/Speech-Acts>) based on political speeches and press releases, from the two major parties — Labor and Liberal — in the 2016 Australian federal election cycle; and (3) we propose a semi-supervised learning approach to the problem by augmenting the training data with in-domain unlabeled text.

## 2 Related Work

The recent adoption of NLP methods has led to significant advances in the field of computational social science (Lazer et al., 2009), including political science (Grimmer and Stewart, 2013). With the increasing availability of datasets and computational resources, large-scale comparative political text analysis has gained the attention of political scientists (Lucas et al., 2015). One task of particular importance is the analysis of the functional

intent of utterances in political text. Though it has received notable attention from many political scientists (see Section 1), the primary focus of almost all work has been to derive insights from manual annotations, and not to study computational approaches to automate the task.

Another related task in the political communication domain is reputation defense, in terms of party credibility. Recently, Duthie and Budzynska (2018) proposed an approach to mine ethos support/attack statements from UK parliamentary debates, while Naderi and Hirst (2018) focused on classifying sentences from Question Time in the Canadian parliament as defensive or not. In this work, our source data is speeches and press releases in the lead-up to a federal election, where we expect there to be rich discourse and interplay between political parties.

Speech act theory is fundamental to study such discourse and pragmatics (Austin, 1962; Searle, 1976). A speech act is an illocutionary act of conversation and reflects shallow discourse structures of language. Due to its predominantly small-data setting, speech act classification approaches have generally relied on bag-of-words models (Qadir and Riloff, 2011; Vosoughi and Roy, 2016), although recent approaches have used deep-learning models through data augmentation (Joty and Hoque, 2016) and learning word representations for the target domain (Joty and Mohiuddin, 2018), outperforming traditional bag-of-words approaches.

Another technique that has been applied to com-

compensate for the sparsity of labeled data is semi-supervised learning, making use of auxiliary unlabeled data, as done previously for speech act classification in e-mail and forum text (Jeong et al., 2009). Zhang et al. (2012) also used semi-supervised methods for speech act classification over Twitter data. They used transductive SVM and graph-based label propagation approaches to annotate unlabeled data using a small seed training set. Joty and Mohiuddin (2018) leveraged out-of-domain labeled data based on a domain adversarial learning approach. In this work, we focus on target based speech act analysis (with a custom class-set) for political campaign text and use a deep-learning approach by incorporating contextualized word representations (Peters et al., 2018) and a cross-view training framework (Clark et al., 2018) to leverage in-domain unlabeled text.

### 3 Problem Statement

Target based speech act classification requires the segmentation of sentences into utterances, and labelling of those utterances according to speech act and target party. In this work we focus primarily on speech act and target party classification.

Our speech act coding schema is comprised of: *assertive*, *commissive*, *directive*, *expressive*, *past-action*, and *verdictive*. An *assertive* commits the speaker to something being the case. With a *commissive*, the speaker commits to a future course of action. Following the work of Artés (2011) and Naurin (2011), we distinguish between *action* and *outcome* commissives. Action commissives (*commissive-action*) are those in which an action is to be taken, while outcome commissives (*commissive-outcome*) can be defined as a description of reality or goals. Secondly, similar to Naurin (2014) we also classify action commissives into vague (*commissive-action-vague*) and specific (*commissive-action-specific*), according to their specificity. This distinction is also related to text specificity analysis work addressed in the news (Louis and Nenkova, 2011) and classroom discussion (Lugini and Litman, 2017) domains. A *directive* occurs when the speaker expects the listener to take action in response. In an *expressive*, the speaker expresses their psychological state, while a *past-action* denotes a retrospective action of the target party, and a *verdictive* refers to an assessment on prospective or retrospective actions.

Examples of the eight speech act classes are

# Doc	# Sent	# Utt	Avg Utterance Length
258	6609	7641	19.3

Table 2: Dataset Statistics: number of documents, number of sentences, number of utterances, and average utterance length

given in Table 1, along with the target party (LABOR, LIBERAL, or NONE), indicating which party the speech act is directed towards, and the “speaker” party making the utterance (information which is provided for every utterance).

#### 3.1 Utterance Segmentation

Sentences are segmented both in the context of speech act and target party — when a sentence has utterances belonging to more than one speech act or/and more than one target. For example, the following sentence conveys a pledge (*commissive-outcome*) followed by the party’s belief (*assertive*), with the utterance boundary indicated by |:

- (2) We will save Medicare | because Medicare is more than just a standard of health.

Further, the following (from the Labor party) has segments comparing LABOR and LIBERAL:

- (3) Our party is united – | the Liberals are not united.

### 4 Election Campaign Dataset

We collected media releases and speeches from the two major Australia political parties — Labor and Liberal — from the 2016 Australian federal election campaign. A statistical breakdown of the dataset is given in Table 2. We compute agreement over 15 documents, annotated by two independent annotators, with disagreements resolved by a third annotator. The remaining documents are annotated by the two main annotators without redundancy. Agreement between the two annotators for utterance segmentation based on exact boundary match using Krippendorff’s alpha ( $\alpha$ ) (Krippendorff, 2011) is 0.84. Agreement statistics for the classification tasks (Cohen, 1960; Carletta, 1996) are given in Tables 3 and 4.

Speech act	%	Kappa ( $\kappa$ )
<i>assertive</i>	40.8	0.85
<i>commissive-action-specific</i>	12.4	0.84
<i>commissive-action-vague</i>	6.6	0.73
<i>commissive-outcome</i>	4.9	0.72
<i>directive</i>	1.7	0.92
<i>expressive</i>	1.9	0.88
<i>past-action</i>	6.3	0.76
<i>verdictive</i>	25.4	0.82

Table 3: Speech act agreement statistics

Target party	%	Kappa ( $\kappa$ )
LABOR	45.9	0.92
LIBERAL	39.1	0.90
NONE	15.0	0.86

Table 4: Target party agreement statistics

## 5 Proposed Approach

Our approach to labeling utterances for speech act and target party classification is as follows. Utterances are first represented as a sequence of word embeddings, and then using a bidirectional Gated Recurrent Unit (“biGRU”: Cho et al. (2014)). The representation of each utterance is set to the concatenation of the last hidden state of both the forward and backward GRUs,  $\mathbf{h}_i = [\overrightarrow{\mathbf{h}}_i, \overleftarrow{\mathbf{h}}_i]$ . After this, the model has a softmax output layer. This network is trained for both the speech act (eight class) and target party (three class) classification tasks, minimizing cross-entropy loss, denoted as  $\mathcal{L}_S$  and  $\mathcal{L}_T$  respectively.

Our approach has the following components:

**ELMo word embeddings** (“biGRU<sub>ELMo</sub>”): As word embeddings we use a 1024d learned linear combination of the internal states of a bidirectional language model (Peters et al., 2018).

**Semi-supervised Learning:** We employ a cross-view training approach (Clark et al., 2018) to leverage a larger volume of unlabeled text. Cross-view training is a kind of teacher–student method, whereby the model “teaches” another “student” model to classify unlabelled data. The student sees only a limited form of the data, e.g., through application of noise (Sajjadi et al., 2016; Wei et al., 2018), or a different view of the input, as used herein. This procedure regularises the learning of the teacher to be more robust, as well as increasing the exposure to unlabeled text.

We augment our dataset with over 36k sentences from Australian Prime Minister candidates’

election speeches.<sup>2</sup> On these unlabeled examples, the model’s probability distribution over targets  $p_\theta(y|s)$  is used to fit auxiliary model(s),  $p_\omega(y|s)$ , by minimising the Kullback-Leibler (KL) divergence,  $\text{KL}(p_\theta(y|s), p_\omega(y|s))$ . This consensus loss component, denoted  $\mathcal{L}_{\text{unsup}}$ , is added to the supervised training objective ( $\mathcal{L}_S$  or  $\mathcal{L}_T$ ).

We evaluate the following auxiliary models:<sup>3</sup>

- a forward GRU (“biGRU<sub>CVT<sub>fwd</sub></sub>”);
- separate forward and backward GRUs (“biGRU<sub>CVT<sub>fwdbwd</sub></sub>”); and
- a biGRU with word-level dropout (“biGRU<sub>CVT<sub>worddrop</sub></sub>”).

The intuition is that the student models only have access to restricted views of the data on which the teacher network is trained, and therefore this acts as a regularization factor over the unlabeled data when learning the teacher model.

**Multi-task Learning** (“biGRU<sub>Multi</sub>”): For speech act classification, target party classification is considered as an auxiliary task, and vice versa. Accordingly, a separate model is built for each task, with the other task as an auxiliary task, in each case using a linearly weighted objective  $\mathcal{L}_S + \alpha\mathcal{L}_T$ , where  $\alpha \geq 0$  is tuned separately in each application. The intuition here is to capture the dependencies between the tasks, e.g., *commissive* is relevant to the Speaker party only.

**Meta-data** (biGRU<sub>Meta</sub>): We concatenate a binary flag encoding the speaker party ( $\mathbf{m}_i$ ) alongside the utterance embedding  $\mathbf{h}_i$ , i.e.,  $[\mathbf{h}_i, \mathbf{m}_i]$ . This representation is passed through a hidden layer with ReLU-activation, then projected onto a output layer with softmax activation for both the classification tasks.

## 6 Evaluation

We compare the models presented in Section 5 with the following baseline approaches:

- Support Vector Machine (“SVM<sub>BoW</sub>”) with unigram term-frequency representation.
- Multi-layer perceptron (“MLP<sub>BoW</sub>”) with unigram term-frequency representation.

<sup>2</sup><https://primeministers.moadoph.gov.au/collections/election-speeches>

<sup>3</sup>Note that auxiliary models share parameters with the corresponding components of main (teacher) model, with the exception of their output layers.

ID	Approach	Speech act		Target party	
		Accuracy	Macro-F1	Accuracy	Macro-F1
1	Meta <sub>naive</sub>	—	—	0.55	0.43
2	SVM <sub>BoW</sub>	0.56	0.41	0.60 <sup>*1</sup>	0.56 <sup>*1</sup>
3	MLP <sub>BoW</sub>	0.60 <sup>*2</sup>	0.47 <sup>*2</sup>	0.61 <sup>*1</sup>	0.57 <sup>*1</sup>
4	DAN <sub>GloVe</sub>	0.53	0.30	0.59	0.54
5	GRU <sub>GloVe</sub>	0.56	0.46	0.58	0.55
6	biGRU <sub>GloVe</sub>	0.57	0.48	0.59	0.56
7	MLP <sub>ELMo</sub>	0.62 <sup>*3</sup>	0.53 <sup>*3</sup>	0.58	0.57
8	biGRU <sub>ELMo</sub>	<b>0.68<sup>*7</sup></b>	<b>0.57<sup>*7</sup></b>	0.63 <sup>*2,3,7</sup>	0.60 <sup>*2,3,7</sup>
9	biGRU <sub>ELMo</sub> + CVT <sub>fwd</sub>	0.66	0.55	0.63	0.58
10	biGRU <sub>ELMo</sub> + CVT <sub>fwd+bwd</sub>	<b>0.68</b>	0.54	0.61	0.56
11	biGRU <sub>ELMo</sub> + CVT <sub>worddrop</sub>	<b>0.69</b>	<b>0.57</b>	0.66 <sup>*8</sup>	0.60
12	biGRU <sub>ELMo</sub> + CVT <sub>worddrop</sub> + Multi	<b>0.69</b>	<b>0.58</b>	0.65	0.60
13	biGRU <sub>ELMo</sub> + CVT <sub>worddrop</sub> + Meta	<b>0.68</b>	<b>0.58</b>	<b>0.71<sup>*11</sup></b>	<b>0.66<sup>*8,11</sup></b>

Table 5: Classification results showing average performance based on 10 runs. \* indicates results significantly better than the indicated approaches (based on ID in the table) according to a paired t-test ( $p < 0.05$ ). Boldface shows the overall best results and results insignificantly different from the best. Meta<sub>naive</sub> is not applicable for speech act classification. Note that all approaches use gold-standard segmentation for evaluation.

Speech act	MLP <sub>ELMo</sub>	Our approach
<i>assertive</i>	0.77	0.80
<i>commissive-action-specific</i>	0.65	0.69
<i>commissive-action-vague</i>	0.45	0.48
<i>commissive-outcome</i>	0.28	0.39
<i>directive</i>	0.58	0.59
<i>expressive</i>	0.55	0.58
<i>past-action</i>	0.45	0.48
<i>verdictive</i>	0.48	0.61

Table 6: Speech act class-wise F1 score.

Target party	biGRU <sub>ELMo</sub>	Our approach
LABOR	0.68	0.74
LIBERAL	0.65	0.75
NONE	0.46	0.48

Table 7: Target party class-wise F1 score.

- Deep Averaging Networks (“DAN<sub>GloVe</sub>”) (Iyyer et al., 2015), GRU (“GRU<sub>GloVe</sub>”), and biGRU (“biGRU<sub>GloVe</sub>”) with pre-trained 300d GloVe embeddings (Pennington et al., 2014).
- MLP with average-pooling over pre-trained ELMo word embeddings (“MLP<sub>ELMo</sub>”).
- Using speaker party as the predicted target party (“Meta<sub>naive</sub>”).

We average results across 10 runs with 90%/10% training/test random splits. Hyperparameters are tuned over a 10% validation

set randomly sampled and held out from the training set. We evaluate using accuracy and macro-averaged F-score, to account for class-imbalance. We compare the baseline approaches against our proposed approach (different components given in Section 5). We evaluate the effect of each component by adding them to the base model (biGRU<sub>ELMo</sub>), e.g., biGRU model with ELMo embeddings and word-level dropout based semi-supervised approach is given as biGRU<sub>ELMo</sub> + CVT<sub>worddrop</sub>. Results for speech act and target party classification are given in Table 5. The corresponding class-wise performance for both speech act and target party tasks with our approach (biGRU<sub>ELMo</sub> + CVT<sub>worddrop</sub> + Meta) compared against the competitive approach from Table 5 is given in Table 6 and Table 7 respectively (and also discussed further in Section 8). All the approaches are evaluated with the gold-standard segmentation. Utterance segmentation is discussed in Section 7.

From the results in Table 5, we observe that the biGRU<sup>4</sup> performs better than the other approaches, and that ELMo contextual embeddings (biGRU<sub>ELMo</sub>) boosts the performance appreciably. Apart from ELMo, the semi-supervised learning methods (biGRU<sub>ELMo</sub> + CVT<sub>worddrop</sub>) provide a boost in performance for the target party

<sup>4</sup>The biGRU model uses ReLU activations, a 128d hidden layer for speech act classification and 64d hidden layer for target party classification, and dropout rate of 0.1.

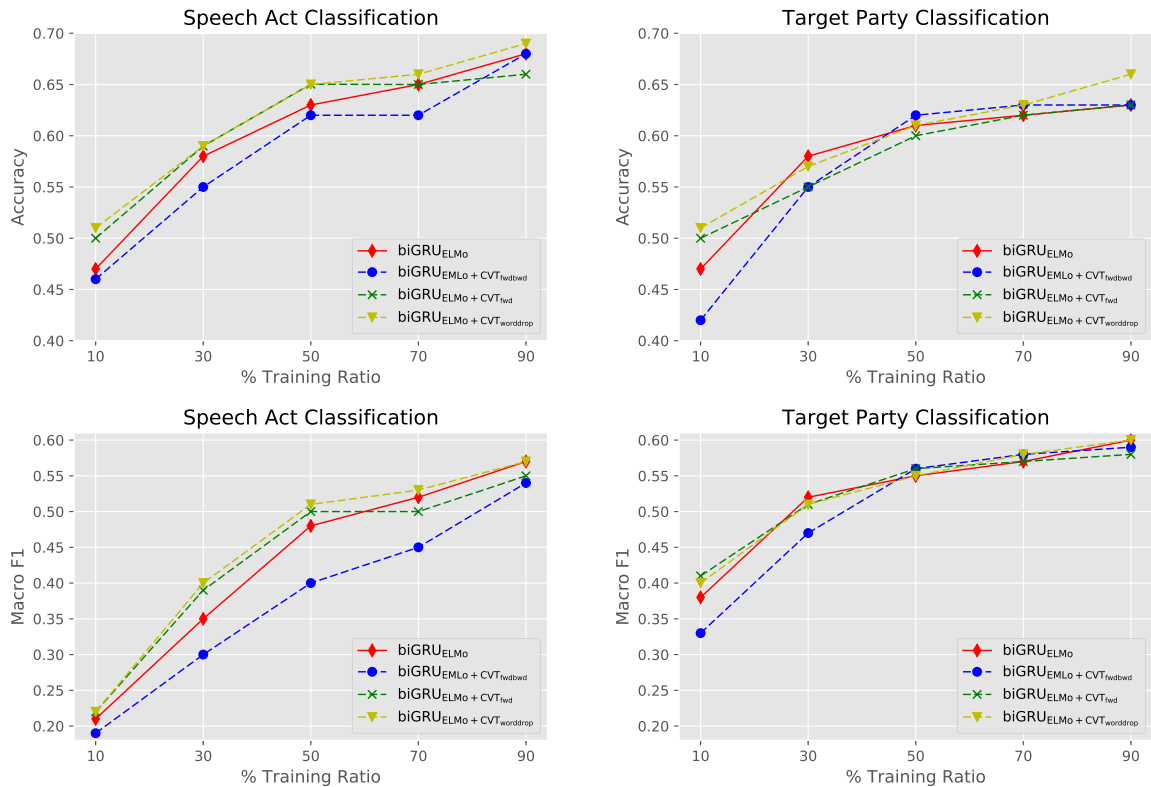


Figure 1: Classification performance across different training ratios. Note that 90% is using all the training data, as 10% is used for validation.

task (wrt accuracy) using all the training data.  $\text{biGRU}_{\text{ELMo}} + \text{CVT}_{\text{worddrop}}$  and  $\text{biGRU}_{\text{ELMo}} + \text{CVT}_{\text{fwd}}$  provide gains in performance for the speech act task, especially with fewer training examples ( $\leq 50\%$  of training data, see Figure 1). Performance of semi-supervised learning models with cross-view training (which leverages in-domain unlabeled text) is compared against  $\text{biGRU}_{\text{ELMo}}$ , which is a supervised approach. Results across different training ratio settings are given in Figure 1. From this, we can see that  $\text{biGRU}_{\text{ELMo}} + \text{CVT}_{\text{worddrop}}$  and  $\text{biGRU}_{\text{ELMo}} + \text{CVT}_{\text{fwd}}$  performs better than  $\text{biGRU}_{\text{ELMo}} + \text{CVT}_{\text{fwbwd}}$  in almost all cases. With a training ratio  $\leq 50\%$ ,  $\text{biGRU}_{\text{ELMo}} + \text{CVT}_{\text{worddrop}}$  achieves a comparable performance to  $\text{biGRU}_{\text{ELMo}} + \text{CVT}_{\text{fwd}}$ .

Multi-task learning ( $\text{biGRU}_{\text{ELMo}} + \text{CVT}_{\text{worddrop}} + \text{Multi}$ ) provides only small improvements for the speech act task. Further, when we add speaker party meta-data ( $\text{biGRU}_{\text{ELMo}} + \text{CVT}_{\text{worddrop}} + \text{Meta}$ ), it provides large gains in performance for the target party task. Overall, the proposed approach ( $\text{biGRU}_{\text{ELMo}} + \text{CVT}_{\text{worddrop}} + \text{Meta}$ ) provides the best performance for the target party task. Its performance is better than the  $\text{biGRU}_{\text{ELMo}} + \text{Meta}$

model, which does not leverage the additional unlabeled text using semi-supervised learning, where it achieves 0.70 accuracy and 0.65 Macro F1. Also, ELMo and semi-supervised methods ( $\text{biGRU}_{\text{ELMo}} + \text{CVT}_{\text{worddrop}}$  and  $\text{biGRU}_{\text{ELMo}} + \text{CVT}_{\text{fwd}}$ ) provide significant improvements for the speech act task, especially under sparse supervision scenarios (see Figure 1, for training ratio  $\leq 50\%$ ).

## 7 Segmentation Results

In the previous experiments, we used gold-standard utterance data, but next we experiment with automatic segmentation. We use sentences as input, based on the NLTK sentence tokenizer (Bird et al., 2009), and automatically segment sentences into utterances based on token-level segmentation, in the form of a BI binary sequence classification task using a CRF model (Hernault et al., 2010).<sup>5</sup> We use the following set of features for each word: token, word shape (capitalization, punctuation, digits), Penn POS tags based on SpaCy, ClearNLP dependency labels (Choi and Palmer, 2012), relative position in the sentence, and features for the

<sup>5</sup>We also experimented with a neural CRF model, but found it to be less accurate.

Utterance	Target party	Speaker
Our new Tourism Infrastructure Fund will bring more visitor dollars and more hospitality jobs to Cairns, Townsville and the regions.	LABOR	LABOR
Just as he sold out 35,000 owner-drivers in his deal with the TWU to bring back the “Road Safety Remuneration Tribunal”.	LABOR	LIBERAL
Then in 2022, we will start construction of the first of 12 regionally superior submarines, the single biggest investment in our military history.	LIBERAL	LIBERAL

Table 8: Scenarios where “Speaker” meta-data benefits the target party classification task.

adjacent words (based on this same feature representation). We compute segmentation accuracy (SA: Zimmermann et al. (2006)), which measures the percentage of segments that are correctly segmented, i.e. both the left and right boundary match the reference boundaries. SA for the CRF model is 0.87. Secondly, to evaluate the effect of segmentation on classification, we compute joint accuracy (JA). It is similar to SA but also requires correctness of the speech act and target party. In cascaded style, JA using the CRF model for segmentation and  $\text{biGRU}_{\text{ELMo} + \text{CVT}_{\text{worddrop}} + \text{Meta}}$  for speech act and target party classification is 0.60 and 0.64 respectively. Here, segmentation errors lead to a small drop in performance.

## 8 Error Analysis

We analyze the class-wise performance and confusion matrix for our best performing approach ( $\text{biGRU}_{\text{ELMo} + \text{CVT}_{\text{worddrop}} + \text{Meta}}$ ). Speech act and target party class-wise performance is given in Tables 6 and 7 respectively. We can see that the proposed approach provides improvement across all classes, while achieving comparable performance for *directive*. Recognizing *commissive-outcome* can be seen to be tougher than other classes. In addition, we analyze the results to identify cases where having “Speaker” party information is beneficial for predicting the target party of sentences. Some of those scenarios are given in Table 8, where the meta-data enables predicting the target party correctly even when there is no explicit reference to the party or leaders.

Confusion matrices for the speech act and target party classification tasks are given in Figure 2. Some observations from the confusion matrices are: (a) *assertive* and *verdictive* are often misclassified as each other; (b) *commissive-action-vague* utterances are often misclassified as *commissive-action-specific*; and (c) LABOR and LIBERAL classes are often misclassified as each

other for the target party classification task.

## 9 Qualitative Analysis

Here we provide the policy-wise speech act distribution for both parties, which indicates the difference in their predilection for the indicated six policy areas (Figure 3). We provide results for the six most frequent policy categories, for each of which, the campaign text is first classified into one of the policy-areas that are relevant to Australian politics, by building a Logistic Regression classifier with data obtained from ABC Fact Check.<sup>6</sup> Some observations (based on Figure 3) are as follows:

- The incumbent government (LIBERAL) uses more *directive*, *expressive*, *verdictive*, and *past-action* utterances than the opposition (LABOR).
- LIBERAL’s text has relatively more pledges (*commissive-action-vague*, *commissive-action-specific* and *commissive-outcome*) on **economy** compared to LABOR, whereas LABOR has more pledges on **social services** and **education**. This is as expected for right- and left-wing parties respectively. Other policy-areas have a comparable number of pledges from both parties. Overall, party-wise salience towards these policy areas correlates highly with the relative breakdowns in the Comparative Manifesto Project (Volkens et al., 2017): where the relative share of sentences from the LABOR and LIBERAL manifestos<sup>7</sup> for **welfare state (health and social services)** is 22:7, **education** is 9:6, **economy** is 11:23, and **technology & infrastructure (communication, infrastructure)** is 17:19.
- Across policy-areas, *specific* pledges are

<sup>6</sup><https://www.abc.net.au/news/factcheck>

<sup>7</sup>[https://manifesto-project.wzb.eu/download/data/2018b/datasets/MPDataset\\_MPDS2018b.csv](https://manifesto-project.wzb.eu/download/data/2018b/datasets/MPDataset_MPDS2018b.csv)




Figure 2: Confusion matrix for speech act and target party classification tasks.

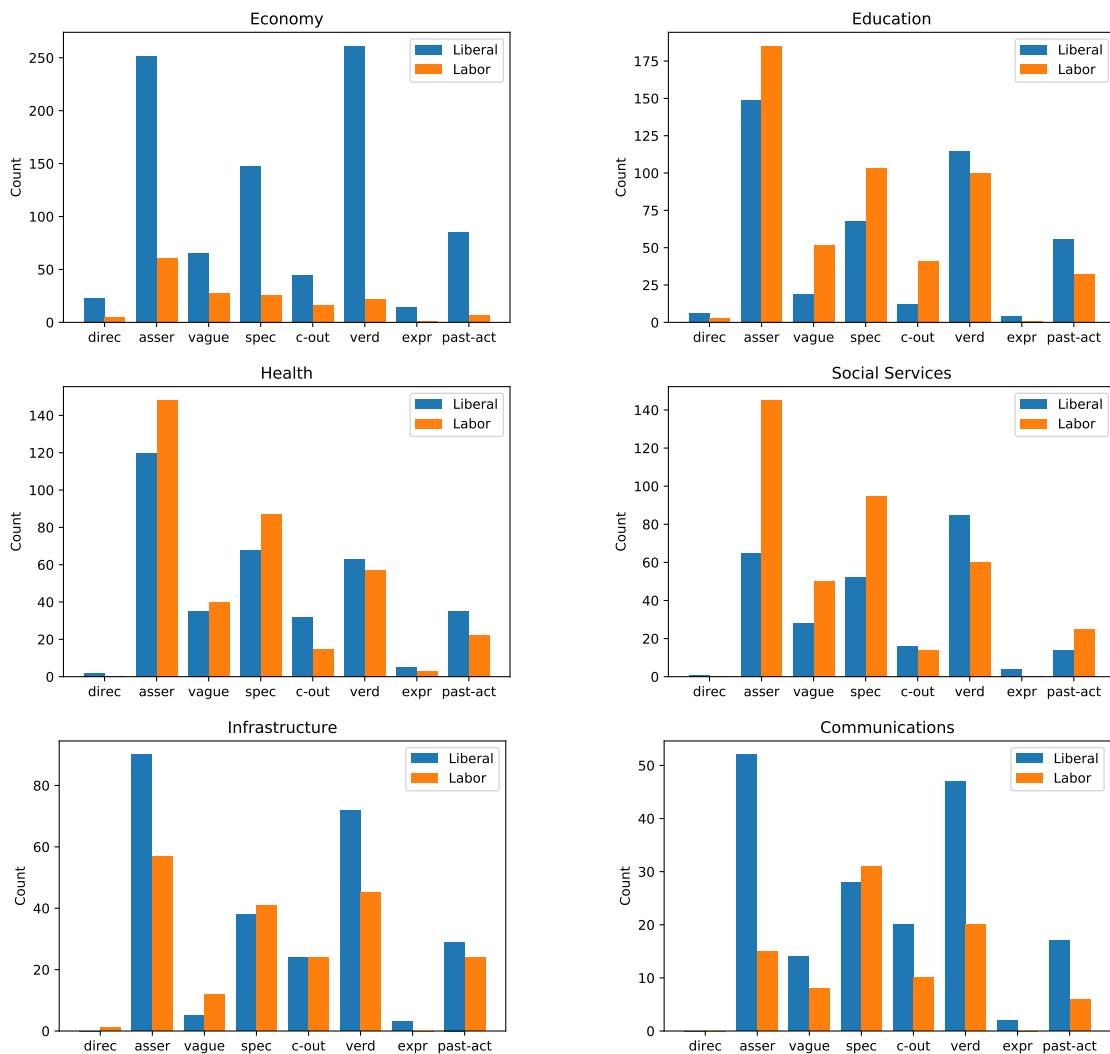


Figure 3: Policy-wise speech act analysis. Classes include: *directive* (“direc”), *assertive* (“asser”), *commissive-action-vague* (“vague”), *commissive-action-specific* (“spec”), *commissive-outcome* (“c-out”), *verdictive* (“verd”), *expressive* (“expr”), and *past-action* (“past-act”)



more frequent than *vague* ones. This aligns with previous studies done by Naurin (2014) and Gibbons et al. (2017).

## 10 Conclusion and Future Work

In this work we present a new dataset of election campaign texts, based on a class schema of speech acts specific to the political science domain. We study the associated problems of identifying the referent political party, and segmentation. We showed that this task is feasible to annotate, and present several models for automating the task. We use a pre-trained language model and also leverage auxiliary unlabeled text with semi-supervised learning approach for the target based speech act classification task. Our results are promising, with the best method being a semi-supervised biGRU with ELMo embeddings for the speech act task, and the model additionally incorporating speaker meta-data for the target party task. We provided qualitative analysis of speech acts across major policy areas, and in future work aim to expand this analysis further with fine-grained policies and ideology-related analysis.

## Acknowledgements

We thank the anonymous reviewers for their insightful comments and valuable suggestions. This work was funded in part by the Australian Government Research Training Program Scholarship, and the Australian Research Council.

## References

- J. Artés. 2011. Do Spanish politicians keep their election promises? *Party Politics*, 19(1):143–158.
- J. L. Austin. 1962. *How to do things with words*. Clarendon Press, Oxford.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit*. O’Reilly Media, Inc.
- Jean Carletta. 1996. Assessing agreement on classification tasks: the kappa statistic. *Computational Linguistics*, 22(2):249–254.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1724–1734.
- Jinho D Choi and Martha Palmer. 2012. Guidelines for the clear style constituent to dependency conversion. *Technical Report 01–12: Institute of Cognitive Science, University of Colorado Boulder*.
- Kevin Clark, Minh-Thang Luong, Christopher D Manning, and Quoc V Le. 2018. Semi-supervised sequence modeling with cross-view training. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1914–1925.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46.
- Rory Duthie and Katarzyna Budzynska. 2018. A deep modular RNN approach for ethos mining. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 4041–4047.
- Pablo Fernandez-Vazquez. 2014. And yet it moves: The effect of election platforms on party policy images. *Comparative Political Studies*, 47(14):1919–1944.
- Andrew Gibbons, Aaron Martin, and Andrea Carson. 2017. Do parties keep their election promises?: A study of the 43rd Australian Parliament (2010–2013). *Australian Political Studies Association*.
- Justin Grimmer and Brandon M Stewart. 2013. Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political analysis*, 21(3):267–297.
- Libby Hemphill and Andrew J Roback. 2014. Tweet Acts: How Constituents Lobby Congress via Twitter. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 1200–1210.
- Hugo Hernault, Danushka Bollegala, and Mitsuru Ishizuka. 2010. A sequential model for discourse segmentation. In *Proceedings of the International Conference on Intelligent Text Processing and Computational Linguistics*, pages 315–326.
- Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber, and Hal Daumé III. 2015. Deep unordered composition rivals syntactic methods for text classification. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 1681–1691.
- Minwoo Jeong, Chin-Yew Lin, and Gary Geunbae Lee. 2009. Semi-supervised speech act recognition in emails and forums. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 1250–1259.
- Shafiq Joty and Enamul Hoque. 2016. Speech act modeling of written asynchronous conversations with task-specific embeddings and conditional structured

- models. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 1746–1756.
- Shafiq Joty and Tasnim Mohiuddin. 2018. Modeling speech acts in asynchronous conversations: A neural-CRF approach. *Computational Linguistics*, pages 1–36.
- Klaus Krippendorff. 2011. Computing krippendorff’s alpha-reliability. *Technical Report, University of Pennsylvania*.
- David Lazer, Alex Sandy Pentland, Lada Adamic, Sinan Aral, Albert Laszlo Barabasi, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, et al. 2009. Life in the network: The coming age of computational social science. *Science*, 323(5915):721.
- Annie Louis and Ani Nenkova. 2011. Automatic identification of general and specific sentences by leveraging discourse annotations. In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 605–613.
- Christopher Lucas, Richard A Nielsen, Margaret E Roberts, Brandon M Stewart, Alex Storer, and Dustin Tingley. 2015. Computer-assisted text analysis for comparative politics. *Political Analysis*, 23(2):254–277.
- Luca Lugini and Diane Litman. 2017. Predicting specificity in classroom discussion. In *Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 52–61.
- Stefan Müller. 2018. Prospective and retrospective rhetoric: A new dimension of party competition and campaign strategies. In *Manifesto Corpus Conference*.
- Nona Naderi and Graeme Hirst. 2018. Using context to identify the language of face-saving. In *Proceedings of the 5th Workshop on Argument Mining*, pages 111–120.
- E. Naurin. 2011. *Election Promises, Party Behaviour and Voter Perception*. Palgrave Macmillan.
- E. Naurin. 2014. Is a promise a promise? Election pledge fulfillment in comparative perspective using Sweden as an example. *West European Politics*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1532–1543.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2227–2237.
- Ashequl Qadir and Ellen Riloff. 2011. Classifying sentences as speech acts in message board posts. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 748–758.
- Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. 2016. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 1163–1171.
- J. R. Searle. 1976. *A Taxonomy of Illocutionary Acts*. Linguistic Agency University of Trier.
- Matthew A. Shapiro, Libby Hemphill, and Jahna Otterbacher. 2018. Updates to congressional speech acts on Twitter. In *Proceedings of the American Political Science Association Annual Meeting*.
- Robert Thomson, Terry Royed, and Elin Naurin. 2010. The Program-to-Policy Linkage: A Comparative Study of Election Pledges and Government Policies in the United States, the United Kingdom, the Netherlands and Ireland. In *Proceedings of the American Political Science Association Annual Meeting*.
- Andrea Volkens, Pola Lehmann, Theres Matthieß, Nicolas Merz, Sven Regel, and Bernhard Weßels. 2017. *The Manifesto Data Collection. Manifesto Project (MRG/CMP/MARPOR). Version 2017b*. Wissenschaftszentrum Berlin für Sozialforschung.
- Soroush Vosoughi and Deb Roy. 2016. Tweet acts: A speech act classifier for Twitter. In *Proceedings of the Tenth International AAAI Conference On Web And Social Media*, pages 711–715.
- Xiang Wei, Boqing Gong, Zixia Liu, Wei Lu, and Liqiang Wang. 2018. Improving the improved training of Wasserstein GANs: A consistency term and its dual effect. In *Proceedings of the International Conference on Learning Representations*.
- Renxian Zhang, Dehong Gao, and Wenjie Li. 2012. Towards scalable speech act recognition in Twitter: tackling insufficient training data. In *Proceedings of the Workshop on Semantic Analysis in Social Media*, pages 18–27.
- Tianyu Zhao and Tatsuya Kawahara. 2017. Joint learning of dialog act segmentation and recognition in spoken dialog using neural networks. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing*, pages 704–712.
- M Zimmermann, A Stolcke, and E Shriberg. 2006. Joint segmentation and classification of dialog acts in multiparty meetings. In *Proceedings of the 2006 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 581–584.