

Our BERT-CNN model outperforms BERT and CNNs for classifying statements from party manifestos into policy domains.

Predicting Policy Domains from Party Manifestos with BERT and Convolutional Neural Networks

Allison Koh
Centre for International Security, Hertie School


Daniel Kai Sheng Boey
Columbia SIPA and Hertie School

Hannah Béchara
Data Science Lab, Hertie School

Background


Automated coding of text using supervised machine learning enables efficient labeling of political texts. We propose using Bidirectional Encoder Representations from Transformers (BERT) with Convolutional Neural Networks to classify statements parsed from party manifestos.

Data



- English text from the Manifesto Project Corpus
- 7 policy domains, 56 policy preferences
- Coding unit: quasi-sentence

Method



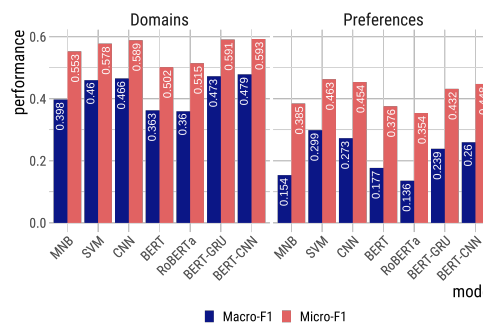
Our BERT-CNN model classifies text by domain (N = 8, incl. "no category") and preference (N = 57).

Experimental Details

We evaluate BERT-CNN against several baselines, which include **Multinomial Naive Bayes**, **Support Vector Machines**, **CNN**, **BERT**, **RoBERTa**, and **BERT-GRU** (BERT with gated recurrent units).

All BERT/RoBERTa models were trained on 10 epochs; others were trained on 100 epochs.

Results



BERT-CNN outperforms baselines for classification into domains, but not preferences.

Analysis

Category	Precision	Recall	F1 Score
Not categorized	0	0	0
External Relations	0.67	0.65	0.66
Freedom and Democracy	0.48	0.46	0.47
Political System	0.5	0.35	0.41
Economy	0.61	0.74	0.67
Welfare and Quality of Life	0.64	0.73	0.68
Fabric of Society	0.56	0.52	0.54
Social Groups	0.57	0.36	0.44

The categories that the model predicted best were: "welfare and quality of life", "economy", and "external relations". The model had difficulty predicting uncategorized statements.

The performance of our methods can benefit from:

- Increasing the number of epochs in training for deep learning models
- Appending training data for "underrepresented" categories

Future Work

Building off of the findings from this research, future work could investigate:

- The performance of other deep language representation models (e.g. RoBERTa)
- Using unsupervised ML to establish variation in features across policy domains.

