Analyzing Political Party Manifestos with Natural Language Processing

Salomon Orellana[0000−0002−9427−5499] and Halil Bisgin[0000−0001−6874−168X]

University of Michigan-Flint, Flint MI 48502, USA
{osalomon,bisgin}@umich.edu

Abstract. This study uses natural language processing (NLP) techniques to analyze the manifestos of political parties. It particularly uses topic modeling and sentiment analysis to analyze the manifestos produced by the political parties of New Zealand, from 1987 to 2017 (a period of significant party system change in New Zealand). Topic modeling may help objectively categorize and visualize the ideas political parties are promoting. Sentiment analysis has the potential to show each political party's attitude towards a policy idea. This study finds that these NLP techniques appear promising, though there is a need for significant fine-tuning.

Keywords: NLP · Sentiment Analysis · Political Texts · Electoral Reforms.

1 Introduction

Text is a common and rich source of data in political science research. Political scientists have analyzed text such as legislative speeches [1], newspapers [2], and party manifestos [3]. However, while there has been a proliferation of studies that use natural language processing (NLP) to analyze political texts, some of the best known and readily available sources of text in the field of political science have yet to be analyzed using NLP.

This study applies NLP techniques to one of the richest sources of text in political science: the manifestos of political parties. Party manifestos have been used by political scientists to objectively estimate various types of information produced by political parties, including: their ideological positions, the positions they take on particular policy areas, and the ideas they are trying to promote. Because manifestos can be gathered over many years and across countries, they also allow researchers to track the behavior of political parties across time and across many countries. The information gleaned from party manifestos has been used to explore many other political phenomena, including: the relation between parties and voters, the role of parties in parliament, and the translation of party programs into policy.

NLP has the potential to contribute to the analysis of manifestos in a variety of ways. It can be used to estimate: the similarity between documents, the topics or entities discussed in a document/corpus, and the sentiment expressed towards

entities. These techniques may help reveal: how party ideology evolves over time, how party positions on specific topics evolve over time, which parties may be most successful in promoting their policy ideas, and how political reforms can influence all of these dynamics. This study applies topic modeling and sentiment analysis to the analysis of manifestos. In later research, these techniques might reveal further insights when applied to other political texts.

This study specifically focuses on the manifestos produced by the political parties of New Zealand, from 1987 to 2017. During this period, New Zealand experienced significant party system change, which provides an opportunity to use a quasi-experimental research design. The study finds that NLP techniques appear promising for the analysis of long political texts, though there remains a need for significant fine-tuning.

2   Background

For most of its history, New Zealand's voters elected their legislators via a single-member-district-plurality (SMDP) system, in which the candidate winning the most votes (not necessarily a majority) wins the district/seat. In the 1990s, however, New Zealand switched to a German-style mixed-member proportional system (MMP) that, as noted by Sartori [4], produces relatively pure proportional results. Since 1996, over half (65 out of 120) of the members of the New Zealand parliament are still elected via SMDP, however, the remaining 55 seats are allocated via compensatory, proportional representation [5]. That is, the remaining seats are awarded to parties according to the percentage of the national vote that they receive, minus the number of seats they win in the single-member districts. Any party that wins at least 5 percent of the vote or that wins at least one seat in the single-member district part of the ballot is entitled to a share of legislative seats roughly proportional to the share of votes they received.

These changes along with factors such as a long-term trend towards party dealignment [6] led to a significant fragmentation of the party system. In 1993 (the last election under the SMDP system), the two largest parties (Labour and National) garnered over 95 percent of the seats, on slightly less than 70 percent of the vote; while in 1996 (the first election under the MMP system), the two largest parties captured only 67.5 percent of parliamentary seats on 62 percent of the vote. Furthermore, while only two parties won at least 5 percent of the seats in 1993, five parties won over 5 percent of the seats in 1996.

The electoral reforms and changes to the party system likely produced several information effects that may be captured by NLP methods. First, electoral and party systems should affect the number of issue areas considered by the political system. More proportional systems with higher numbers of parties should encourage political competition along a broader range of issue areas [7], especially because they make it easier for smaller parties to introduce new and sometimes controversial ideas. Second, proportional/multiparty systems should represent a wider range of positions on a left-right scale and in most other policy areas than do less-proportional/fewer-party systems [8].

Topic modeling may capture how the diversity of topics varies across electoral/party systems, while sentiment analysis may reveal the diversity of stances parties take on each topic.

3   Materials and Methods

3.1   Data

A significant advantage of analyzing manifestos is that they are readily accessible via the Manifesto Project Database. This database grew out of the efforts of the Manifesto Research Group (MRG 1979-1989) and the Comparative Manifestos Project (CMP 1989-2009), and is now currently maintained and expanded by Manifesto Research on Political Representation (MARPOR) [9]. The Manifesto Project Database makes available the manifestos of over 1000 political parties, from over 50 countries, from 1945 to the present. The database also makes available the results of hand coded content analyses for every manifesto. This hand coding required the support of coders from over 50 countries and can be leveraged to gauge the performance of NLP analyses.

This study particularly analyzes a subset of the manifesto corpus: the political party manifestos produced by New Zealand's political parties, for the elections that took place between the years 1987 and 2017. New Zealand holds elections every three years, and that amounts to 12 elections, during that period. Roughly 11 parties produced 61 manifestos during that time period. Prior to the electoral reforms (in 1996), only 3 parties were producing manifestos - that were collected by researchers. By 2017, roughly 6 parties were producing manifestos.

These manifestos were analyzed using the following NLP techniques: identifying the topics discussed in documents and sentiment analysis. Topic modeling techniques are used to detect the topics parties promote and if there were more topics discussed by the party system as whole after the electoral reforms. Sentiment analysis techniques are used to gauge a party's position on the topics discovered by topic modeling. The methodology of each of these techniques is briefly discussed below.

3.2   Documents as Vectors

NLP techniques require that documents first be converted to mathematical objects, i.e., vectors. There are multiple ways to turn documents into vectors. One vector type commonly used is the TF-IDF (term-frequency times inverse document frequency) vector, which is a frequency vector that adjusts the count of each word by the number of documents in which the word occurs.

Another vector type is Word2vec, which accounts for the words around each word, how the neighbors of a word affect the overall meaning of a statement. It learns the meaning of words by processing a large corpus of unlabeled text; and it searches for words that tend to be near each other. All words in a corpus are represented by numerical vectors containing scores associated with dimensions of meaning.

Another vector representation rising in popularity is BERT (Bidirectional Encoder Representation from Transformers). BERT was developed by Google and is considered a state-of-the-art NLP technique. It is trained on unlabeled text including Wikipedia and Book corpuses. BERT uses transformer architecture (an attention model) to learn embeddings for words. It involves two pre-training steps: Masked Language Modelling (MLM) and Next Sentence Prediction (NSP). The training text is represented via three embeddings: token embeddings, segment embeddings, and position embeddings.

3.3   Topic Modeling

This study applies NLP techniques to identify the topics political parties are discussing. It is expected that the topics discussed and emphasized vary across political parties and across time, and thus, unsupervised topic modeling techniques are required. This study specifically applies the topic modeling technique developed by Grootendorst [10], which leverages sentence transformers and c-TF-IDF (a class-based variant of TF-IDF) to create dense clusters that yield interpretable topics.[1]

In this technique, BERT transformer embeddings are used to convert documents into numerical data. BERT extracts word embeddings based on the context of the word. Many pre-trained models are available, which helps reduce processing time and can produce more accurate representations of words and sentences. To find topics, documents with similar topics are clustered together, using the UMAP (Uniform Manifold Approximation and Projection) dimensionality reduction algorithm. UMAP stores a high portion of the information of the high-dimensional local structure in a structure of lower dimensionality.

After dimensionality reduction, the documents are clustered using HDB-SCAN (Hierarchical Density-Based Spatial Clustering of Applications with Noise), a density-based algorithm. HDBSCAN does not force all data points to cluster, it treats some as outliers. The parameters of HDBSCAN can be tweaked to control the number of topics returned, which may be necessary when the system is returning hundreds of topics.

To derive topics from clustered documents (to extract what makes each set of documents unique compared to the other), a class-based variant of TF-IDF (c-TF-IDF) is used. TF-IDF algorithms basically compare the importance of words between documents. Here, a cluster is treated as a single document, to which a TF-IDF algorithm is applied. The resulting TF-IDF score would demonstrate the important words in a topic. The higher the score, the more representative it should be of its topic – the score is a proxy of information density. The top 10 to 20 words (with the highest scores) can be selected to visualize what may be the focus of a topic. Topic sizes can be used to gauge how frequently certain topics appear.

---

[1] The topic modeling presented here was carried out using Grootendorst's Python pacakge, BERTopic

3.4   Sentiment Analysis

Many sentiment analysis techniques attempt to determine the sentiment (positive or negative) of documents or sentences, and even parts of sentences (aspects). Much of the ABSA research conducted to date focuses on the analysis of product reviews, which tend to be short and relatively structured – at least, relative to longer documents, such as the manifestos analyzed in this study. Thus, this study currently focuses on detecting sentiment in sentences, particularly by applying the VADER (Valence Aware Dictionary for sEntiment Reasoning) algorithm to each sentence in a manifesto that contains at least one of the words returned by the topic modeling analysis.

VADER is a rule-based sentiment analysis algorithm developed by Hutto and Gilbert [11]. VADER sentimental analysis relies on a dictionary that maps lexical features to emotion intensities known as sentiment scores. Hutto maintains the dictionary and Python package used in this analysis.[2]

The algorithm finds keywords in a document and maps each one to numerical scores or weights in the dictionary (or "mapping"). Stems, lemmas, or n-gram tokens can be used in the dictionary. Vader recognizes that words such as 'love', 'enjoy', 'happy', 'like' convey a positive sentiment. It is also intelligent enough to understand the basic context of these words, particularly modifying words that can change the sentiment. It can recognize that a statement such as "did not love" is a negative statement. It also understands the emphasis implied by capitalization and punctuation (such as "ENJOY"), as well as emoticons, such as happy faces.

Scores are summed for each keyword in a document. The sentiment score of a text can be obtained by summing up the intensity of each word in the text. Thus, VADER is able to address both polarity (positive/negative) and intensity of emotion.

4   Results

4.1   Topic Modeling of Manifestos

Topic modeling can help reveal what policy issues the political parties wish to emphasize. As mentioned, one expectation is that more proportional systems with higher numbers of parties should encourage political competition along a broader range of issue areas. Thus, we should expect that New Zealand's party system brought attention to more issues after its reforms than before. This is where the quasi-experimental analysis is more directly applied.

This section compares the number of issues covered by the party system in the three elections before the reform (1987, 1990, 1993) and the three elections after the reform (1999, 2002, 2005). Topic modeling analysis was performed on all the manifestos appearing during each of these periods. The 1996 election is left out, because it was the first election under the new MMP system. Thus, it

---

[2] The Python package is called vaderSentiment.

is regarded here as a transition period. This design thus matches Campbell and Stanley's description of a one-group pretest-posttest quasi-experimental design [12], where a group (New Zealand's party system) is observed prior to being exposed to an event (major electoral reform) and then it is observed again after the event.

In figures 1 and 2, it can first be observed that the topic modeling approach returned many topic areas, for each of the two periods. It can also be observed that more distinct topics were discussed after the reforms than before. In the three elections before the reform, roughly 104 issues were discussed by the political parties. In the three elections after the reforms, roughly 170 issues were discussed by the political parties. When overlapping is taken into account, the conclusion is more pronounced. In the three elections before the reform, roughly 18 issues were discussed by the political parties. In the three elections after the reforms, roughly 34 issues were discussed by the political parties.

The figures also provide more detail about the largest (most frequent) topics. The largest topics are numbered in the figure and their word composition is displayed in the caption below the figure. The largest topics in figure 1 include health care, families, fishing, and so on. The largest topics in figure 2 include transport, taxation, housing, and so on.

The topics produced by this topic modeling approach appear reasonable. Table 1 shows a topic cluster that appears to be about transportation and more specifically about environmentally friendly transportation:

Table 1. Example Topic

Top 10 Words	Importance Value
'transport'	0.065317291
'rail'	0.04883337
'electric'	0.03538137
'trains'	0.027568078
'buses'	0.027568078
'cars'	0.024513584
'cycling'	0.023804225
'auckland'	0.023682579
'safer'	0.02330147
'vehicles'	0.022399063

The full set of words captures a focus on transportation. A subset of the words – rail, trains, buses, cycling – suggests this topic is more about public transportation and more specifically about environmentally friendly forms of transportation.
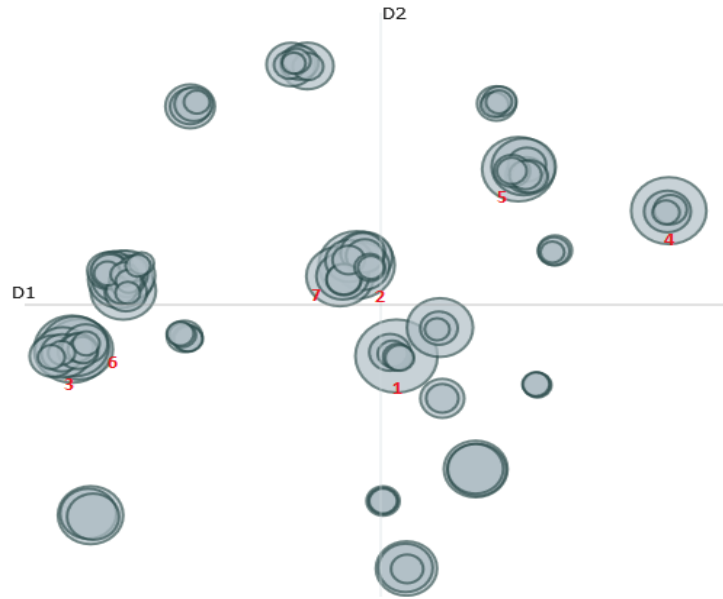
Fig. 1. Intertopic Distance Map: Pre-reform Topics
Topic 1: 'health', 'hospitals', 'public', 'care', 'hospital', 'system', 'patients', 'services', 'quality', 'national' Topic 2: 'families', 'children', 'family', 'parents', 'childcare', 'child', 'support', 'social', 'subsidy', 'stress' Topic 3: 'fishing', 'fisheries', 'fish', 'quota', 'aqua-culture', 'recreational', 'management', 'commercial', 'industry', 'percent' Topic 4: 'rela-tions', 'industrial', 'unions', 'employment', 'union', 'workers', 'bargaining', 'employers', 'contracts', 'equity' Topic 5: 'trade', 'export', 'markets', 'imports', 'products', 'tar-iffs', 'marketing', 'market', 'overseas', 'goods' Topic 6: 'environmental', 'environment', 'conservation', 'natural', 'management', 'sustainable', 'resource', 'resources', 'accounts', 'mineral' Topic 7: 'tax', 'taxation', 'small', 'taxes', 'businesses', 'business', 'duties', 'ap-plies', 'definition', 'which'.

4.2   Sentiment Analysis of Manifesto Sentences

As mentioned, this study focuses on detecting sentiment in sentences, particu-larly by applying the VADER algorithm to each sentence in a manifesto that contains at least one of the words returned by the topic modeling analysis. Here it is applied to the sentences containing the words associated with environmen-tally friendlier transportation: rail, trains, buses, cycling. That is, to estimate the support that each political party shows towards environmentally friendlier forms of transportation, the polarity of each sentence containing at least one of these words is calculated. Then, the average is calculated for all the sentences appearing in each party's manifesto.

Table 2 shows the scores returned by VADER for a small sample of sentences from the Green Party's manifesto. Each sentence contains at least one of the key
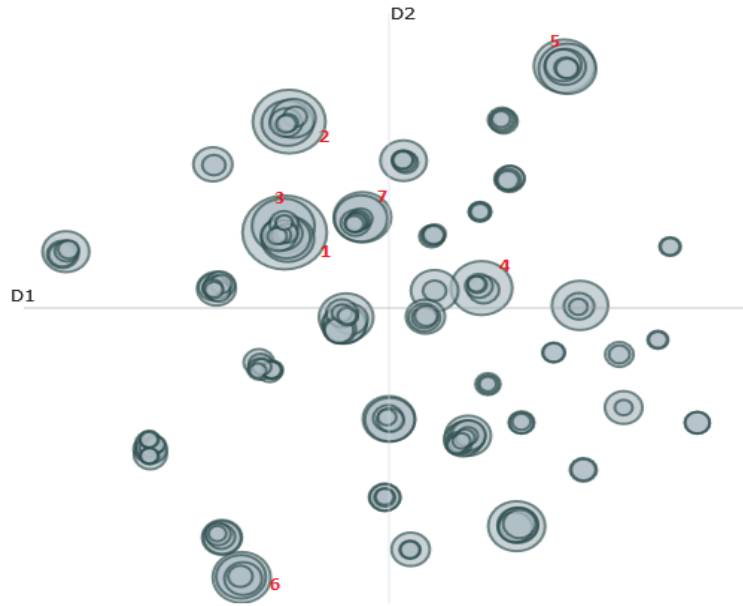
Fig. 2. Intertopic Distance Map: Post-reform Topics
Topic 1: 'transport', 'road', 'rail', 'roads', 'roading', 'safety', 'traffic', 'infrastructure', 'congestion', 'transit' Topic 2: 'tax', 'taxation', 'taxes', 'rate', 'income', 'lower', 'company', 'corporate', 'personal', 'higher' Topic 3: 'housing', 'tenants', 'houses', 'state', 'income', 'low', 'rental', 'rent', 'accommodation', 'house' Topic 4: 'defence', 'force', 'forces', 'maritime', 'air', 'capability', 'allies', 'navy', 'peace', 'army' Topic 5: 'they', 'but', 'we', 'delivered', 'it', 'what', 'let', 'now', 'more', 'win' Topic 6: 'radio', 'broadcasting', 'television', 'tvnz', 'spectrum', 'fm', 'air', 'music', 'tv', 'nz' Topic 7: 'export', 'business', 'value', 'small', 'growth', 'economy', 'exports', 'commerce', 'wealth', 'businesses'.

words (rail, trains, buses, or cycling). The polarity appears relatively positive, as may be expected from the Green Party (GP) on this issue.

To get a sense of the total support each party shows on this issue, an average score is calculated across all the sentences containing the key words. The average score for each party, along with the number of sentences containing the key words, is reported in table 3. The average scores appear reasonable. The Green Party appears to show by far the highest level of support for environmentally friendlier forms of transportation. It has the highest number of related sentences (65) and the highest average polarity (0.36). It is much less of an emphasis for other parties, and they have considerably less positivity towards that issue.

While these results appear quite reasonable, there are some concerns left to address. The current approach likely suffices to capture a party's attitude (polarity) towards an issue, particularly when the goal is to compare the polarity of several parties, but the accuracy is likely reduced by applying VADER to complete sentences – which may contain more than one target topic. Problems

Table 2. Sample of VADER Scores for Each Sentence

Party	Text	Polarity Vader
GP	Support locating clusters near transport hubs (rail lines, ports etc).	0.4019
GP	Expand the nationwide network of cycle/pedestrian trails.	0.3182
GP	Promote rail as a great way to travel and seek to make it more available and reliable.	0.7717
GP	All goods and services produced or sold in New Zealand to meet quality and sustainability standards (eg energy and recycling standards).	0.2732
GP	Fast, electric rail lines eliminate pollution and create healthier, congestion-free cities.	0.2732
GP	Safe walking and cycling for kids.	0.4404

Table 3. Average VADER Score for Each Party on Environmentally Friendlier Transportation

Party	Number of Sentences	Avg Polarity Vader
ACT	-	-
GP	65	0.359523
LP	18	0.1463
Maori	2	0.08895
NP	29	0.15811
NZF	7	0.025414

such as multiple targets, are likely addressed more effectively by aspect based sentiment analysis (ABSA) methods. However, most ABSA techniques have been developed for shorter, more structured documents, such as product reviews. Thus most ABSA techniques underperform the VADER method used in this study. This is a promising area to continue researching.

5   Conclusion

Applying NLP to key texts analyzed by political scientists appears promising. It was possible to observe that topic modeling can reveal the policy issues the parties talk about and how the diversity of the political information environment is affected by major electoral reforms. It was also possible to see that sentiment analysis can capture the support parties show towards particular policy ideas.

The analyses performed here, however, can likely benefit from significant fine-tuning. First, each analysis likely could benefit from more appropriate vector representations, especially vector representations that more accurately captured a party's attitude towards various policy issues. More fine-tuned ABSA techniques, for example, might improve the accuracy of measurements of a party's attitude (polarity) on an issue.

Manifestos are also limited in the insights they can provide about the political information environment. Thus, it is hoped that the techniques applied to

manifestos in this study will lead to more substantive insights when applied to other texts in the political realm.

References

1. K. T. Poole and H. Rosenthal, "D-nominate after 10 years: A comparative update to congress: A political-economic history of roll-call voting," *Legislative Studies Quarterly*, pp. 5–29, 2001.
2. J. Hayward and C. Rudd, "Metropolitan newspapers and the election," *Left Turn: The New Zealand General Election of*, pp. 89–104, 1999.
3. A. Volkens, J. Bara, I. Budge, M. D. McDonald, and H.-D. Klingemann, *Mapping policy preferences from texts: statistical solutions for manifesto analysts*, vol. 3. OUP Oxford, 2013.
4. G. Sartori, *Comparative constitutional engineering: an inquiry into structures, incentives, and outcomes*. New York University Press, 1994. Giovanni Sartori.¡br/¿22 cm.
5. R. J. Johnston and C. J. Pattie, "Campaigning and split-ticket voting in new electoral systems: the first mmp elections in new zealand, scotland and wales," *Electoral Studies*, vol. 21, p. 583, 2002.
6. J. Vowles, "The politics of electoral reform in new zealand," *International Political Science Review*, vol. 16, pp. 95–115, 1995.
7. A. Lijphart, *Patterns of democracy: government forms and performance in thirty-six countries*. Yale University Press, 1999. Arend Lijphart.¡br/¿22 cm.
8. R. S. Katz, *Democracy and elections*. Oxford University Press, 1997.
9. A. Volkens, T. Burst, W. Krause, P. Lehmann, T. Matthieß, N. Merz, B. Weßels, and L. Zehnter, "The manifesto data collection. manifesto project (mrg/cmp/marpor)," 2021.
10. M. Grootendorst, "Bertopic: Leveraging bert and c-tf-idf to create easily interpretable topics," *Zenodo*, 2020.
11. C. Hutto and E. Gilbert, "Vader: A parsimonious rule-based model for sentiment analysis of social media text," vol. 8, pp. 216–225, 2014.
12. D. Campbell and J. Stanley, *Experimental and Quasi-experimental Designs for Research*. Houghton Mifflin Company, 1963.