

Chapter 7

Natural Language Processing for Policymaking



Zhijing Jin and Rada Mihalcea

Abstract Language is the medium for many political activities, from campaigns to news reports. Natural language processing (NLP) uses computational tools to parse text into key information that is needed for policymaking. In this chapter, we introduce common methods of NLP, including text classification, topic modelling, event extraction, and text scaling. We then overview how these methods can be used for policymaking through four major applications including data collection for evidence-based policymaking, interpretation of political decisions, policy communication, and investigation of policy effects. Finally, we highlight some potential limitations and ethical concerns when using NLP for policymaking.

7.1 Introduction

Language is an important form of data in politics. Constituents express their stances and needs in text such as social media and survey responses. Politicians conduct campaigns through debates, statements of policy positions, and social media. Government staff needs to compile information from various documents to assist in decision-making. Textual data is also prevalent through the documents and debates in the legislation process, negotiations and treaties to resolve international conflicts, and media such as news reports, social media, party platforms, and manifestos.

Natural language processing (NLP) is the study of computational methods to automatically analyse text and extract meaningful information for subsequent analysis. The importance of NLP for policymaking has been highlighted since the

Z. Jin

Max Planck Institute for Intelligent Systems, Tübingen, Germany

ETH Zürich, Zürich, Switzerland

e-mail: zhij.jin@gmail.com

R. Mihalcea (✉)

University of Michigan, Ann Arbor, MI, USA

e-mail: mihalcea@umich.edu

last century (Gigley, 1993). With the recent success of NLP and its versatility over tasks such as classification, information extraction, summarization, and translation (Brown et al., 2020; Devlin et al., 2019), there is a rising trend to integrate NLP into the policy decisions and public administrations (Engstrom et al., 2020; Misuraca et al., 2020; Van Roy et al., 2021). Main applications include extracting useful, condensed information from free-form text (Engstrom et al., 2020), and analysing sentiment and citizen feedback by NLP Biran et al. (2022) as in many projects funded by EU Horizon projects (European Commission, 2017). Driven by the broad applications of NLP (Jin et al., 2021a), the research community also starts to connect NLP with various social applications in the fields of computational social science (Engel et al., 2021; Lazer et al., 2009; Luz, 2022; Shah et al., 2015) and political science in particular (Glavaš et al., 2019; Grimmer & Stewart, 2013).

We show an overview of NLP for policymaking in Fig. 7.1. According to this overview, the chapter will consist of three parts. First, we introduce in Sect. 7.2 NLP methods that are applicable to political science, including text classification, topic modelling, event extraction, and score prediction. Next, we cover a variety of cases where NLP can be applied to policymaking in Sect. 7.3. Specifically, we cover four stages: analysing data for evidence-based policymaking, improving policy communication with the public, investigating policy effects, and interpreting political phenomena to the public. Finally, we will discuss limitations and ethical considerations when using NLP for policymaking in Sect. 7.4.

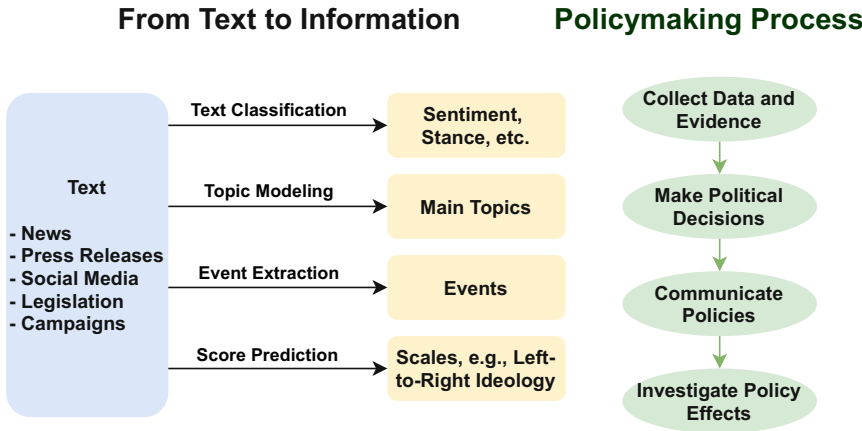


Fig. 7.1 Overview of NLP for policymaking

7.2 NLP for Text Analysis

NLP brings powerful computational tools to analyse textual data (Jurafsky & Martin, 2000). According to the type of information that we want to extract from the text, we introduce four different NLP tools to analyse text data: text classification (by which the extracted information is the *category* of the text), topic modelling (by which the extracted information is the *key topics* in the text), event extraction (by which the extracted information is the list of *events* mentioned in the text), and score prediction (where the extracted information is a *score* of the text). Table 7.1 lists each method with the type of information it can extract and some example application scenarios, which we will detail in the following subsections.

7.2.1 Text Classification

As one of the most common types of text analysis methods, text classification reads in a piece of text and predicts its category using an NLP text classification model, as in Fig. 7.2.

Table 7.1 Four common NLP methods, the type of information extracted by each of them, and example applications

| NLP method | Information to extract | Example applications |
|---------------------|------------------------|--|
| Text classification | Category of text | Identify the sentiment, stance, etc. |
| Topic modelling | Key topics in text | Summarize topics in political agenda |
| Event extraction | List of events | Extract news events, international conflicts |
| Score prediction | Score | Text scaling |

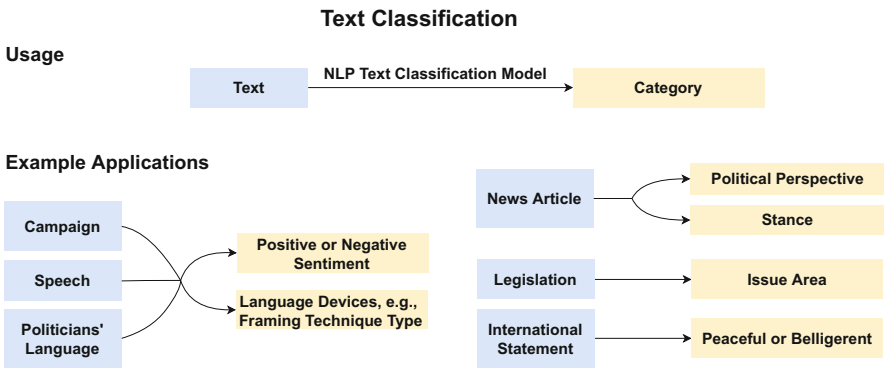


Fig. 7.2 The usage and example applications of text classification on political text

There are many off-the-shelf existing tools for text classification (Brown et al., 2020; Loria, 2018; Yin et al., 2019) such as the implementation¹ using the Python package `transformers` (Wolf et al., 2020). A well-known subtask of text classification is sentiment classification (also known as sentiment analysis or opinion mining), which aims to distinguish the subjective information in the text, such as positive or negative sentiment (Pang & Lee, 2007). However, the existing tools only do well in categories that are easy to predict. If the categorization is customized and very specific to a study context, then there are two common solutions. One is to use dictionary-based methods, by a list of frequent keywords that correspond to a certain category (Albaugh et al., 2013) or using general linguistic dictionaries such as the Linguistic Inquiry and Word Count (LIWC) dictionary (Pennebaker et al., 2001). The second way is to adopt the data-driven pipeline, which requires human hand coding of documents into a predetermined set of categories, then train an NLP model to learn the text classification task (Sun et al., 2019), and verify the performance of the NLP model on a held-out subset of the data, as introduced in Grimmer and Stewart (2013). An example of adapting the state-of-the-art NLP models on a customized dataset is demonstrated in this guide.²

Using the text classification method, we can automate many types of analyses in political science. As listed in the examples in Fig. 7.2, researchers can detect political perspective of news articles (Huguet Cabot et al., 2020), the stance in media on a certain topic (Luo et al., 2020), whether campaigns use positive or negative sentiment (Ansolabehere & Iyengar, 1995), which issue area is the legislation about (Adler & Wilkerson, 2011), topics in parliament speech (Albaugh et al., 2013; Osnabrügge et al., 2021), congressional bills (Collingwood & Wilkerson, 2012; Hillard et al., 2008) and political agenda (Karan et al., 2016), whether the international statement is peaceful or belligerent (Schrodt, 2000), whether a speech contains positive or negative sentiment (Schumacher et al., 2016), and whether a US Circuit Courts case decision is conservative or liberal (Hausladen et al., 2020). Moreover, text classification can also be used to categorize the type of language devices that politicians use, such as what type of framing the text uses (Huguet Cabot et al., 2020), and whether a tweet uses political parody (Maronikolakis et al., 2020).

7.2.2 Topic Modelling

Topic modelling is a method to uncover a list of frequent topics in a corpus of text. For example, news articles that are against vaccination might frequently mention the topic “autism”, whereas news articles supporting vaccination will be more likely to mention “immune” and “protective”. One of the most widely used models is the

¹ <https://discuss.huggingface.co/t/new-pipeline-for-zero-shot-text-classification/681>.

² <https://skimai.com/fine-tuning-bert-for-sentiment-analysis/>.




Fig. 7.3 Given a collection of text documents, topic modelling generates a list of topic clusters

Latent Dirichlet Allocation (LDA) (Blei et al., 2001) which is available in the Python packages NLTK and Gensim, as in this guide.³

Specifically, LDA is a probabilistic model that models each topic as a mixture of words, and each textual document can be represented as a mixture of topics. As in Fig. 7.3, given a collection of textual documents, LDA topic modelling generates a list of topic clusters, for which the number N of topics can be customized by the analyst. In addition, if needed, LDA can also produce a representation of each document as a weighted list of topics. While often the number of topics is predetermined by the analyst, this number can also be dynamically determined by measuring the perplexity of the resulting topics. In addition to LDA, other topic modelling algorithms have been used extensively, such as those based on principal component analysis (PCA) (Chung & Pennebaker, 2008).

Topic modelling, as described in this section, can facilitate various studies on political text. Previous studies analysed the topics of legislative speech (Quinn et al., 2006, 2010), Senate press releases (Grimmer, 2010a), and electoral manifestos (Menini et al., 2017).

7.2.3 Event Extraction

Event extraction is the task of extracting a list of events from a given text. It is a subtask of a larger domain of NLP called information extraction (Manning et al., 2008). For example, the sentence “Israel bombs Hamas sites in Gaza” expresses an event “*Israel* $\xrightarrow{\text{bombs}}$ *Hamas sites*” with the location “*Gaza*”. Event extraction usually incorporates both entity extraction (e.g. Israel, Hamas sites, and Gaza in the previous example) and relation extraction (e.g. “bombs” in the previous example).

Event extraction is a handy tool to monitor events automatically, such as detecting news events (Mitamura et al., 2017; Walker et al., 2006) and detecting international conflicts (Azar, 1980; Trappl, 2006). To foster research on event extraction, there are tremendous efforts into textual data collection (McClelland,

³ <https://skimai.com/fine-tuning-bert-for-sentiment-analysis/>.

1976; Merritt et al., 1993; Raleigh et al., 2010; Schrodtt & Hall, 2006; Sundberg & Melander, 2013), event coding schemes to accommodate different political events (Bond et al., 1997; Gerner et al., 2002; Goldstein, 1992), and dataset validity assessment (Schrodtt & Gerner, 1994).

As for event extraction models, similar to text classification models, there are off-the-shelf tools such as the Python packages *stanza* (Qi et al., 2020) and *spaCy* (Honnibal et al., 2020). In case of customized sets of event types, researchers can also train NLP models on a collection of textual documents with event annotations (Hogenboom et al., 2011; Liu et al., 2020, *inter alia*).

7.2.4 Score Prediction

NLP can also be used to predict a score given input text. A useful application is political text scaling, which aims to predict a score (e.g. left-to-right ideology, emotionality, and different attitudes towards the European integration process) for a given piece of text (e.g. political speeches, party manifestos, and social media posts) (Gennaro & Ash, 2021; Laver et al., 2003; Lowe et al., 2011; Slapin & Proksch, 2008, *inter alia*).

Traditional models for text scaling include Wordscores (Laver et al., 2003) and WordFish (Lowe et al., 2011; Slapin & Proksch, 2008). Recent NLP models represent the text by high-dimensional vectors learned by neural networks to predict the scores (Glavaš et al., 2017b; Nanni et al., 2019). One way to use the NLP models is to apply off-the-shelf general-purpose models such as InstructGPT (Ouyang et al., 2022) and design a prompt to specify the type of the scaling to the API,⁴ or borrow existing, trained NLP models if the same type of scaling has been studied by previous researchers. Another way is to collect a dataset of text with hand-coded scales, and train NLP models to learn to predict the scale, similar to the practice in Gennaro and Ash (2021); Slapin and Proksch (2008), *inter alia*.

7.3 Using NLP for Policymaking

In the political domain, there are large amounts of textual data to analyse (NEUENDORF & KUMAR, 2015), such as parliament debates (Van Aggelen et al., 2017), speeches (Schumacher et al., 2016), legislative text (Baumgartner et al., 2006; Bevan, 2017), database of political parties worldwide (Döring & Regel, 2019), and expert survey data (Bakker et al., 2015). Since it is tedious to hand-code all textual data, NLP provides a low-cost tool to automatically analyse such massive text.

⁴ <https://beta.openai.com/docs/introduction>.

In this section, we will introduce how NLP can facilitate four major areas to help policymaking: before policies are made, researchers can use NLP to analyse data and extract key information for evidence-based policymaking (Sect. 7.3.1); after policies are made, researchers can interpret the priorities among and reasons behind political decisions (Sect. 7.3.2); researchers can also analyse features in the language of politicians when communicating the policies to the public (Sect. 7.3.3); and finally, after the policies have taken effect, researchers can investigate the effectiveness of the policies (Sect. 7.3.4).

7.3.1 *Analysing Data for Evidence-Based Policymaking*

A major use of NLP is to extract information from large collections of text. This function can be very useful for analysing the views and needs of constituents, so that policymakers can make decisions accordingly.

As in Fig. 7.4, we will explain how NLP can be used to analyse data for evidence-based policymaking from three aspects: data, information to extract, and political usage.

Data Data is the basis of such analyses. Large amounts of textual data can reveal information about constituents, media outlets, and influential figures. The data can come from a variety of sources, including social media such as Twitter and Facebook, survey responses, and news articles.

Information to Extract Based on the large textual corpora, NLP models can be used to extract information that are useful for political decision-making, ranging from information about people, such as sentiment (Rosenthal et al., 2015; Thelwall et al., 2011), stance (Gottipati et al., 2013; Luo et al., 2020; Stefanov et al., 2020; Thomas et al., 2006), ideology (Hirst et al., 2010; Iyyer et al., 2014; Preoțiuc-Pietro et al., 2017), and reasoning on certain topics (Camp et al., 2021; Demszky et al.,

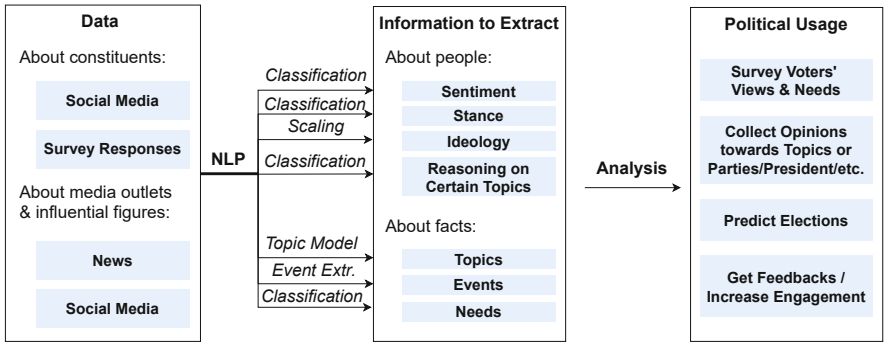


Fig. 7.4 NLP to analyse data for evidence-based policymaking

2019; Egami et al., 2018), to factual information, such as main topics (Gottipati et al., 2013), events (Ding & Riloff, 2018; Ding et al., 2019; Mitamura et al., 2017; Trappl, 2006), and needs (Crayton et al., 2020; Paul & Frank, 2019; Sarol et al., 2020) expressed in the data. The extracted information cannot only be about people but also about political entities, such as the left-right political scales of parties and political actors (Glavaš et al., 2017b; Slapin & Proksch, 2008), which claims are raised by which politicians (Blessing et al., 2019; Padó et al., 2019), and the legislative body's vote breakdown for state bills by backgrounds such as gender, rural-urban, and ideological splits (Davoodi et al., 2020).

To extract such information from text, we can often utilize the main NLP tools introduced in Sect. 7.2, including text classification, topic modelling, event extraction, and score prediction (especially text scaling to predict left-to-right ideology). In NLP literature, social media, such as Twitter, is a popular source of textual data to collect public opinions (Arunachalam & Sarkar, 2013; Pak & Paroubek, 2010; Paltoglou & Thelwall, 2012; Rosenthal et al., 2015; Thelwall et al., 2011).

Political Usage Such information extracted from data is highly valuable for political usage. For example, voters' sentiment, stance, and ideology are important supplementary for traditional polls and surveys to gather information about the constituents' political leaning. Identifying the needs expressed by people is another important survey target, which helps politicians understand what needs they should take care of and match the needs and availabilities of resources (Hiware et al., 2020).

Among more specific political uses is to understand the public opinion on parties/president, as well as on certain topics. The public sentiment towards parties (Pla & Hurtado, 2014) and president (Marchetti-Bowick & Chambers, 2012) can serve as a supplementary for the traditional approval rating survey, and stances towards certain topics (Gottipati et al., 2013; Luo et al., 2020; Stefanov et al., 2020) can be important information for legislators to make decisions on debatable issues such as abortion, taxes, and legalization of same-sex marriage. Many existing studies use NLP on social media text to predict election results (Beverungen & Kalita, 2011; Mohammad et al., 2015; O'Connor et al., 2010; Tjong Kim Sang & Bos, 2012; Unankard et al., 2014). In general, big data-driven analyses can facilitate decision-makers to collect more feedback from people and society, enabling policymakers to be closer to citizens, and increase transparency and engagement in political issues (Arunachalam & Sarkar, 2013).

7.3.2 *Interpreting Political Decisions*

After policies are made, political scientists and social scientists can use textual data to interpret political decisions. As in Fig. 7.5, there are two major use cases: mining political agendas and discovering policy responsiveness.




Fig. 7.5 NLP to interpret political decisions

Mining Political Agendas Researchers can use textual data to infer a political agenda, including the topics that politicians prioritize, political events, and different political actors’ stances on certain topics. Such data can come from press releases, legislation, and electoral campaigns. Examples of previous studies to analyse the topics and prioritization of political bodies include the research on the prioritization each senator assigns to topics using press releases (Grimmer, 2010b), topics in different parties’ electoral manifestos (Glavaš et al., 2017a), topics in EU parliament speeches (Lauscher et al., 2016) and other various types of text (Grimmer, 2010a; Hopkins & King, 2010; King & Lowe, 2003; Roberts et al., 2014), as well as political event detection from congressional text and news (Nanni et al., 2017).

Research on politicians’ stances include identifying policy positions of politicians (Laver et al., 2003; Lowe et al., 2011; Slapin & Proksch, 2008; Winter & Stewart, 1977, inter alia), how different politicians agree or disagree on certain topics in electoral campaigns (Menini & Tonelli, 2016), and assessment of political personalities (Immelman, 1993).

Further studies look into how political interests affect legislative behaviour. Legislators tend to show strong personal interest in the issues that come before their committees (Fenno, 1973), and Mayhew (2004) identifies that senators replying on appropriations secured for their state have a strong incentive to support legislations that allow them to secure particularistic goods.

Discovering Policy Responsiveness Policy responsiveness is the study of how policies respond to different factors, such as how changes in public opinion lead to responses in public policy (Stimson et al., 1995). One major direction is that politicians tend to make policies that align with the expectations of their constituents, in order to run for successful re-election in the next term (Canes-Wrone et al., 2002). Studies show that policy preferences of the state public can be a predictor of future state policies (Caughey & Warshaw, 2018). For example, Lax and Phillips (2009) show that more LGBT tolerance leads to more pro-gay legislation in response.

A recent study by Jin et al. (2021b) uses NLP to analyse over 10 million COVID-19-related tweets targeted at US governors; using classification models to obtain the public sentiment, they study how public sentiment leads to political decisions of COVID-19 policies made by US governors. Such use of NLP on massive textual

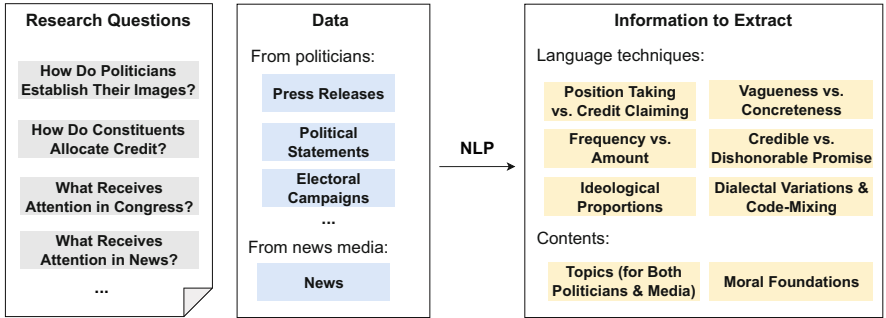


Fig. 7.6 NLP to analyse policy communication

data contrasts with the traditional studies of policy responsiveness which span over several decades and use manually collected survey results (Caughey & Warshaw, 2018; Lax & Phillips, 2009, 2012).

7.3.3 Improving Policy Communication with the Public

Policy communication is the study to understand how politicians present the policies to their constituents. As in Fig. 7.6, common research questions in policy communication include how politicians establish their images (Fenno, 1978) such as campaign strategies (Petrocik, 1996; Sigelman & Buell Jr, 2004; Simon, 2002), how constituents allocate credit, what receives attention in Congress (Sulkin, 2005), and what receives attention in news articles (Armstrong et al., 2006; McCombs & Valenzuela, 2004; Semetko & Valkenburg, 2000).

Based on data from press releases, political statements, electoral campaigns, and news articles,⁵ researchers usually analyse two types of information: the language techniques politicians use and the contents such as topics and underlying moral foundations in these textual documents.

Language Techniques Policy communication largely focuses on the types of languages that politicians use. Researchers are interested in first analysing the language techniques in political texts, and then, based on these techniques, researchers can dive into the questions of why politicians use them and what are the effects of such usage.

For example, previous studies analyse what portions of political texts are position-taking versus credit-claiming (Grimmer, 2013; Grimmer et al., 2012),

⁵ Other data sources used in policy communication research include surveys of senate staffers (Cook, 1988), newsletters that legislators send to constituents (Lipinski, 2009), and so on.

whether the claims are vague or concrete (Baerg et al., 2018; Eichorst & Lin, 2019), the frequency of credit-claiming messages versus the actual amount of contributions (Grimmer et al., 2012), and whether politicians tend to make credible or dishonourable promises (Grimmer, 2010b). Within the political statements, it is also interesting to check the ideological proportions (Sim et al., 2013) and how politicians make use of dialectal variations and code-mixing (Sravani et al., 2021).

The representation styles usually affect the effectiveness of policy communication, such as the role of language ambiguity in framing the political agenda (Campbell, 1983; Page, 1976) and the effect of credit-claiming messages on constituents' allocation of credit (Grimmer et al., 2012).

Contents The contents of policy communication include the topics in the political statements, such as what senators discuss in floor statements (Hill & Hurley, 2002) and what presidents address in daily speeches (Lee, 2008), and also the moral foundations used by politicians underlying their political tweets (Johnson & Goldwasser, 2018).

Using the extracted content information, researchers can explore further questions such as whether competing politicians or political elites emphasize the same issues (Gabel & Scheve, 2007; Petrocik, 1996) and how the priorities politicians articulate co-vary with the issues discussed in the media (Bartels, 1996). Another open research direction is to analyse the interaction between newspapers and politicians' messages, such as how often newspapers cover a certain politician's message and in what way and how such coverage affects incumbency advantage.

Meaningful Future Work Apart from analysing the language of existing political texts that aims to maximize political interests, an advanced question that is more meaningful to society is how to improve policy communication to steer towards a more beneficial future for society as a whole. There is relatively little research on this, and we welcome future work on this meaningful topic.

7.3.4 *Investigating Policy Effects*

After policies are taken into effect, it is important to collect feedback or evaluate the effectiveness of policies. Existing studies evaluate the effects of policies along different dimensions: one dimension is the change in public sentiment, which can be analysed by comparing the sentiment classification results before and after policies, following a similar paradigm in Sect. 7.3.1. There are also studies on how policies affect the crowd's perception of the democratic process (Miller et al., 1990).

Another dimension is how policies result in economic changes. Calvo-González et al. (2018) investigate the negative consequences of policy volatility that harm long-term economic growth. Specifically, to measure policy volatility, they first obtain main topics by topic modelling on presidential speeches and then analyse how the significance of topics changes over time.

7.4 Limitations and Ethical Considerations

There are several limitations that researchers and policymakers need to take into consideration when using NLP for policymaking, due to the data-driven and black-box nature of modern NLP. First, the effectiveness of the computational models relies on the quality and comprehensiveness of the data. Although many political discourses are public, including data sources such as news, press releases, legislation, and campaigns, when it comes to surveying public opinions, social media might be a biased representation of the whole population. Therefore, when making important policy decisions, the traditional polls and surveys can provide more comprehensive coverage. Note that in the case of traditional polls, NLP can still be helpful in expediting the processing of survey answers.

The second concern is the black-box nature of modern NLP models. We do not encourage decision-making systems to depend fully on NLP, but suggest that NLP can assist human decision-makers. Hence, all the applications introduced in this chapter use NLP to compile information that is necessary for policymaking instead of directly suggesting a policy. Nonetheless, some of the models are hard to interpret or explain, such as text classification using deep learning models (Brown et al., 2020; Yin et al., 2019), which could be vulnerable to adversarial attacks by small paraphrasing of the text input (Jin et al., 2020). In practical applications, it is important to ensure the trustworthiness of the usage of AI. There could be a preference for transparent machine learning models if they can do the work well (e.g. LDA topic models and traditional classification methods using dictionaries or linguistic rules) or tasks with well-controlled outputs such as event extraction to select spans of the given text that mention events. In cases where only the deep learning models can provide good performance, there should be more detailed performance analysis (e.g. a study to check the correlation of the model decisions and human judgments), error analysis (e.g. different types of errors, failure modes, and potential bias towards certain groups), and studies about the interpretability of the model (e.g. feature attribution of the model, visualization of the internal states of the model).

Apart from the limitations of the technical methodology, there are also ethical considerations arising from the use of NLP. Among the use cases introduced in this chapter, some applications of NLP are relatively safe as they mainly involve analysing public political documents and fact-based evidence or effects of policies. However, others could be concerning and vulnerable to misuse. For example, although effective, truthful policy communication is beneficial for society, it might be tempting to overdo policy communication and by all means optimize the votes. As it is highly important for government and politicians to gain positive public perception, overly optimizing policy communication might lead to propaganda, intrusion of data privacy to collect more user preferences, and, in more severe cases, surveillance and violation of human rights. Hence, there is a strong need for policies to regulate the use of technologies that influence public opinions and pose a challenge to democracy.

7.5 Conclusions

This chapter provided a brief overview of current research directions in NLP that provide support for policymaking. We first introduced four main NLP tasks that are commonly used in text analysis: text classification, topic modelling, event extraction, and text scaling. We then showed how these methods can be used in policymaking for applications such as data collection for evidence-based policymaking, interpretation of political decisions, policy communication, and investigation of policy effects. We also discussed potential limitations and ethical considerations of which researchers and policymakers should be aware.

NLP holds significant promise for enabling data-driven policymaking. In addition to the tasks overviewed in this chapter, we foresee that other NLP applications, such as text summarization (e.g. to condense information from large documents), question answering (e.g. for reasoning about policies), and culturally adjusted machine translation (e.g. to facilitate international communications), will soon find use in policymaking. The field of NLP is quickly advancing, and close collaborations between NLP experts and public policy experts will be key to the successful use and deployment of NLP tools in public policy.

References

- Adler, E. Scott, & Wilkerson, J. (2011). Congressional bills project. *NSF 00880066 and 00880061*. <http://www.congressionalbills.org/>
- Albaugh, Q., Sevenans, J., Soroka, S., & Loewen, P. J. (2013). The automated coding of policy agendas: A dictionary-based approach. In *The 6th Annual Comparative Agendas Conference, Antwerp, Belgium*.
- Ansolabehere, S., & Iyengar, S. (1995). *Going negative: How political advertisements shrink and polarize the electorate* (Vol. 95). New York: Simon & Schuster.
- Armstrong, E. M., Carpenter, D. P., & Hojnacki, M. (2006). Whose deaths matter? Mortality, advocacy, and attention to disease in the mass media. *Journal of Health Politics, Policy and Law*, 31(4), 729–772.
- Arunachalam, R., & Sarkar, S. (2013). The new eye of government: citizen sentiment analysis in social media. In *Proceedings of the IJCNLP 2013 Workshop on Natural Language Processing for Social Media (SocialNLP)* (pp. 23–28). Nagoya, Japan: Asian Federation of Natural Language Processing. <https://www.aclweb.org/anthology/W13-4204>
- Azar, E. E. (1980). The conflict and peace data bank (COPDAB) project. *Journal of Conflict Resolution*, 24(1), 143–152.
- Baerg, N., Duell, D., & Lowe, W. (2018). Central bank communication as public opinion: experimental evidence. *Work in Progress*.
- Bakker, R., De Vries, C., Edwards, E., Hooghe, L., Jolly, S., Marks, G., Polk, J., Rovny, J., Steenbergen, M., & Vachudova, M. A. (2015). Measuring party positions in Europe: The Chapel Hill expert survey trend file, 1999–2010. *Party Politics*, 21(1), 143–152.
- Bartels, L. M. (1996). Politicians and the press: Who leads, who follows. In *Annual Meeting of the American Political Science Association* (pp. 1–60).
- Baumgartner, F. R., Green-Pedersen, C., & Jones, B. D. (2006). Comparative studies of policy agendas. *Journal of European Public Policy*, 13(7), 959–974.

- Bevan, S. (2017). Gone fishing: The creation of the comparative agendas project master codebook. *Comparative Policy Agendas: Theory, Tools, Data*. <http://sbevan.%20com/cap-master-codebook.html>
- Beverungen, G., & Kalita, J. (2011). Evaluating methods for summarizing Twitter posts. In *Proceedings of the 5th AAAI ICWSM*.
- Biran, O., Feder, O., Moatti, Y., Kiourtis, A., Kyriazis, D., Manias, G., Mavrogiorgou, A., Sgouros, N. M., Barata, M. T., Oldani, I., Sanguino, M. A. & Kranas, P. (2022). PolicyCLOUD: A prototype of a cloud serverless ecosystem for policy analytics. *CoRR*, abs/2201.06077. <https://arxiv.org/abs/2201.06077>
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2001). Latent dirichlet allocation. In T. G. Dietterich, S. Becker, & Z. Ghahramani (eds.), *Advances in Neural Information Processing Systems 14 [Neural Information Processing Systems: Natural and Synthetic, NIPS 2001, December 3–8, 2001, Vancouver, British Columbia, Canada]* (pp. 601–608). MIT Press. <https://proceedings.neurips.cc/paper/2001/hash/296472c9542ad4d4788d543508116cbc-Abstract.html>
- Blessing, A., Blokker, N., Haunss, S., Kuhn, J., Lapesa, G., & Padó, S. (2019). An environment for relational annotation of political debates. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations* (pp. 105–110). Florence, Italy: Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-3018>. <https://aclanthology.org/P19-3018>
- Bond, D., Jenkins, J. C., Taylor, C. L., & Schock, K. (1997). Mapping mass political conflict and civil society: issues and prospects for the automated development of event data. *Journal of Conflict Resolution*, 41(4), 553–579.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., . . . , Amodei, D. (2020). Language models are fewshot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, & H.-T. Lin (Eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, Neurips 2020, December 6–12, 2020, Virtual*. <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfbcb4967418bfb8ac142f64a-Abstract.html>
- Calvo-González, O., Eizmendi, A., & Reyes, G. J. (2018). Winners never quit, quitters never grow: Using text mining to measure policy volatility and its link with long-term growth in latin America. *World Bank Policy Research Working Paper* (8310).
- Camp, N. P., Voigt, R., Jurafsky, D., & Eberhardt, J. L. (2021). The thin blue waveform: racial disparities in officer prosody undermine institutional trust in the police. *Journal of Personality and Social Psychology*, 121, 1157–1171.
- Campbell, J. E. (1983). Ambiguity in the issue positions of presidential candidates: A causal analysis. *American Journal of Political Science*, 27, 284–293.
- Canes-Wrone, B., Brady, D. W., & Cogan, J. F. (2002). Out of step, out of office: Electoral accountability and house members' voting. *American Political Science Review*, 96, 127–140.
- Caughey, D., & Warshaw, C. (2018). Policy preferences and policy change: Dynamic responsiveness in the American states, 1936–2014. *American Political Science Review*, 112, 249–266.
- Chung, C. K., & Pennebaker, J. W. (2008). Revealing dimensions of thinking in OpenEnded self-descriptions: An automated meaning extraction method for natural language. *Journal of Research in Personality*, 42(1), 96–132.
- Collingwood, L., & Wilkerson, J. (2012). Tradeoffs in accuracy and efficiency in supervised learning methods. *Journal of Information Technology & Politics*, 9(3), 298–318.
- Cook, T. E. (1988). Press secretaries and media strategies in the house of representatives: Deciding whom to pursue. *American Journal of Political Science*, 32, 1047–1069.
- Crayton, A., Fonseca, J., Mehra, K., Ng, M., Ross, J., Sandoval-Castañeda, M., & von Gnechten, R. (2020). Narratives and needs: Analyzing experiences of cyclone amphan using Twitter discourse. *CoRR*, abs/2009.05560. <https://arxiv.org/abs/2009.05560>
- Davoodi, M., Waltenburg, E., & Goldwasser, D. (2020). Understanding the language of political agreement and disagreement in legislative texts. In *Proceedings of the 58th Annual*

- Meeting of the Association for Computational Linguistics* (pp. 5358–5368). Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.476>. <https://aclanthology.org/2020.acl-main.476>
- Demszky, D., Garg, N., Voigt, R., Zou, J., Shapiro, J., Gentzkow, M., & Jurafsky, D. (2019). Analyzing polarization in social media: method and application to tweets on 21 mass shootings. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, volume 1 (Long and Short Papers)* (pp. 2970–3005). Minneapolis, Minnesota: Association for Computational Linguistics. <https://doi.org/10.18653/v1/N19-1304>. <https://aclanthology.org/N19-1304>
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In J. Burstein, C. Doran, & T. Solorio (Eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2–7, 2019, Volume 1 (Long and Short Papers)* (pp. 4171–4186). Association for Computational Linguistics. <https://doi.org/10.18653/v1/n19-1423>
- Ding, H., & Riloff, E. (2018). Human needs categorization of affective events using labeled and unlabeled data. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long papers)* (pp. 1919–1929). New Orleans, Louisiana: Association for Computational Linguistics. <https://doi.org/10.18653/v1/N18-1174>. <https://aclanthology.org/N18-1174>
- Ding, H., Riloff, E., & Feng, Z. (2019). Improving human needs categorization of events with semantic classification. In *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (*SEM 2019)* (pp. 198–204). Minneapolis, Minnesota: Association for Computational Linguistics. <https://doi.org/10.18653/v1/S19-1022>. <https://aclanthology.org/S19-1022>
- Döring, H., & Regel, S. (2019). Party facts: A database of political parties worldwide. *Party Politics*, 25(2), 97–109.
- Egami, N., Fong, C. J., Grimmer, J., Roberts, M. E., & Stewart, B. M. (2018). How to make causal inferences using texts. *CoRR*, abs/1802.02163. <http://arxiv.org/abs/1802.02163>
- Eichorst, J., & Lin, N. C. N. (2019). Resist to commit: Concrete campaign statements and the need to clarify a partisan reputation. *The Journal of Politics*, 81(1), 15–32.
- Engel, U., Quan-Haase, A., Liu, S. X., & Lyberg, L. (2021). *Handbook of computational social science* (Vol. 2). Taylor & Francis.
- Engstrom, D. F., Ho, D. E., Sharkey, C. M., & Cuéllar, M. (2020). Government by algorithm: Artificial intelligence in federal administrative agencies. *NYU School of Law, Public Law Research Paper* (20–54).
- European Commission (2017). COM(2011) 808 Final: Horizon 2020 — the framework programme for research and innovation. In <https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%5C%3A52011PC0809> (15 May, 2022).
- Fenno, R. F. (1973). *Congressmen in committees*. In Boston: Little Brown & Company.
- Fenno, R. F. (1978). *Home style: House members in their districts*. Boston: Addison Wesley.
- Gabel, M., & Scheve, K. (2007). Estimating the effect of elite communications on public opinion using instrumental variables. *American Journal of Political Science*, 51(4), 1013–1028.
- Gennaro, G., & Ash, E. (2021). Emotion and reason in political language. *The Economic Journal*, 132(643), 1037–1059. <https://doi.org/10.1093/ej/ueab104>
- Gerner, D. J., Schrod, P. A., Yilmaz, O., & Abu-Jabr, R. (2002). Conflict and mediation event observations (cameo): A new event data framework for the analysis of foreign policy interactions. In *International Studies Association, New Orleans*.
- Gigley, H. M. (1993). Projected government needs in human language technology and the role of researchers in meeting them. In *Human Language Technology: Proceedings of a Workshop Held at Plainsboro, New Jersey, March 21–24, 1993*. <https://aclanthology.org/H93-1056>
- Glavaš, G., Nanni, F., & Ponzetto, S. P. (2017a). Cross-lingual classification of topics in political texts. In *Proceedings of the Second Workshop on NLP and Computational Social Science*

- (pp. 42–46). Vancouver, Canada: Association for Computational Linguistics. <https://doi.org/10.18653/v1/W17-2906>. <https://aclanthology.org/W17-2906>
- Glavaš, G., Nanni, F., & Ponzetto, S. P. (2017b). Unsupervised cross-lingual scaling of political texts. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers* (pp. 688–693). Valencia, Spain: Association for Computational Linguistics. <https://aclanthology.org/E17-2109>.
- Glavaš, G., Nanni, F., & Ponzetto, S. P. (2019). Computational analysis of political texts: bridging research efforts across communities. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts* (pp. 18–23). Florence, Italy: Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-4004>. <https://aclanthology.org/P19-4004>
- Goldstein, J. S. (1992). A conflict-cooperation scale for weis events data. *Journal of Conflict Resolution*, 36(2), 369–385.
- Gottipati, S., Qiu, M., Sim, Y., Jiang, J., & Smith, N. A. (2013). Learning topics and positions from Debatepedia. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing* (pp. 1858–1868). Seattle, Washington, USA: Association for Computational Linguistics. <https://aclanthology.org/D13-1191>
- Grimmer, J. (2010a). A Bayesian hierarchical topic model for political texts: Measuring expressed agendas in Senate press releases. *Political Analysis*, 18(1), 1–35.
- Grimmer, J. (2013). Appropriators not position takers: The distorting effects of electoral incentives on congressional representation. *American Journal of Political Science*, 57(3), 624–642.
- Grimmer, J., Messing, S., & Westwood, S. J. (2012). How words and money cultivate a personal vote: The effect of legislator credit claiming on constituent credit allocation. *American Political Science Review*, 106(4), 703–719.
- Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, 21(3), 267–297.
- Grimmer, J. R. (2010b). *Representational style: The central role of communication in representation*. Harvard University.
- Hausladen, C. I., Schubert, M. H., & Ash, E. (2020). Text classification of ideological direction in judicial opinions. *International Review of Law and Economics*, 62, 105903. <https://doi.org/10.1016/j.irl.2020.105903>. <https://www.sciencedirect.com/science/article/pii/S014818819303667>
- Hill, K. Q., & Hurley, P. A. (2002). Symbolic speeches in the us senate and their representational implications. *Journal of Politics*, 64(1), 219–231.
- Hillard, D., Purpura, S., & Wilkerson, J. (2008). Computer-assisted topic classification for mixed-methods social science research. *Journal of Information Technology & Politics*, 4(4), 31–46.
- Hirst, G., Riabinin, Y., & Graham, J. (2010). Party status as a confound in the automatic classification of political speech by ideology. In *Proceedings of the 10th International Conference on Statistical Analysis of Textual Data (JADT 2010)* (pp. 731–742)
- Hiware, K., Dutt, R., Sinha, S., Patro, S., Ghosh, K., & Ghosh, S. (2020). NARMADA: Need and available resource managing assistant for disasters and adversities. In *Proceedings of the Eighth International Workshop on Natural Language Processing for Social Media* (pp. 15–24). Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.socialnlp-1.3>. <https://aclanthology.org/2020.socialnlp-1.3>
- Hogenboom, F., Frasinicar, F., Kaymak, U., & de Jong, F. (2011). An overview of event extraction from text. In M. van Erp, W. R. van Hage, L. Hollink, A. Jameson, & R. Troncy (Eds.), *Proceedings of the Workshop on Detection, Representation, and Exploitation of Events in the Semantic Web (Derive 2011)*, Bonn, Germany, October 23, 2011 (CEUR Workshop Proceedings) (Vol. 77, pp. 948–57). CEUR-WS.org. http://ceur-ws.org/Vol-779/derive2011%5C_submission%5C_1.pdf
- Honnibal, M., Montani, I., Van Landeghem, S., & Boyd, A. (2020). spaCy: Industrial-strength Natural Language Processing in Python. <https://doi.org/10.5281/zenodo.1212303>
- Hopkins, D. J., & King, G. (2010). A method of automated nonparametric content analysis for social science. *American Journal of Political Science*, 54(1), 229–247.

- Huguet Cabot, P.-L., Dankers, V., Abadi, D., Fischer, A., & Shutova, E. (2020). The pragmatics behind politics: Modelling metaphor, framing and emotion in political discourse. In *Findings of the association for computational linguistics: emnlp 2020* (pp. 4479–4488). Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.findings-emnlp.402>. <https://aclanthology.org/2020.findings-emnlp.402>
- Immelman, A. (1993). The assessment of political personality: A psychodiagnostically relevant conceptualization and methodology. *Political Psychology*, 14, 725–741.
- Iyyer, M., Enns, P., Boyd-Graber, J., & Resnik, P. (2014). Political ideology detection using recursive neural networks. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 1113–1122). Baltimore, Maryland: Association for Computational Linguistics. <https://doi.org/10.3115/v1/P14-1105>. <https://aclanthology.org/P14-1105>
- Jin, D., Jin, Z., Zhou, J. T., & Szolovits, P. (2020). Is BERT really robust? A strong baseline for natural language attack on text classification and entailment. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, the Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, the Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7–12, 2020* (pp. 8018–8025). AAAI Press. <https://aaai.org/ojs/index.php/AAAI/article/view/6311>
- Jin, Z., Chauhan, G., Tse, B., Sachan, M. & Mihalcea, R. (2021a). How good is NLP? A sober look at NLP tasks through the lens of social impact. In *Findings of the association for computational linguistics: ACL-IJCNLP* (pp. 3099–3113). Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.findings-acl.273>
- Jin, Z., Peng, Z., Vaidhya, T., Schoelkopf, B., & Mihalcea, R. (2021b). Mining the cause of political decision-making from social media: A case study of COVID-19 policies across the US states. In *Findings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021*. Association for Computational Linguistics.
- Johnson, K., & Goldwasser, D. (2018). Classification of moral foundations in microblog political discourse. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 720–730). Melbourne, Australia: Association for Computational Linguistics. <https://doi.org/10.18653/v1/P18-1067>. <https://aclanthology.org/P18-1067>
- Jurafsky, D., & Martin, J. H. (2000). *Speech and language processing An introduction to natural language processing, computational linguistics, and speech recognition*. Prentice Hall Series in Artificial Intelligence. Prentice Hall.
- Karan, M., Šnajder, J., Širinić, D., & Glavaš, G. (2016). Analysis of policy agendas: Lessons learned from automatic topic classification of Croatian political texts. In *Proceedings of the 10th SIGHUM Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities* (pp. 12–21). Berlin, Germany: Association for Computational Linguistics. <https://doi.org/10.18653/v1/W16-2102>. <https://aclanthology.org/W16-2102>
- King, G., & Lowe, W. (2003). An automated information extraction tool for international conflict data with performance as good as human coders: A rare events evaluation design. *International Organization*, 57(3), 617–642.
- Lauscher, A., Fabo, P. R., Nanni, F., & Ponzetto, S. P. (2016). Entities as topic labels: Combining entity linking and labeled lda to improve topic interpretability and evaluability. *IJCoL. Italian Journal of Computational Linguistics*, 2(2–2), 67–87.
- Laver, M., Benoit, K., & Garry, J. (2003). Extracting policy positions from political texts using words as data. *American Political Science Review*, 97(2), 311–331.
- Lax, J. R., & Phillips, J. H. (2009). Gay rights in the states: Public opinion and policy responsiveness. *American Political Science Review*, 103(3), 367–386.
- Lax, J. R., & Phillips, J. H. (2012). The democratic deficit in the states. *American Journal of Political Science*, 56(1), 148–166.
- Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A.-L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., & Jebara, T. (2009). Computational social science. *Science*, 323(5915), 721–723.

- Lee, F. E. (2008). Dividers, not uniters: Presidential leadership and senate partisanship, 1981–2004. *The Journal of Politics*, 70(4), 914–928.
- Lipinski, D. (2009). *Congressional communication: Content and consequences*. University of Michigan Press.
- Liu, K., Chen, Y., Liu, J., Zuo, X., & Zhao, J. (2020). Extracting events and their relations from texts: A survey on recent research progress and challenges. *AI Open*, 1, 22–39. <https://doi.org/10.1016/j.aiopen.2021.02.004>. <https://www.sciencedirect.com/science/article/pii/S266665102100005X>
- Loria, S. (2018). TextBlob documentation. *Release 0.15.2*.
- Lowe, W., Benoit, K., Mikhaylov, S., & Laver, M. (2011). Scaling policy preferences from coded political texts. *Legislative Studies Quarterly*, 36(1), 123–155.
- Luo, Y., Card, D., & Jurafsky, D. (2020). Detecting stance in media on global warming. In *Findings of the association for computational linguistics: EMNLP 2020* (pp. 3296–3315). Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.findings-emnlp.296>. <https://aclanthology.org/2020.findings-emnlp.296>
- Luz, S. (2022). Computational linguistics and natural language processing. English. In F. Zanettin & C. Rundle (Eds.), *The Routledge handbook of translation and methodology*. United States: Routledge.
- Manning, C. D., Raghavan, P., & Schütze, H. (2008). *Introduction to information retrieval*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511809071>. <https://nlp.stanford.edu/IR-book/pdf/irbookprint.pdf>
- Marchetti-Bowick, M., & Chambers, N. (2012). Learning for microblogs with distant supervision: political forecasting with Twitter. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics* (pp. 603–612). Avignon, France: Association for Computational Linguistics. <https://www.aclweb.org/anthology/E12-1062>
- Maronikolakis, A., Villegas, D. S., Preotiuc-Pietro, D., & Aletras, N. (2020). Analyzing political parody in social media. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 4373–4384). Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.403>. <https://aclanthology.org/2020.acl-main.403>
- Mayhew, D. R. (2004). *Congress: The electoral connection*. Yale University Press.
- McClelland, C. A. (1976). *World event/interaction survey codebook*.
- McCombs, M., & Valenzuela, S. (2004). *Setting the agenda: Mass media and public opinion*. Wiley.
- Menini, S., Nanni, F., Ponzetto, S. P., & Tonelli, S. (2017). Topic-based agreement and disagreement in US electoral manifestos. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 2938–2944). Copenhagen, Denmark: Association for Computational Linguistics. <https://doi.org/10.18653/v1/D17-1318>. <https://aclanthology.org/D17-1318>
- Menini, S., & Tonelli, S. (2016). Agreement and disagreement: Comparison of points of view in the political domain. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers* (pp. 2461–2470). Osaka, Japan: The COLING 2016 Organizing Committee. <https://aclanthology.org/C16-1232>
- Merritt, R. L., Muncaster, R. G., & Zinnes, D. A. (1993). *International event-data developments: DDIR phase II*. University of Michigan Press.
- Miller, W. L., Clarke, H. D., Harrop, M., LeDuc, L., & Whiteley, P. F. (1990). *How voters change: The 1987 british election campaign in perspective*. Oxford University Press.
- Misuraca, G., van Noordt, C., & Boukli, A. (2020). The use of AI in public services: Results from a preliminary mapping across the EU. In Y. Charalabidis, M. A. Cunha, & D. Sarantis (Eds.), *ICEGOV 2020: 13th International Conference on Theory and Practice of Electronic Governance, Athens, Greece, 23–25 September, 2020* (pp. 90–99). ACM. <https://doi.org/10.1145/3428502.3428513>
- Mitamura, T., Liu, Z., & Hovy, E. H. (2017). Events detection, coreference and sequencing: what's next? Overview of the TAC KBP 2017 event track. In *Proceedings of*

- the 2017 Text Analysis Conference, TAC 2017, Gaithersburg, Maryland, USA, November 13–14, 2017. NIST. https://tac.nist.gov/publications/2017/additional.papers/TAC2017.KBP%5C_Event%5C_Nugget%5C_overview.proceedings.pdf
- Mohammad, S. M., Zhu, X., Kiritchenko, S., & Martin, J. D. (2015). Sentiment, emotion, purpose, and style in electoral tweets. *Information Processing & Management*, 51(4), 480–499. <https://doi.org/10.1016/j.ipm.2014.09.003>
- Nanni, F., Glavas, G., Ponzetto, S. P., & Stuckenschmidt, H. (2019). Political text scaling meets computational semantics. *CoRR*, abs/1904.06217. <http://arxiv.org/abs/1904.06217>
- Nanni, F., Ponzetto, S. P., & Dietz, L. (2017). Building entitycentric event collections. In *2017 ACM/IEEE Joint Conference on Digital Libraries, JCDL 2017, Toronto, ON, Canada, June 19–23, 2017* (pp. 199–208). IEEE Computer Society. <https://doi.org/10.1109/JCDL.2017.7991574>
- Neuendorf, K. A., & Kumar, A. (2015). Content analysis. *The International Encyclopedia of Political Communication*, 8, 1–10.
- O'Connor, B., Balasubramanyan, R., Routledge, B. R., & Smith, N. A. (2010). From tweets to polls: Linking text sentiment to public opinion time series. In W. W. Cohen & S. Gosling (Eds.), *Proceedings of the Fourth International Conference on Weblogs and Social Media, ICWSM 2010, Washington, DC, USA, May 23–26, 2010*. The AAAI Press. <http://www.aaai.org/ocs/index.php/ICWSM/ICWSM10/paper/view/1536>
- Osnabrügge, M., Ash, E., & Morelli, M. (2021). Cross-domain topic classification for political texts. *Political Analysis*, 1–22. <https://doi.org/10.1017/pan.2021.37>
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P. F., Leike, J., & Lowe, R. (2022). Training language models to follow instructions with human feedback. *CoRR*, abs/2203.02155. <https://doi.org/10.48550/arXiv.2203.02155>
- Padó, S., Blessing, A., Blokker, N., Dayanik, E., Haunss, S., & Kuhn, J. (2019). Who sides with whom? Towards computational construction of discourse networks for political debates. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 2841–2847). Florence, Italy: Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1273>. <https://aclanthology.org/P19-1273>
- Page, B. I. (1976). The theory of political ambiguity. *American Political Science Review*, 70(3), 742–752.
- Pak, A., & Paroubek, P. (2010). Twitter as a corpus for sentiment analysis and opinion mining. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*. Valletta, Malta: European Language Resources Association (ELRA). http://www.lrec-conf.org/proceedings/lrec2010/pdf/385_Paper.pdf
- Paltoglou, G., & Thelwall, M. (2012). Twitter, myspace, digg: Unsupervised sentiment analysis in social media. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(4), 66:1–66:19. <https://doi.org/10.1145/2337542.2337551>
- Pang, B., & Lee, L. (2007). Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval*, 2(1–2), 1–135. <https://doi.org/10.1561/15000000011>
- Paul, D., & Frank, A. (2019). Ranking and selecting multi-hop knowledge paths to better predict human needs. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (pp. 3671–3681). Minneapolis, Minnesota: Association for Computational Linguistics. <https://doi.org/10.18653/v1/N19-1368>. <https://aclanthology.org/N19-1368>
- Pennebaker, J. W., Francis, M. E., & Booth, R. J. (2001). Linguistic inquiry and word count: Liwc 2001. *Mahway: Lawrence Erlbaum Associates*, 71(2001). 2001.
- Petrocik, J. R. (1996). Issue ownership in presidential elections, with a 1980 case study. *American Journal of Political Science*, 40, 825–850.
- Pla, F., & Hurtado, L.-F. (2014). Political tendency identification in Twitter using sentiment analysis techniques. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers* (pp. 183–192). Dublin, Ireland: Dublin City

- University & Association for Computational Linguistics. <https://www.aclweb.org/anthology/C14-1019>
- Preoțiuc-Pietro, D., Liu, Y., Hopkins, D., & Ungar, L. (2017). Beyond binary labels: Political ideology prediction of Twitter users. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 729–740). Vancouver, Canada: Association for Computational Linguistics. <https://doi.org/10.18653/v1/P17-1068>. <https://aclanthology.org/P17-1068>
- Qi, P., Zhang, Y., Zhang, Y., Bolton, J., & Manning, C. D. (2020). Stanza: a Python natural language processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. <https://nlp.stanford.edu/pubs/qi2020stanza.pdf>
- Quinn, K. M., Monroe, B. L., Colaresi, M., Crespin, M. H. & Radev, D. R. (2006). An automated method of topic-coding legislative speech over time with application to the 105th–108th US Senate. In *Midwest political science association meeting* (pp. 1–61).
- Quinn, K. M., Monroe, B. L., Colaresi, M., Crespin, M. H., & Radev, D. R. (2010). How to analyze political attention with minimal assumptions and costs. *American Journal of Political Science*, 54(1), 209–228.
- Raleigh, C., Linke, A., Hegre, H., & Karlsen, J. (2010). Introducing ACLED-Armed conflict location and event data. *Journal of Peace Research*, 47(5), 651–660. <https://journals.sagepub.com/doi/10.1177/0022343310378914>
- Roberts, M. E., Stewart, B. M., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S. K., Albertson, B., & Rand, D. G. (2014). Structural topic models for open-ended survey responses. *American Journal of Political Science*, 58(4), 1064–1082.
- Rosenthal, S., Nakov, P., Kiritchenko, S., Mohammad, S., Ritter, A., & Stoyanov, V. (2015). Semeval-2015 task 10: Sentiment analysis in Twitter. In D. M. Cer, D. Jurgens, P. Nakov, & T. Zesch (Eds.), *Proceedings of the 9th International Workshop on Semantic Evaluation, semeval@naacl-hlt 2015, Denver, Colorado, USA, June 4–5, 2015* (pp. 451–463). The Association for Computer Linguistics. <https://doi.org/10.18653/v1/s15-2078>
- Sarol, M. J., Dinh, L., Rezapour, R., Chin, C.-L., Yang, P., & Diesner, J. (2020). An empirical methodology for detecting and prioritizing needs during crisis events. In *Findings of the association for computational linguistics: EMNLP 2020*, 4102–4107. Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.findings-emnlp.366>. <https://aclanthology.org/2020.findings-emnlp.366>
- Schrodt, P. A. (2000). Pattern recognition of international crises using Hidden Markov Models. In *Political Complexity: Nonlinear Models of Politics*, 296–328, University of Michigan Press.
- Schrodt, P. A., & Gerner, D. J. (1994). Validity assessment of a machinecoded event data set for the middle east, 1982–92. *American Journal of Political Science*, 38, 825–854.
- Schrodt, P. A., & Hall, B. (2006). Twenty years of the kansas event data system project. *The Political Methodologist*, 14(1), 2–8.
- Schumacher, G., Schoonvelde, M., Traber, D., Dahiya, T., & Vries, E. D. (2016). EUSpeech: A new dataset of EU elite speeches. In *Proceedings of the International Conference on the Advances in Computational Analysis of Political Text (Poltext 2016)* (pp. 75–80).
- Semetko, H. A., & Valkenburg, P. M. (2000). Framing European politics: A content analysis of press and television news. *Journal of Communication*, 50(2), 93–109.
- Shah, D. V., Cappella, J. N., & Neuman, W. R. (2015). Big data, digital media, and computational social science: possibilities and perils. *The ANNALS of the American Academy of Political and Social Science*, 659(1), 6–13.
- Sigelman, L., & Buell Jr., E. H. (2004). Avoidance or engagement? Issue convergence in us presidential campaigns, 1960–2000. *American Journal of Political Science*, 48(4), 650–661.
- Sim, Y., Acree, B. D. L., Gross, J. H., & Smith, N. A. (2013). Measuring ideological proportions in political speeches. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing* (pp. 91–101). Seattle, Washington, USA: Association for Computational Linguistics. <https://aclanthology.org/D13-1010>

- Simon, A. F. (2002). *The winning message: Candidate behavior, campaign discourse, and democracy*. Cambridge University Press.
- Slapin, J. B., & Proksch, S.-O. (2008). A scaling model for estimating time-series party positions from texts. *American Journal of Political Science*, 52(3), 705–722.
- Stravani, D., Kameswari, L., & Mamidi, R. (2021). Political discourse analysis: A case study of code mixing and code switching in political speeches. In *Proceedings of the Fifth Workshop on Computational Approaches to Linguistic Code-Switching* (pp. 1–5). Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.calcs-1.1>. <https://aclanthology.org/2021.calcs-1.1>
- Stefanov, P., Darwish, K., Atanasov, A., & Nakov, P. (2020). Predicting the topical stance and political leaning of media using tweets. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 527–537). Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.50>. <https://aclanthology.org/2020.acl-main.50>
- Stimson, J. A., MacKuen, M. B., & Erikson, R. S. (1995). Dynamic representation. *American Political Science Review*, 89, 543–565.
- Sulkin, T. (2005). *Issue politics in congress*. Cambridge University Press.
- Sun, C., Qiu, X., Xu, Y., & Huang, X. (2019). How to fine-tune BERT for text classification? In M. Sun, X. Huang, H. Ji, Z. Liu & Y. Liu (Eds.), *Chinese Computational Linguistics - 18th China National Conference, CCL 2019, Kunming, China, October 18–20, 2019, Proceedings*, Lecture Notes in Computer Science (Vol. 11856, pp. 194–206). Springer. https://doi.org/10.1007/978-3-030-32381-3%5C_16
- Sundberg, R., & Melander, E. (2013). Introducing the ucdp georeferenced event dataset. *Journal of Peace Research*, 50(4), 523–532.
- Thelwall, M., Buckley, K., & Paltoglou, G. (2011). Sentiment in Twitter events. *Journal of the American Society for Information Science and Technology*, 62(2), 406–418. <https://doi.org/10.1002/asi.21462>
- Thomas, M., Pang, B., & Lee, L. (2006). Get out the vote: Determining support or opposition from congressional floor-debate transcripts. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing* (327–335). Sydney, Australia: Association for Computational Linguistics. <https://aclanthology.org/W06-1639>
- Tjong Kim Sang, E., & Bos, J. (2012). Predicting the 2011 Dutch senate election results with Twitter. In *Proceedings of the Workshop on Semantic Analysis in Social Media* (pp. 53–60). Avignon, France: Association for Computational Linguistics. <https://www.aclweb.org/anthology/W12-0607>
- Trapp, R. (2006). *Programming for Peace: Computer-Aided Methods for International Conflict Resolution and Prevention* (Vol. 2). Springer Science & Business Media.
- Unanue, S., Li, X., Sharaf, M. A., Zhong, J., & Li, X. (2014). Predicting elections from social networks based on sub-event detection and sentiment analysis. In B. Benatallah, A. Bestavros, Y. Manolopoulos, A. Vakali & Y. Zhang (Eds.), *Web Information Systems Engineering - WISE 2014 - 15th International Conference, Thessaloniki, Greece, October 12–14, 2014, Proceedings, Part II*. Lecture Notes in Computer Science (Vol. 8787, pp. 1–16). Springer. https://doi.org/10.1007/978-3-319-11746-1%5C_1.
- Van Aggelen, A., Hollink, L., Kemman, M., Kleppe, M. & Beunders, H. (2017). The debates of the European Parliament as linked open data. *Semantic Web*, 8(2), 271–281.
- Van Roy, V., Rossetti, F., Perset, K., & Galindo-Romero, L. (2021). *AI watch - national strategies on artificial intelligence: A European perspective, 2021 edition*. Scientific Analysis or Review, Policy Assessment, Country report KJ-NA-30745-EN-N (online). Luxembourg (Luxembourg). <https://doi.org/10.2760/069178>(online)
- Walker, C., Strassel, S., Medero, J., & Maeda, K. (2006). Ace 2005 multilingual training corpus. *Linguistic Data Consortium, Philadelphia*, 57.
- Winter, D. G., & Stewart, A. J. (1977). Content analysis as a technique for assessing political leaders. In *A psychological examination of political leaders* (pp. 27–61).

- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Scao, T. L., Gugger, S., ..., Rush, A. (2020). Transformers: state-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations* (pp. 38–45). Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-demos.6>. <https://aclanthology.org/2020.emnlp-demos.6>
- Yin, W., Hay, J., & Roth, D. (2019). Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP/IJCNLP)* (pp. 3914–3923). Hong Kong, China: Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-1404>. <https://aclanthology.org/D19-1404>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

