Research Master Thesis

Decoding Populism: Analyzing Lexical Choice and Linguistic Simplicity in Tweets

Marije R. Brandsma

a thesis submitted in partial fulfillment of the requirements for the degree of

RMA Linguistics
(Human Language Technology)

Vrije Universiteit Amsterdam

Computational Lexicology and Terminology Lab
Department of Language and Communication
Faculty of Humanities

Supervised by: Prof. dr. A.S. Fokkens and dr. P.J.M. Sommerauer
2^{nd} reader: dr. H. D. van der Vliet

Submitted: June 30, 2023

Abstract

This thesis explores the effectiveness of lexical choice and linguistic simplicity in distinguishing between populist and non-populist party tweets through Natural Language Processing. It demonstrates the importance of lexical choice in distinguishing between populist and non-populist language, while in this thesis linguistic simplicity did not turn out to be a distinguishable language characteristic. Pronouns, quantifiers, and adverbs emerged as particularly influential function word categories. Populist parties seemed to prefer emotionally charged words as opposed to non-populist parties. However, the overlap between populist discourse and right-wing political ideologies highlights the complexity of language analysis and the need for a nuanced understanding of the political context in which these words are utilized. The findings contribute to a better understanding of the role of lexical choice and linguistic simplicity in differentiating between populist and non-populist speech, shedding light on the complexities and nuances of political language.

Declaration of Authorship

I, Marije Roelofje Brandsma, declare that this thesis, titled *Decoding Populism: Analyzing Lexical Choice and Linguistic Simplicity in Tweets* and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a degree degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Date: June 30th, 2023

Signed:

Acknowledgments

I would like to express my gratitude to the following people who have played a significant role in the completion of my thesis:

First and foremost, I would like to thank my classmates for going on our academic journey together. Their companionship, humor, shared knowledge (and drinks) have made these past two years a lot more fun than it would have been on my own.

I am also grateful to Hasan Shahoud for the discussions we had regarding our theses and methodologies. Our exchange of ideas and perspectives has helped me in writing my thesis.

My appreciation also goes to my family, particularly my husband and mother, who read through all of my writings and provided valuable feedback, despite not being interested in the topic in the slightest. Thank you for all the feedback, and I promise to stop talking about this topic now.

I would like to thank the CLTL team at the VU for their guidance and teaching over the past few years. I am particularly grateful to the teachers who were present at my presentation and offered some very useful last-minute advice regarding my research.

Finally, special thanks are due to my supervisors, Antske and Pia. Their expertise, constant availability, and willingness to meet with me have been essential in finishing my thesis as it is now. I learned a lot from them. Especially the in-depth feedback from Pia has been incredibly helpful, not only to this thesis but to my academic writing and thinking in general.

List of Figures

List of Tables

Contents

Chapter 1

Introduction

In recent years, populist parties have gained traction worldwide (Algan et al., 2017). Prominent figures such as Donald Trump, Marine Le Pen, Geert Wilders, Nigel Farage, Matteo Salvini, but also left-wing politicians such as Bernie Sanders and Hugo Chávez have successfully attracted supporters with their distinctive "us" versus "them" narrative. Populist parties typically emphasize a dichotomy between two societal groups: "the pure people" versus "the corrupt elite" (Mudde, 2004, p. 543). Their communication style often involves emotive language and conveys a collective identity of "we" in opposition to established powers.

Specific rhetorical characteristics can assist in categorizing politicians and political parties as populist. Scholars have interpreted these common characteristics as either a strategic approach (Weyland, 2001) or a distinctive style (Moffitt and Tormey, 2014). Alternatively, some researchers perceive populism as an ideology (Mudde, 2007), while others find it most aptly described as a discourse (Hawkins, 2009). In this thesis, I explore populism as a style.

This thesis aims to contribute to the understanding of the language characteristics of the populist style. This is done by examining two stylometric characteristics of populism: lexical choice and linguistic simplicity. Furthermore, it explores different political categorizations such as left- and right-wing ideologies. Through this, this thesis seeks to avoid treating populism as an ideology exclusively associated with the right wing, recognizing its presence across the political spectrum (albeit predominantly observed on the left and right fringes). Through an analysis of textual features, the nature of the populist style is explored.

1.1 Researching populist language through NLP

Political ideology is known to influence language, including populist language. Previous research in the field of Natural Language Processing (NLP) has primarily focused on automatically classifying political orientation, rather than delving into the underlying language phenomena that drive such classification. Using NLP helps to understand these linguistic aspects and contributes to a deeper comprehension of the current political landscape and the growing appeal of populist parties and their rhetoric among voters.

Research about populist style has been done before, but not often using NLP methods. Usually, it consists of a political science comparison of two different politicians, resulting in an analysis of personal style and not a general conclusion on populist language. One such paper where NLP has been used to inspect populist language is Hirzalla (2019). However, their research was limited in scope, focusing on three right-wing parties in the Netherlands, two of which were not represented in the House of Representatives or the Senate. By focusing solely on these specific populist parties, a comparison between populist and non-populist party language was not conducted. Because of this, this study cannot conclusively state that the features it identified are specific to populist discourse.

This thesis expands on existing research by examining a broader range of political parties and conducting a comprehensive analysis that compares populist and non-populist party language. Moreover, this research does not isolate right-wing populism but explores the phenomenon of left-wing populism as well. By adopting this inclusive approach, a more comprehensive understanding of the language characteristics associated with populism in Dutch political discourse is achieved.

1.2 Goal and research questions

This thesis employs NLP techniques to investigate the characteristics of the Dutch populist language and how it differs from the non-populist language. This analysis is conducted by examining tweets from various political parties. Two primary stylistic distinctions are examined: lexical choice and linguistic complexity. By undertaking this study, the thesis makes the following contributions to the field:

- Establishment of a comprehensive dataset comprising tweets published by Dutch political parties and politicians.
- Thorough descriptions of the role of lexical choice and lexical complexity in populist and non-populist text, and their resulting role in the populist style.

Through these contributions, this thesis provides an enhanced understanding of how populist parties communicate their messages. This understanding sheds light on the strategies and tactics employed to appeal to specific segments of the population. Such insights help researchers, policymakers, and the general public better comprehend the nature and appeal of populism in political discourse. By examining both left-wing and right-wing populism, this thesis contributes to a more nuanced understanding of populism beyond its association with a particular political ideology. Understanding the similarities and differences in language characteristics between different ideological variants of populism and non-populism provides valuable insights into the underlying drivers and dynamics of populist discourse across the political spectrum.

This thesis reveals that lexical choice demonstrates strong predictive capability in distinguishing populist discourse, while linguistic complexity does not. Moreover, it highlights the challenges associated with exclusively attributing certain language features to populism, as they may overlap with other political party characteristics.

1.3 Outline

Chapter 2 delves into the theoretical background and previous research, examining the influence of political ideology on language and discussing existing studies on the language characteristics of populism and its classification in NLP. Chapter 3 outlines the methodology employed, including feature selection and extraction. Chapter 4 is a summary of the data collection process and gives an overview of the used data.

Chapter 5 presents the results of the analysis, showcasing the performance of the model and the effectiveness of the selected features in classifying populist and non-populist tweets.

Chapter 6 provides a comprehensive conclusion, summarizing the key findings and their implications. Finally, Chapter 7 concludes the paper, highlighting the significance of the study, contributions and observations, and potential avenues for future research.

Chapter 2

Theoretical background

2.1 Populism in political discourse

Populism is not new. It is a phenomenon with historical roots that can be traced back to different periods and contexts, as described by Rodriguez-Pose (2022). Populism has manifested in various forms, often emerging in response to societal, economic, and political crises that grip societies. According to Rodriguez-Pose, populism mostly seems to attract people who have known better times, but are now experiencing difficulties. Throughout history, populism has garnered attention for its ability to mobilize and appeal to segments of the population dissatisfied with the status quo.

While populism has a long-standing presence, the Brexit vote in the United Kingdom and the election of Donald Trump as the President of the United States, both in 2016, have meaningfully contributed to its surge (Rodriguez-Pose, 2022). The Brexit vote reflected deep-seated concerns about national identity, sovereignty, and immigration, while Trump's election campaign tapped into populist rhetoric by challenging the political establishment and promising to prioritize the interests of the ordinary citizen. These events have served as powerful catalysts, showing populist parties to be potential competitors of the current establishment. They contributed to propelling populist sentiments and movements to the forefront of political discourse.

In the contemporary political landscape, anti-establishment parties can be found all over the developed world. These parties position themselves as representing the working class person and immediately countering the elite, established power. Mudde describes the ideology of populist parties through their distinction between two groups in society: 'the pure people' versus 'the corrupt elite' (Mudde, 2004, p. 543).

The language and discourse employed by populist parties often exhibit distinct rhetorical characteristics, drawing upon popular grievances and presenting themselves as the voice of the disenfranchised which, according to the populist parties, all working-class, non-elite civilians are. Szilagy (2017) and Engesser et al. (2017) summarize the characteristics as follows: they praise 'the people', use personal pronouns to associate themselves with this group, call out the elites and the established power, use metonymies, and highlight the victimhood of their target audience. These points often correlate heavily with lexical choice.

2.2 Populist lexical choice

The characteristics proposed by Szilagy (2017) and supported by Engesser et al. (2017) mostly revolve around lexical choice, specifically the words employed by populist parties to communicate their message.

One of the fundamental features proposed is the division between 'the pure people' and 'the corrupt elite', as mentioned before. Populist rhetoric frequently highlights the distinction between these two groups and positions themselves as advocating for and being a part of 'the people'. This core message is exemplified in a tweet posted by the international account of Forum voor Democratie (*Forum for Democracy*, FvD) on June 19th, 2021: "We are facing an unprecedented restriction of individual liberties. A power grab by globalist elites who aim to radically change our world, our democracies, and our Western way of life."[1] The individual is placed opposite of the elite, who consciously attempt to remove their liberties and change their way of life. This is reflected in words such as "globalist elites".

Furthermore, populists often employ pronouns to emphasize a sense of 'us' versus 'them'. Populist text often refers to the elite as 'you', 'them', and 'they', while identifying themselves and their potential voters ('the people') as 'we' and 'I'. In the aforementioned example from FvD, this distinction is evident in the use of 'we' and 'our'.

Another characteristic of populist rhetoric is the tendency to launch frequent attacks on the existing elite, whether it be individuals in positions of power, wealthy civilians, or those deemed privileged in a different way. A tweet by Rashida Tlaib, a member of the American left-wing populist party Justice Democrats, serves as an illustration of such attacks. In her tweet, Tlaib accuses oil companies ('big oil') of 'price gouging' and 'exploitation' of the American people, highlighting the narrative of corporations taking advantage of ordinary citizens: "Last year, Big Oil raked in record-breaking profits while price gouging our families. (...) That's exploitation."[2]

Populist discourse also makes use of metonymies or alternative words to describe certain entities. For instance, 'Brussels' may be used to refer to the European Union, and in the case of Rashida Tlaib's tweet, the term 'Big Oil' is used to refer to oil companies.

Finally, a sense of victimhood of 'the people' is a prominent element emphasized in populist rhetoric. The manifestation of victimhood is less straightforward to capture through lexical choice alone. Both the tweets by Rashida Tlaib and FvD convey a sense of victimhood. Phrases such as 'taking advantage of' and 'unprecedented restriction' indicate this.

Limited research has been conducted on the impact of lexical choice for classifying populist and non-populist texts. However, studies have explored the role of lexical choice in distinguishing between different political ideologies, such as left-wing and right-wing political language. Researchers such as Lin et al. (2006) have found that a speaker's lexical choice can provide insights into their political ideas. They used lexical choice as a key feature for distinguishing political ideologies, achieving notable success in their models. Their models could correctly predict ideology with an accuracy of 70%, and later even between 80.1% and 86% depending on the party (Yu et al., 2008). Due to this, lexical choice was assumed to be a good predictor of political ideology.

---

[1] `https://twitter.com/FVD_Intl/status/1406230859015933953`

[2] `https://twitter.com/RepRashida/status/1625857244285005825`

However, Hirst et al. (2010) discovered that party status rather than ideology played a decisive role in classification models. Instead of recognizing a party's political polarity, it measured whether the party was in the current government or in the opposition. The lexical choice of parties almost completely swapped when analyzing a time when the other polarity of ideology was in power.

These findings highlight the importance of understanding what a model is truly distinguishing. Exploring alternative characteristics that the model could be recognizing instead is crucial not only for studying left- and right-wing ideologies, but for differentiating between other political distinctions, such as populism and non-populism, as well. By considering multiple factors and characteristics, it is possible to develop more accurate and nuanced models and to interpret them relative to each other.

2.3 Linguistic simplicity

Besides the lexical choices of populist parties, linguistic simplicity has been recognized as an important characteristic of populist speech for over two decades. Researchers examining language complexity in populism focus on factors such as the difficulty of words, lexicon, and sentence length to determine the level of reading simplicity, for example through the Flesch-Kincaid test for English (Kincaid et al., 1975). It is argued that populist parties prefer simple language to distinguish themselves from the more formal language associated with 'elite' politicians (Block and Negrine, 2017; Canovan, 1999; Rooduijn, 2014).

For example, Viser (2015) conducted research on the announcement speeches of United States presidential candidates Donald Trump and Hillary Clinton in 2015. They found that Donald Trump, widely considered a populist politician, had a Flesch-Kincaid score of 4.1, corresponding to a language level that a fourth-grader could understand. In contrast, Hillary Clinton had a score of 7.1, reflecting a language level appropriate for a seven-grader. This finding was supported by Kayam (2018), who analyzed debates and interviews with Trump.

However, linguistic simplicity as an important characteristic of populism has not been without criticism. McDonnell and Ondelli (2020) examined four different populist politicians in four different countries and incorporated various language simplicity tests, such as the Dale-Chall test (Chall and Dale, 1995) and the Flesch-Kincaid test (Kincaid et al., 1975). Their findings revealed some surprising results. Contrary to previous claims, they did not consistently find that populist language was less complex than the language of their opponents. In fact, they sometimes found the opposite effect, with the populist language being more complex.

While linguistic simplicity is an interesting direction and highlights one of the tactics assumed to make populist rhetoric appealing, McDonnell and Ondelli (2020) argue for a greater focus on vocabulary choice and rhetorical devices, particularly concerning the ideas, concepts, and frames that commonly occur in populism, such as 'the people' and 'the elites'. This thesis will do both: explore to what extent linguistic simplicity plays a role in distinguishing populist from non-populist text while at the same time analyzing lexical choice.

2.4 Introduction to Dutch politics

Since this thesis focuses on Dutch populist parties, it is important to provide an introduction to the Dutch political system and its populist parties.

The Netherlands is a parliamentary democracy (for an overview, see House of Representatives (2011)). In the Dutch political system, citizens participate in this democracy by voting for different parties to represent them in the House of Representatives (*Tweede Kamer der Staten-Generaal*) and the Senate (*Eerste Kamer der Staten-Generaal*). Elections are held every four years. The political landscape in the Netherlands is dynamic, with new parties frequently emerging, and the composition of parties in the House and Senate changing after (almost) every election.

The Dutch political parties can be divided into two groups: those that are part of the coalition, and those in the opposition. Parties belonging to the coalition are responsible for forming the government and delivering ministers for the Cabinet, while opposition parties monitor the actions of the government. It is worth noting that coalition parties sometimes vote in alignment with each other to support the government, even if it contradicts their own ideological positions.

The Dutch election system operates on a representative basis rather than allocating seats directly to parties. This means that individual representatives hold seats, and they have the flexibility to distance themselves from their party if they choose, even though this is generally frowned upon. If this does happen, a representative can either give up their seat or continue as an independent politician. Sometimes, multiple representatives may split from their party simultaneously. An example of this occurred in 2021 inside the FvD party, resulting in the creation of the Groep van Haga (*Van Haga Group*).

Currently, the House of Representatives in the Netherlands comprises twenty parties, including four officially independent individuals or groups. The coalition consists of four parties: VVD (*People's Party for Freedom and Democracy*), D66 (*Democrats 1966*), CDA (*Christian Democratic Appeal*), and ChristenUnie (*Christian Union*). The remaining sixteen parties make up the opposition.

According to an expert survey conducted by (Meijers and Zaslove, 2018) (published in (Meijers and Zaslove, 2021)), parties were placed along a spectrum of populism. Four Dutch parties were identified as populist: PVV (*Party of Freedom*), FvD, SP *(Socialist Party*), and 50Plus (*50 plus*). However, the political landscape has since evolved, particularly due to the COVID-19 crisis. FvD, in particular, has undergone changes and has been associated with anti-elite and anti-globalist positions, sometimes accused of promoting conspiracy theories (Ottens, 2020; FvD, 2023a). Additionally, a new populist party, BBB (*Farmer Citizen Party*), has emerged. Following the 2021 elections, 50Plus was left with only one seat, and the representative on that seat left the party to continue independently. As a result, 50Plus currently has minimal influence in Dutch politics.

According to Meijers and Zaslove, the least populist parties at the time of the survey were D66, VVD, ChristenUnie, GroenLinks (*Green Left*), SGP (*Reformed Political Party*), and PvdA (*Labour Party*).

According to (Tunderman, 2022), the Netherlands provides fertile ground for populist parties due to various socio-economic factors. The study highlights increasing financial inequalities, deteriorating working conditions, and soaring housing prices as important issues that contribute to the appeal of populist parties in the Netherlands. In this context, three parties are identified by Tunderman as relevant Dutch populist parties: the PVV, FvD, and the SP.

The PVV, led by Geert Wilders, has gained prominence with its anti-Islam stance, emphasizing concerns about immigration, national identity, and security PVV (2021). It has capitalized on public discontent regarding cultural integration and has been critical of the European Union. It is the largest opposition party and the third largest party in the Netherlands, with 11% of the seats in parliament.

The FvD, under the leadership of Thierry Baudet, emerged as a populist party challenging the political establishment FvD (2023b). Initially focused on Euroscepticism and nationalist themes, the party later experienced internal divisions and shifts in its ideological positioning, leaving a set of core members that are often regarded as being on the far-right side of the political spectrum. They are outspoken believers in *The Great Reset* as a movement against the average citizen, are very critical of the way the government dealt with the Covid pandemic, and speak out against immigration, climate change 'hysteria', and corrupt parties in the coalition.

The SP, a left-wing populist party, has focused on social and economic issues, advocating for improved working conditions, social equality, and increased social welfare. It has positioned itself as a defender of the working class and has criticized neoliberal economic policies SP (2023).

These populist parties have managed to attract support by tapping into public concerns and grievances related to economic inequality, immigration, national identity, and dissatisfaction with mainstream politics. The specific issues they prioritize and their policy proposals may differ, but they share a common emphasis on challenging the established political order and appealing to segments of the population that feel left behind or unheard, as (Tunderman, 2022) describes.

Chapter 3

Method

This thesis builds on the assumption that language carries meaningful information that can differentiate between populist and non-populist language, as discussed in chapter 1 and 2. A classification model was trained to distinguish between these two groups based on the features lexical choice and linguistic simplicity. The goal is to explore the relevance of these features in distinguishing populist and non-populist tweets.

The features related to lexical choice were further divided into the subfeatures *topic words* and *function words* (section 3.2.1). Similarly, linguistic simplicity was divided into the following subfeatures: *average number of words per sentence*, the Dutch readability formula *Leesindex A*, the *Flesch readability formula*, and the adjusted version for Dutch known as *Flesch-Douma* (section 3.2.2).

In order to train and evaluate the classification models, a series of experiments were conducted to assess the impact of different features on the model's performance. These experiments involved analyzing the changes in performance by adding or removing specific features, through addition and ablation studies.[1]

By conducting these analyses and examining the impact of different subfeatures within lexical choice and linguistic simplicity, the aim was to identify the most informative aspects of language that contribute to distinguishing between populist and non-populist language.

This chapter is divided into three sections. First, the manner of classification is described in section 3.1. Section 3.2 discusses the features that were used to train the models. Finally, the control experiments are explained in section 3.3. These experiments aim to verify whether the feature relevance found in the populist versus non-populist classification is specific to this particular distinction.

3.1 Classification

To accomplish the task of recognizing a two-class split in the data, such as populist or non-populist, coalition or opposition, left-wing or right-wing, and progressive or conservative, a binary Support Vector Machine (SVM) was chosen as the classification algorithm.

Binary SVM Support Vector Machines are widely used for binary classification tasks due to their effectiveness in separating data points into distinct classes. The goal of a

[1] All code and results can be found on `https://github.com/mrbrandsma/master-thesis-populism`.

binary SVM is to find an optimal hyperplane that separates the data into two different classes. The model learns to do this during the training phase, where a hyperplane is found that minimizes classification errors. This hyperplane is then applied to new, unseen data.

In this case, the binary SVM model was trained to classify text samples into the categories of interest: populist or non-populist language, coalition or opposition affiliation, left-wing or right-wing ideology, and progressive or conservative stance. By training the SVM on a labeled dataset with language features such as lexical choice, it learns the patterns and characteristics associated with each class and can then classify unseen texts based on those patterns.

Feature performance Central to this thesis is the question of whether certain features are helpful in distinguishing populist versus non-populist language. To analyze this, the precision, recall, and F1-score of a model are calculated and inspected. If a feature is useful, the model will perform above chance. If it is not, it is expected that the model has a bad performance, not being able to outperform the majority class baseline. This would suggest that the model did not find the features it was trained on helpful and has instead learned to rely on the statistical likelihood of the majority class. Generally, this will be reflected in a high recall score for the majority class, while reflecting a low recall for the minority class.

As can be seen in chapter 4, the model has been trained on a dataset where the non-populist parties make up the majority class. This may lead to predicting the non-populist class more frequently during classification, especially if the tested feature is not helpful.

3.2 Features

3.2.1 Lexical features

To examine the word usage in tweets, Term Frequency-Inverse Document Frequency (TF-IDF) has been used as the main feature to train a binary SVM classifier on lexical choice. TF-IDF is a numerical representation of the importance of a term. It is widely used in NLP to measure the significance of words or terms in a collection of texts. TF-IDF measures how often a term appears in a text relative to the total number of terms in that text (Term Frequency), and reflects the rarity or uniqueness of a term in the entire dataset (Inverse Document Frequency). Because of this, it is able to reduce the impact of common words and can account for text length.

Besides returning a performance score of the model when trained on the TF-IDF representation, this classification method can also provide insight into the specific words that contribute to the classification model by ranking the terms based on their TF-IDF scores. The 10 terms with the highest positive TF-IDF score and the 10 terms with the lowest negative TF-IDF score were taken. These correspond to the two classes that were classified in the model.

These 20 terms were then presented to 29 human participants who were asked which terms they associated with different groups of political parties. The goal of this survey was to check whether the words that were relevant to the model are intuitively relevant to the average person as well. If these terms are specific to populist language, it is expected that people should have an intuitive understanding of this. They had four options: D66, PvdA, GroenLinks and ChristenUnie (left non-populist), SP (left populist), VVD and SGP (right non-populist), and PVV and FvD (right populist). Participants could choose multiple party groups, with a minimum of one. The participants were not informed of the reason for these party groupings, nor did they know the reason behind the choice of words or the goal of the research, to avoid influencing their categorizations. Because of this, it is unlikely that they associated the words with whether a party is considered populist or not.

Subfeature: Topic words Topic words were both explored as a subfeature. Topic words are words that are indicative of specific topics or themes. These topic words were chosen by investigating the 100 terms that most impacted the model and removing the topic words in the model training procedure. This process was repeated until there were no topic words in the most frequent terms list anymore. Topic words that were removed generally had to do with climate change, migration, economy, and education. A list can be found in Appendix A. The role of topic words was explored through an ablation and addition study by comparing models with topic words and without topic words.

Subfeature: Function words Function words were also removed for some of the models to inspect their impact. These are words that serve grammatical or structural purposes in a sentence, as opposed to content words. The following types of function words were inspected: articles, prepositions, quantifiers, conjunctions, pronouns, auxiliary verbs, adverbs, modifiers, and interjections. The impact of these different types of function words was studied through an ablation and addition study.

Pronouns especially were of interest, as these belong to the proposed strategy of populist party language (see section 2.2).

Word removal In the process of selecting words for removal, the main criterion was to exclude words that were strongly associated with specific political parties but were not relevant to the research goal. These words could be indicative of certain parties due to their frequent usage by those parties (such as FvD referring to their online program Forum Inside), or the presence of prominent politician's names. Since the objective was to distinguish between populist and non-populist language rather than identifying specific party affiliations, it was necessary to remove these party-specific words from consideration. These were not treated as a feature, and thus were removed from all models.

Politician's names were removed, but terms such as 'hi' and 'hello' were removed as well, as it was indicative of one specific party's communication on Twitter (D66). Words that were frequently used by one specific party, such as the word 'inside' in the example before, were removed as well. A full list can be found in Appendix A, including the reasoning behind the removal of those words.

3.2.2 Linguistic simplicity

To measure the linguistic simplicity of the tweets, features that showed to be effective in classifying linguistic simplicity in Vandeghinste and Bulté (2019) have been used. This research identified average words per sentence (w/s), Leesindex A, Flesch, and Flesch-Douma as a few of the most informative simplicity features for Dutch.

Leesindex A is an older readability formula that provides a score that indicates the ease of understanding a text (Brouwer, 1963). A higher score suggests a more straightforward and accessible text. The formula is as follows:

$$A = 195 - \left(2 \times \frac{\text{total words}}{\text{total sentences}}\right) - \left(66.7 \times \frac{\text{total syllables}}{\text{total words}}\right)$$

The Flesch formula, introduced by Flesch (1948), is applicable to both English and Dutch (Vandeghinste and Bulté, 2019). This formula has been used in research regarding populist language before and found to be effective. A higher Flesch index score corresponds to a more readable text. The Flesch formula is expressed as:

$$I = 206.835 - \left(1.015 \times \frac{\text{total words}}{\text{total sentences}}\right) - \left(84.6 \times \frac{\text{total syllables}}{\text{total words}}\right)$$

The Flesch-Douma formula is an adaptation of the Flesch formula for Dutch (Douma, 1960).

$$I = 206.835 - \left(0.93 \times \frac{\text{total words}}{\text{total sentences}}\right) - \left(77 \times \frac{\text{total syllables}}{\text{total words}}\right)$$

For each tweet, these scores are calculated, and the resulting scores are provided as features of the model. In addition to the scores of the model, the averages and standard deviations of the different reading features are calculated as well. These are based on the scores of the training dataset.

3.3 Control experiments

In addition to classifying populist and non-populist parties, this thesis investigates alternative political categorizations. As described in chapter 2, it is possible that a model measures a different kind of political distinction instead of populism. To check for this, the analysis explores the party's role (opposition/coalition), as well as two other measurements of political ideology: left/right and progressive/conservative. All these classifications have been taken from Meijers and Zaslove (2018). The characteristics of parties based on these distinctions are presented in Table 3.1.

As can be seen, there are no coalition parties that are considered populist. The absence of populist parties in the coalition implies that any potential linguistic distinctions related to party role and coalition participation cannot be directly tested using the available data. When considering political ideology, both populist and non-populist parties can be found across the categories.

The same analysis used to classify populist and non-populist tweets was applied to analyze the other categories. By employing this approach, we can determine whether the features used in the analysis exclusively predict populist or non-populist categorizations or if they might be building on other differences, such as the role of the party within parliament.

Party	Role	Left/right	Prog/cons
Populist			
- PVV	Opposition	Right	Conservative
- SP	Opposition	Left	Progressive
- FvD	Opposition	Right	Conservative
Non-populist			
- VVD	Coalition	Right	Conservative
- D66	Coalition	Left	Progressive
- PvdA	Opposition	Left	Progressive
- GroenLinks	Opposition	Left	Progressive
- ChristenUnie	Coalition	Left	Conservative
- SGP	Opposition	Right	Conservative

Table 3.1: Characteristics of parties.

Chapter 4

Data

4.1 Data collection

For this thesis, tweets spanning from October 26, 2017, to March 29, 2023 have been collected. The chosen time frame corresponds to the initiation of cabinet Rutte III and extends into the cabinet period of Rutte IV, during which the parties VVD, D66, CDA, and ChristenUnie have engaged in collaborative governance. It was chosen to only use tweets from the most recent 6 years because time is known to play a role in the language used in politics (Németh, 2023). Because of this, the role division between the political parties discussed in this thesis has remained unchanged throughout the collection of tweets. As a result, the analysis covers a six-year tweet period characterized by the consistent roles played by the respective parties. This permits annotating all parties the same throughout the dataset when it comes to their role in parliament. It is important to note that none of the populist parties in the dataset have held a position in the coalition in the past.[1]

The selection and classification of populist and non-populist parties are based on Meijers and Zaslove's survey 2018. Parties that were closer to an average score and thus harder to definitively classify or were not included in the survey have been left out. Consequently, the following populist parties (in random order) have been chosen: Socialistische Partij (*Socialist Party*, SP), Forum voor Democratie (*Forum for Democracy*, FvD), and Partij voor de Vrijheid (*Party for Freedom*, PVV). This corresponds to the identification of populist parties by Tunderman (2022). Likewise, the following non-populist parties (in random order) have been selected: Democraten '66 (*Democrats '66*, D66), Staatkundig Gereformeerde Partij (*Reformed Political Party*, SGP), Volkspartij voor Vrijheid en Democratie (*People's Party for Freedom and Democracy*, VVD), Partij van de Arbeid (*Labour Party*, PvdA), GroenLinks (*Green Left*) and ChristenUnie (*Christian Union*).

In general, all tweets posted by the selected parties since the start date have been included in the analysis. However, for D66, a cut-off point of 5600 tweets was applied due to the exceptionally large volume of tweets by this party during the specified time period. This threshold of 5600 was set as it represents the number of tweets posted by the PVV, with the second-largest number of tweets posted by a party. The resulting distribution of tweets after undergoing preprocessing (further details of which will be provided in the corresponding paragraph) can be found in Table 4.2.

[1]The PVV has formally supported the coalition in cabinet Rutte I (2010-2012) but was never an official part of it and did not deliver ministers for the Cabinet.

Party	Authors	Avg characters p.t.	Avg words p.t.
PVV	1	158	24
SP	+	183	30
FvD	Unknown	220	35
VVD	Unknown	135	22
D66	+	181	29
PvdA	9	167	27
GroenLinks	1+	181	29
ChristenUnie	+	155	25
SGP	+	177	29

Table 4.1: Information on the twitter data per party (p.t. = per tweet). A '+' indicates a team of writers consisting of an unknown amount, while a '1+' indicates one writer, but with input and feedback from (multiple) other people. All datasets (train, dev, test) are used to calculate the averages.

4.2 Authors

Given the focus on author-specific features in this thesis, the number of authors per tweet is important to consider. This can explain whether style in a tweet is the result of one writer, or a group of writers and hence a better representation of conscious political style. Table 4.1 provides an overview of the number of writers. It can be observed that, for the parties that have provided information on their tweeting process, multiple writers are typically involved. PVV is the exception, however. Unlike other parties, the PVV does not own a Twitter account. Instead, their party leader, Geert Wilders, uses his personal account for tweeting on behalf of the party. Consequently, there is only one writer associated with the PVV's tweets. Because of this, the PVV's tweets can be seen as a direct representation of Geert Wilders' personal views and writing style. With a single writer, it is more likely that the tweets from the PVV will exhibit a consistent style and tone. Because the PVV is such an important populist party in the Netherlands, however, this party was still included. More on the writing process of the tweets per party can be found in appendix B.

4.3 Data splitting

The dataset has been divided into three subsets: a training set comprising 60% of the data, a development set consisting of 15%, and a test set consisting of 15% of the original tweets. This data splitting ensures that the model is trained on a substantial amount of data while also allowing for robust evaluation and fine-tuning. For this division of data, the tweets have been entirely randomized, ensuring a representative distribution in all three datasets. Details regarding the distribution of tweets across these sets can be observed in Table 4.2.

As shown in Figures 4.1 and 4.2, the dataset used in this thesis is slightly imbalanced, with 56.7% of the tweets belonging to the non-populist class and 43.3% belonging to the populist class. This class imbalance may introduce a bias towards the majority class, potentially affecting the model's predictions. This has been discussed in section 3.1.

The dataset consists of tweets from different political parties, with variations in

	Train data	Dev data	Test data	
Populist	**9.409**	**1.998**	**2.004**	***13.441***
- PVV	3.626	764	285	*4.675*
- SP	2.386	505	460	*3.351*
- FvD	3.397	729	733	*4.859*
Non-populist	**12.264**	**2.630**	***2.628***	***17.522***
- VVD	1.485	328	285	*2.098*
- D66	3.911	815	827	*5.553*
- PvdA	2.352	533	481	*3.366*
- GroenLinks	3.038	665	686	*4.389*
- ChristenUnie	594	119	132	*845*
- SGP	884	170	132	*1.186*
	21.673	4628	4632	30.933

Table 4.2: Amount of tweets per party.

the number of tweets. Notably, parties like ChristenUnie (CU), SGP, and VVD have a relatively small percentage of tweets in the dataset. On the other hand, the three populist parties PVV, FvD, and SP are relatively similar in terms of the number of tweets represented in the dataset. This distribution is advantageous as it reduces the likelihood of the model favoring a specific party's language patterns over general populist patterns.

Furthermore, the distribution of tweets in the testing data aligns roughly with that of the training data, ensuring consistency in the evaluation process. By maintaining similar class distributions in both the training and testing datasets, the model's performance can be assessed in a more reliable and representative manner.

4.4 Preprocessing

To safeguard privacy, all usernames within the dataset have been substituted with a placeholder. Similarly, any website links present in the tweets have also been replaced to ensure anonymity. Additionally, clear party giveaways such as party names and hashtags have been substituted with placeholders as well.

In order to enhance the quality and relevance of the dataset, a filtering step was implemented to remove the tweets containing less than five words. This step was done after the initial data splitting, as it turned out to be necessary later in the process. Typically, these shorter tweets consisted of brief responses to other users and often lacked substantial or meaningful content. As a result of this cleaning process, a total of 108 tweets (2.28%) were removed from the test set, and 488 tweets (2.02%) were removed from the training set. Due to the similar relative amount of tweets removed in both datasets, this did not lead to any issues.

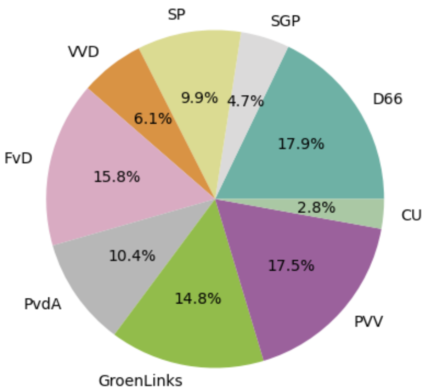

(a) Party division

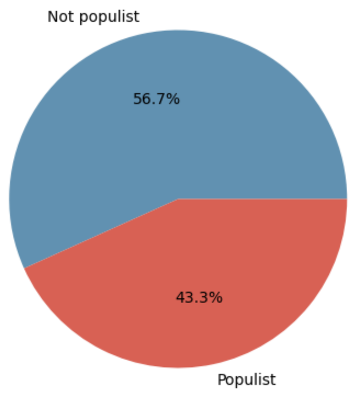

(b) Populism division

Figure 4.1: The division between political parties (4.1a) and populist versus non-populist parties (4.1b) in the testing dataset.

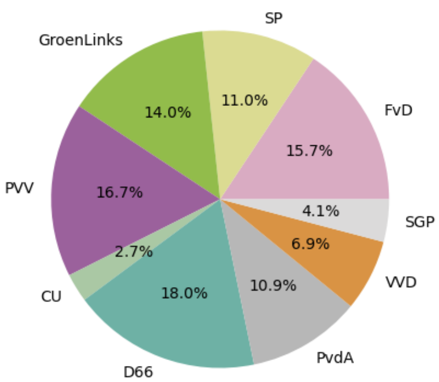

(a) Party division

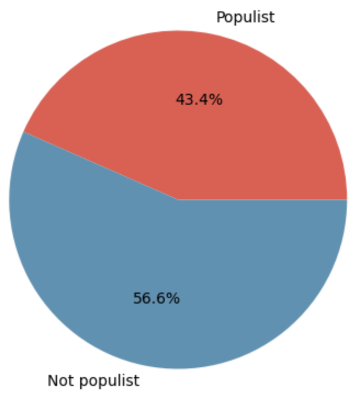

(b) Populism division

Figure 4.2: The division between political parties (4.2a) and populist versus non-populist parties (4.2b) in the training dataset.

Chapter 5

Results

This chapter presents the results of the experiments that aimed to assess the impact of lexical choice and linguistic simplicity on the classification model's ability to distinguish between populist and non-populist language. Section 5.1 discusses the experiments that relate to lexical choice, and section 5.2 refers to linguistic simplicity. All results can be found in Appendix D.

5.1 Lexical choice

The experiments explored here aimed to assess the impact of lexical choice on the classification model's ability to distinguish between populist and non-populist language. The main feature used in these experiments was TF-IDF, which measures the importance of words in a tweet by considering their frequency in the tweet and rarity across the entire dataset.

The experiments further explored two subfeatures within the lexical choice category: topic words and function words. By investigating the performance of the classification model with and without these subfeatures, the experiments aimed to determine their contribution to distinguishing between populist and non-populist language.

5.1.1 Effect of lexical choice

As can be seen in Table 5.1, words are a useful feature when predicting whether a tweet is posted by a populist or non-populist party. The table also illustrates the effects of removing function words and topic words. After removing these types of words, the score is lower but still relatively high. Notably, the removal of function words has a larger impact on the model compared to topic words. It is important to note that while all function words were removed, only a selected subset of the most impactful topic words were excluded from the analysis.

5.1.2 Function words

According to the results presented in Table 5.1, the model's performance is slightly affected when function words are removed. However, a more detailed analysis using feature ablation and addition reveals that individual categories of function words do not substantially impact the model. It is likely that these small effects that each category has, lead to the lower overall score. However, there are a few categories of function

Included words	**Precision**	**Recall**	**F1-score**
All words	0.797	0.789	0.792
(-) Function words	*-0.021*	*-0.020*	*-0.021*
(-) Topic words	*-0.010*	*-0.010*	*-0.010*
No topic/content words	0.767	0.761	0.763

Table 5.1: TF-IDF macro averages of words as features. The drop in performance after the ablation study is shown for function words and topic words. The total score without topic and content words is represented in the bottom row. (-) indicates ablation.

words that impact the model more than the other categories. An overview of these scores can be found in Table 5.2.

Features	**Precision**	**Recall**	**F1-Score**
Content words	0.767	0.761	0.763
(+) quantifiers	*+0.003*	*+0.002*	*+0.003*
(+) pronouns	*+0.006*	*+0.006*	*+0.007*
(+) adverbs	*+0.004*	*+0.003*	*+0.003*
Content and function words	0.787	0.779	0.782
(-) prepositions	*-0.004*	*-0.003*	*-0.004*
(-) quantifiers	*-0.003*	*-0.003*	*-0.003*
(-) pronouns	*-0.008*	*-0.008*	*-0.008*
(-) auxiliary verbs	*+0.004*	*+0.004*	*+0.004*

Table 5.2: Precision, Recall, and F1-Score for TF-IDF, with different function word types added (+) and removed (-) (drop or raise in macro averages). For all of these scores, topic words were removed. Only the scores that have a difference of at least 0.003 from the original scores are shown. For all scores, see Appendix D.3 and D.4.

Specifically, pronouns and quantifiers demonstrate a notable impact on the model's predictive performance. The model shows improved performance when these two types of function words are added as features and reduced performance when they are removed. Thus, it is likely that pronouns and quantifiers influence the model's classification capabilities the most.

5.1.3 Topic words

Class	**Precision**	**Recall**	**F1-score**
Scores with topic words			
Non-populist	0.786	0.835	0.810
Populist	0.765	0.703	0.732
Scores without topic words			
Non-populist	0.780	0.828	0.803
Populist	0.754	0.695	0.723

Table 5.3: TF-IDF without function words, with and without topic words.

Next, when removing the most important topic words (Table 5.3), the performance of the model does not decrease a lot. The impact on the populist class is slightly bigger than on the non-populist class. Due to this, it is likely that topic words do not play a large role in classifying populist and non-populist tweets.

5.1.4 Most impactful terms

When investigating what words the model uses to predict whether a tweet comes from a populist party, the words as shown in Table 5.4 in order of importance are identified. Interestingly, the populist terms are more frequently emotional in nature, while the non-populist words are generally more neutral.

Populist	**Non-populist**
waanzin (*madness*)	europarlementariër (*MEP*)
schandalig (*outrageous*)	twijfel (*doubt*)
rechtvaardigheid (*justice*)	thema (*theme/topic*)
verzet (*resistance*)	bijdrage (*contribution*)
tuig (*scum*)	check (*check*)
massaal (*en masse*)	manifest (*manifest*)
volk (*people/nation*)	wetsvoorstel (*bill (law proposal)*)
leugen (*lie*)	bijvoorbeeld (*for example*)
fantastisch (*fantastic*)	keuze (*choice*)
kamerdebat (*debate*)	rechts (*right (wing)*)

Table 5.4: Most important words to the model.

The results of the participant survey regarding the classification of words into different groups are presented in Figure 5.1. A detailed breakdown of the scores for each word can be found in AppendixC. Overall, participants demonstrated a good ability to recognize which words belong to which group (populist or non-populist).

Of particular interest are the words "rechtvaardigheid" (*justice*), "fantastisch" (*fantastic*), and "kamerdebat" (*debate in parliament*) in the populist category, as well as"manifest" (*manifest*) and "rechts" (*right (wing)*) In the non-populist category. These words were not correctly recognized by participants as belonging to their respective categories as consistently as the other words.

For the words "rechtvaardigheid" and "manifest", all possible categories (left-wing populism, left-wing non-populism, right-wing populism, and left-wing populism) were chosen, indicating uncertainty in their classification. The same was true for the Dutch term for "debate in parliament", although here there was a preference for non-populist parties. "Fantastisch" was predominantly considered non-populist, with 31% of participants assigning it to left-wing non-populism and 82.8% assigning it to right-wing non-populism.

An interesting finding is that the word "tuig" (*scum*) was exclusively associated with right-wing parties, with a majority (93.1%) recognizing it as a populist word. Only 24.1% of participants thought it might be a right-wing non-populist term. Conversely, the word "rechts" was mostly assigned to left-wing parties, although not as convincingly.

Table 5.5 provides a list of words that were found to be important for multiple classification tasks. As these words were the most important, they are expected to be an indicator of populist or non-populist speech. It shows that there is an overlap among the words belonging to the populist class. These words commonly correspond to opposition, right-wing, and conservative parties. It is not surprising, given that all three

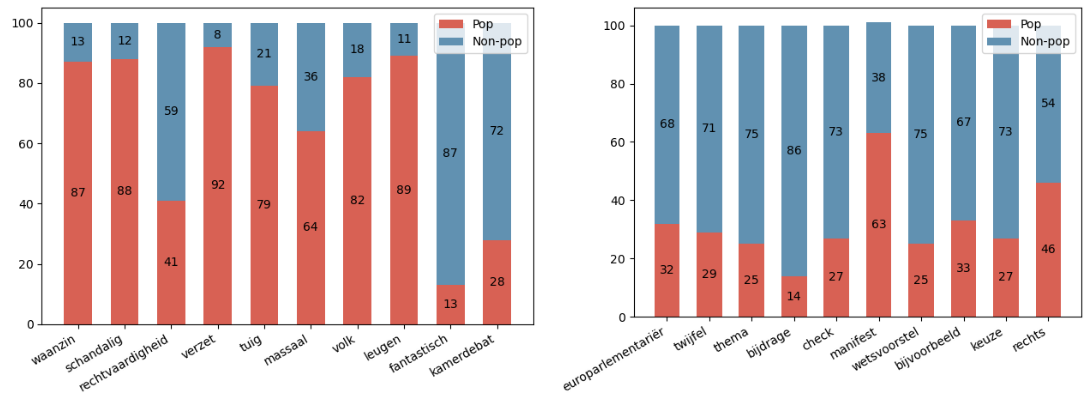

(a) Populist words

(b) Non-populist words

Figure 5.1: The associations that participants (n=29) have with the words identified by the model as the most relevant in distinguishing populist versus non-populist tweets.

populist parties are in the opposition, and two of them are right-wing and conservative. However, this shows that caution should be exercised in drawing conclusions about these words being exclusively populist. As some of them exclusively appear in the populist vs. non-populist distinction, this suggests that these words might be specific to this distinction, but due to the large overlap with other categorizations, this cannot be exclusively stated.

Another noteworthy observation is that the left/right distinction and the progressive/conservative distinction yield the same set of words. This alignment can be attributed to the Dutch political landscape, as discussed in Chapter 3 (Table 3.1), where left-wing parties are typically progressive, while right-wing parties tend to be conservative. The exception to this pattern is the ChristenUnie, which is left-wing but conservative. This party had a relatively small number of tweets compared to the other parties and thus might not have influenced the model much.

	coal	**opp**	**left**	**right**	**prog**	**cons**
Populist words						
waanzin (*madness*)		x		x		x
rechtvaardigheid (*justice*)			x		x	
verzet (*resistance*)		x		x		x
tuig (*scum*)		x		x		x
volk (*people/nation*)				x		x
kamerdebat (*debate in parliament*)		x		x		x
Non-populist words						
twijfel (*doubt*)			x		x	
check (*check*)			x			

Table 5.5: Overlap of words that are important to the model for different classes.

5.1.5 Control experiments

All in all, the relevance of lexical choice seems to be a good feature for all classification tasks, as can be seen in 5.6. The only classification that seems to work less well than others is the model that distinguishes between coalition and opposition. Of all the different distinctions, lexical choice works the best for populism vs. non-populism. The most important topic words were removed for populism vs. non-populism but not for the other categories, so it is interesting that that does not seem to increase their score compared to populist vs. non-populist categorization. This might mean that either topic words are not as important as previously thought, or the manner of removing topic words as done in this thesis was not sufficient to test their impact.

	Precision	**Recall**	**F1-score**
Pop vs. non-pop	0.767	0.761	0.763
Coal vs. opp	0.714	0.647	0.665
Left vs. right	0.755	0.751	0.753
Coal vs. Prog	0.754	0.752	0.752

Table 5.6: Macro averages of lexical choice as a feature for different classification tasks. Function words and topic words are left out for all models.

5.2 Linguistic simplicity

The experiments discussed in this paragraph explored the impact of linguistic simplicity on the classification model's performance in distinguishing between populist and non-populist language. This was done through exploring average word amount per tweet, Leesindex A, Flesch, and Flesch-Douma as features.

5.2.1 Effect of linguistic simplicity

The analysis of linguistic simplicity as a feature reveals that the model does not perform well when trained on these features, as shown in Table 5.7. The model demonstrates a preference for the non-populist class, which suggests that the score corresponds to the majority class, and thus linguistic simplicity is not a good indicator of populist vs. non-populist texts.

	Precision	**Recall**	**F1-score**
Non-populist	0.594	0.848	0.699
Populist	0.547	0.241	0.335
accuracy			0.585
macro avg	0.571	0.544	0.517
weighted avg	0.574	0.585	0.541

Table 5.7: Precision, Recall and F1-Score for linguistic simplicity.

5.2.2 Effect of linguistic simplicity subfeatures

To further assess the influence of linguistic simplicity features, an ablation analysis was conducted, as presented in Table 5.8. The results indicate that removing individual features has little effect on the model's performance, suggesting that these subfeatures do not substantially contribute to the classification task.

Features	**Precision**	**Recall**	**F1-Score**
All features	0.571	0.544	0.517
(-) w/s	*+0.001*	*+0.002*	*+0.001*
(-) Leesindex A	*-0.001*	*0*	*-0.001*
(-) Flesch	*-0.001*	*0*	*-0.001*
(-) Flesch-Douma	*0*	*+0.001*	*+0.001*

Table 5.8: Macro averages of linguistic simplicity as features. The drop in performance after the ablation study is shown for individual subfeatures. (-) indicates ablation.

Table 5.9 provides mean scores and standard deviation scores for the different features. It can be observed that there is a slightly higher score for non-populist parties, but the overall scores are relatively similar.

To give more insight into the linguistic simplicity scores, mean scores and standard deviation scores of the different features are shown in Table 5.9. This shows a slightly higher simplicity score for non-populist parties, which would indicate an easier text, but overall the scores are pretty similar. This suggests that, on average, the linguistic simplicity of tweets from both populist and non-populist parties is comparable. This is reflected in the model scores and the limited impact of the features on the model.

	Populist		**Non-populist**	
	mean	*SD*	*mean*	*SD*
w/s	14.39	7.83	13.75	7.22
Leesindex A	63.86	23.83	71.26	23.9
Flesch	57.52	25.32	66.11	25.15
Flesch-Douma	70.86	23.06	78.68	22.9

Table 5.9: Mean and standard deviation scores for the different simplicity features.

5.2.3 Control experiments

	Precision	**Recall**	**F1-score**
Pop vs. non-pop	0.571	0.544	0.517
Coal vs. opp	0.380	0.500	0.432
Left vs. right	0.588	0.545	0.502
Cons vs. prog	0.581	0.560	0.738

Table 5.10: Macro averages of linguistic simplicity as a feature for different classification tasks.

When examining different classification tasks trained with these features, it appears that linguistic simplicity is not an effective feature for classification (see Table 5.10). The model's performance seems to build on the statistical likelihood of the majority class.

When we investigate different classification tasks trained with the same features, we find that linguistic simplicity as a feature is not an effective feature.

Chapter 6

Conclusion

This study aimed to investigate the impact of lexical choice and linguistic complexity on the classification of populist and non-populist language. Through a series of experiments, the effectiveness of these two language characteristics was examined.

The results indicate that lexical choice is indeed a valuable feature for predicting whether a tweet is written by a populist or non-populist party. Even after removing function words and topic words, the classification model still achieved relatively high scores.

Further analysis of function words reveals that individual categories of function words do not have a big influence on the model. However, pronouns and quantifiers show to be notable contributors to the model's predictive performance. Their inclusion as features improved the model's classification capabilities, while their removal resulted in reduced performance.

In the case of topic words, the removal of the most impactful ones had a relatively small effect on the model's performance. The impact was slightly more pronounced for the populist class. Thus, topic words seem to play a limited role in the populist style.

The investigation of the most important terms used by the model highlighted an interesting pattern. Populist words tend to be more emotional in nature, while non-populist words are generally more neutral. The participant survey supports these distinctions in lexical choice.

Additionally, the analysis reveals an overlap of words that are important for multiple classification tasks. This overlap suggests that caution should be exercised in drawing conclusions about the exclusivity of certain words to the populist or non-populist category. Furthermore, the alignment between left/right and progressive/conservative distinctions in the Dutch political landscape is important to keep in mind when interpreting these results.

The analysis of linguistic simplicity as a feature reveals that the model did not perform well when trained on these features. The model exhibited a preference for the non-populist class, suggesting that the score corresponded to the majority class and, thus, linguistic simplicity is not a reliable indicator of populist and non-populist tweets.

Furthermore, an ablation analysis was conducted to assess the influence of individual linguistic simplicity subfeatures. The results show that removing these subfeatures had a minimal effect on the model's performance, suggesting that they do not contribute to the recognition of populist tweets in an important way.

Control experiments were also conducted to assess the effectiveness of linguistic simplicity as a feature for different classification tasks. The results indicate that linguistic simplicity is not an effective feature in other classifications either.

This study demonstrates the importance of lexical choice in distinguishing between populist and non-populist language. Linguistic simplicity does not noticeably assist in this classification.

Chapter 7

Discussion

The findings in this thesis highlight the interconnectedness of different political dimensions and the nuanced nature of political language. It underscores the importance of considering multiple factors when analyzing and interpreting the associations between words and political categories. It is also a step towards identifying markers of populism (in this case lexical choice). These markers can aid in the automatic classification and detection of populist discourse in large-scale text data, contributing to the development of tools and methods for tracking and monitoring populism in political and social media contexts. This can assist in identifying trends, assessing the spread of populist ideas, and monitoring the impact of populist rhetoric on public opinion.

In this chapter, some of the limitations, observations, and directions for future research are discussed.

7.1 Lexical choice and emotion

Upon examining the words that exhibited the strongest indications of populist or non-populist speech to the model, it became apparent that these words evoke strong emotions and convey a sense of intensity commonly associated with populist rhetoric (Mudde, 2004). In contrast, non-populist speech was associated with terms that were more neutral in nature.

In general, human participants displayed a notable level of accuracy in recognizing and categorizing this pattern as either populist or non-populist. This suggests that there is an intuitive understanding among individuals regarding the distinct linguistic patterns and connotations associated with populist discourse as found in this thesis.

Because of this, focusing on sentiment as a characteristic of populist speech is an intriguing direction for future research. Sentiment analysis involves determining the emotional tone or polarity of a text, which can provide valuable insights into the language used by populist parties and their supporters. Furthermore, investigating sentiment can shed light on the relationship between populist discourse and public opinion. Understanding how populist messages elicit specific emotions and resonate with different segments of the population can give insights into the appeal and influence of populist ideologies. The importance of this relationship is highlighted by Rodriguez-Pose (2022) as well.

7.2 Quantifiers, pronouns, and adverbs

The role of quantifiers, pronouns, and adverbs in populist speech, as identified in this thesis, also presents an interesting avenue for future research. These types of words may carry particular significance and rhetorical functions in the stylistic choices of populist parties. Understanding how these types of words are used in the context of populism can shine a light on the linguistic and persuasive techniques employed by populist leaders and parties. While the relevance of pronouns in populist speech has been discussed in literature before, this has not yet been done as extensively for quantifiers and adverbs (Szilagy, 2017; Engesser et al., 2017).

It would be valuable to explore their rhetorical purposes, such as framing, emotional appeals, or identity construction, and examine how they contribute to the overall persuasive strategies employed. Additionally, investigating the semantic and pragmatic nuances of these words in different populist contexts and cultural settings can shed light on their specific roles in shaping populist narratives and the way they mobilize public support.

Furthermore, comparative studies such as this one could be conducted to analyze the use of quantifiers, pronouns, and adverbs in non-populist political speech, allowing for a better understanding of the distinctive features that set populist language apart.

7.3 Topic words

A limitation in this thesis was the removal of topic words. When testing the impact of topic words, all topic words should be removed to test their influence on the model. However, that was difficult to do and instead only the most impacting words were removed. Because of this, a conclusive statement on the role of topics in populist speech can not be given. The model might have learned to associate certain topics with certain parties and thus interfered with the analysis. Future research should look into this issue.

One possible approach would be to develop a way to identify topic words in political speech, through either a term list or a classification task. While a term list would require a lot of manual labor, a classification task could automate this process. This could involve training a machine learning model to classify words as either topic words or non-topic words. It is important to note that the success of this classification task would depend on the availability of a well-labeled training dataset that accurately captures the topic words specific to the political context under study.

7.4 Overlap in characteristics

The findings in this thesis raise an important consideration: the identified lexical choices, as explored in this thesis, may not be exclusively representative of populist speech alone but may also capture elements of other party characteristics. The overlap between these categories, as demonstrated in section 5.1.5, suggests a potential intertwining of features, blurring the boundaries between (especially) populist discourse and right-wing political ideologies. This underscores the complexity of language analysis and the need for a nuanced understanding of the political context in which these words are utilized.

This emphasizes the necessity for a nuanced understanding of the political context in which these words are employed. Themes and narratives can be shared by, for example, both populist and right-wing political parties: certain words related to immigration or cultural identity may be used by both types of parties. This overlap between linguistic features suggests that a simplistic categorization of language features as exclusively populist or right-wing ideology may oversimplify political language. It highlights the need for a more nuanced understanding of the contextual factors that shape language use and the ideologies that underpin it. Further research and analysis are required to disentangle the extent to which the observed lexical patterns are attributable to populist rhetoric specifically or are indicative of broader ideological tendencies.

7.5 Twitter as a medium

When it comes to linguistic simplicity, it is important to consider the unique characteristics of the tweet format when interpreting the results related to linguistic simplicity. While simplicity measures can provide insights into the ease of understanding and accessibility of longer texts, their applicability to the condensed nature of tweets should be carefully considered. The format of tweets, being short and concise online messages, can influence linguistic characteristics. The limited character count and the fast-paced nature of social media platforms often require users to convey their message within a restricted space, resulting in the use of shorter sentences and condensed language. This condensed format might lead to a simplified expression compared to other forms of political communication, such as party programs or statements in debates.

Future studies could explore alternative methods or adaptations of linguistic simplicity measures to better capture the nuances of linguistic simplicity in the context of social media platforms. Taking another source of populist texts is another way to further research the linguistic simplicity of populist discourse. Studying longer political texts such as party manifestos, policy documents, or speeches, can offer a more comprehensive understanding of populist language.

7.6 Other limitations and future research

It is worth noting that the results found in this thesis may be specific to the dataset and methodology used in this study. Further research and exploration of additional features are needed to gain a deeper understanding of the relationship between lexical choice, linguistic simplicity, and political discourse.

Exploring role changes and their impact on populist language is an intriguing direction for future research. Investigating whether the language use of populist parties undergoes notable changes when their role in parliament shifts can provide valuable insights into the dynamics of populist rhetoric and different kinds of this rhetoric. This could shed light on how populist parties adapt their communication strategies to different contexts and power dynamics, and whether there are distinct linguistic patterns associated with these role changes.

Expanding the scope of the analysis beyond the three populist parties studied in this thesis is another promising direction for future research. A focus on balancing the data when it comes to parties and their political ideologies (such as left- and right-wing) is important as well to limit unwanted and informative correlations. Including a broader range of populist parties, both within and across different countries and languages, enables a more comprehensive and comparative analysis of populist language. This approach can help identify common patterns, cross-cultural variations, and underlying linguistic mechanisms that transcend specific contexts. The correlation between right-wing conservative parties and populism could be explored more in-depth as well.

Bibliography

Y. Algan, S. Guriev, E. Papaioannou, and E. Passari. The european trust crisis and the rise of populism. *C.E.P.R. Discussion Papers*, 12444, November 2017.

E. Block and R. Negrine. The populist communication style: toward a critical framework. *International Journal of Communication*, 11:178–197, 2017.

R. Brouwer. nderzoek naar de leesmoeilijkheden van nederlands proza. *Pedagogische Studiën*, 40:454–464, 1963.

M. Canovan. Trust the people! populism and the two faces of democracy. *political studies*, 47(1):2–16, 1999.

J. S. Chall and E. Dale. *Readability Revisited: The New Dale-Chall Readability Formula.* Brookline Books, Cambridge, MA, 1995.

W. Douma. De leesbaarheid van landbouwbladen: een onderzoek naar en een toepassing van leesbaarheidsformules. *Bulletin*, 17, 1960.

S. Engesser, N. Ernst, F. Esser, and F. Büchel. Populism and social media: how politicians spread a fragmented ideology. *Information, Communication & Society*, 20(8):1109–1126, 2017.

R. Flesch. A new readability yardstick. *Journal of Applied Psychology*, 32:221–233, 1948.

FvD. Forum voor democratie). `https://fvd.nl`, 2023a. Accessed: February 3, 2023.

FvD. Standpunten: Waar staat forum voor democratie (fvd) in het politieke spectrum? fvd.nl, 2023b. URL `https://fvd.nl/standpunten`.

K. Hawkins. Is chávez populist? measuring populist discourse in comparative perspective. *Comparative Political Studies*, 42(8):1040–1067, 2009.

G. Hirst, Y. Riabinin, and J. Graham. 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)*, Rome, june 2010.

N. Hirzalla. A natural language processing based text analysis of populist rhetoric in social media text messages. *LingUU*, 3(2), november 2019.

House of Representatives. Democracy in the netherlands. `https://www.houseofrepresentatives.nl/how-parliament-works/democracy-netherlands`, 2011. Accessed: February 3, 2023.

O. Kayam. The readability and simplicity of donald trump's language. *Political Studies Review*, 16(1):73–88, 2018.

P. J. Kincaid, R. P. Fishburne, R. L. Rogers, and B. S. Chissom. Derivation of new readability formulas (automated readability index, fog count, and flesch reading ease formula) for navy enlisted personnel. In *Research Branch Report*, volume 56, pages 8–75. Institute for Simulation and Training, 1975. URL `https://stars.library.ucf.edu/istlibrary/56/`.

W. H. Lin, T. Wilson, J. Wiebe, and A. Hauptmann. Which side are you on? identifying perspectives at the document and sentence levels. In *Proceedings of the 10th Conference on Natural Language Learning (CoNLL-X)*, pages 109–116, 2006.

D. McDonnell and S. Ondelli. The language of right-wing populist leaders: Not so simple. *Perspectives on Politics*, 20(3):828–841, 2020. doi: https://doi.org/10.1017/S1537592720002418.

M. Meijers and A. Zaslove. Populism and political parties expert survey 2018. POPPA, 2018. URL `www.poppa-data.eu`.

M. J. Meijers and A. Zaslove. Measuring populism in political parties: Appraisal of a new approach. *Comparative Political Studies*, 54(2):372–407, 2021. doi: 10.1177/0010414020938081. URL `https://doi.org/10.1177/0010414020938081`.

B. Moffitt and S. Tormey. Rethinking populism: politics, mediatisation and political style. *political studies*, 62(2):381–397, 2014.

C. Mudde. The populist zeitgeist. *Government and Opposition*, 39:541–563, 2004.

C. Mudde. *Populist Radical Right Parties in Europe*. Stanford University Press, Stanford, CA, 2007.

R. Németh. A scoping review on the use of natural language processing in research on political polarization: trends and research prospects. *Journal of Computational Social Science*, 6:289–313, 2023. URL `https://doi.org/10.1007/s42001-022-00196-2`.

N. Ottens. Dutch far-right forum for democracy implodes. `https://atlanticsentinel.com/2020/11/dutch-far-right-forum-for-democracy-implodes/`, 2020.

PVV. Het gaat om u: Verkiezingsprogramma 2021-2025. pvv.nl, 2021. URL `https://www.pvv.nl/images/09012020/verkiezingen2020/0acxyuew34z/VerkiezingsProgramma2021-Final.pdf`.

A. Rodriguez-Pose. The rise of populism and the revenge of places that don't matter. In A. Velasco and I. Bucelli, editors, *Populism: Origins and Alternative Policy Responses*, chapter 4, pages 79–103. LSE Press, London, 2022.

M. Rooduijn. The nucleus of populism: in search of the lowest common denominator. *government and opposition*, 49(4):573–599, 2014.

SP. Standpunten. sp.nl, 2023. URL `https://www.sp.nl/standpunten`.

A. Szilagy. How to talk like a populist. VICE News, 2017. URL `https://www.vice.com/en/article/gyd98y/how-to-talk-like-a-populist`.

S. Tunderman. The political economy of populist reason in the netherlands. *New Political Science*, 44(4):507–523, 2022. doi: 10.1080/07393148.2022.2129923. URL `https://doi.org/10.1080/07393148.2022.2129923`.

V. Vandeghinste and B. Bulté. Linguistic proxies of readability: Comparing easy-to-read and regular newspaper dutch. *Computational Linguistics in the Netherlands*, 9: 81–100, 2019.

M. Viser. For presidential hopefuls, simples language resonates. *Boston Globe*, 2015.

K. Weyland. Clarifying a contested concept: populism in the study of latin american politics. *Comparative Politics*, 34(1):1–22, 2001.

B. Yu, S. Kaufmann, and D. Diermeier. Classifying party affiliation from political speech. *Journal of Information Technology in Politics*, 5(1):33–48, 2008.

Appendix A

List of removed words

Word	Notes
ha	One party starts replies with a greeting. These have been removed.
hoi	One party starts replies with a greeting. These have been removed.
hi	One party starts replies with a greeting. These have been removed.
inside	The name of an online show of one of the parties (Forum Inside).
^(NAME)	One party ends replies with the name of the person who wrote the reply. These have been removed.
Baudet	Name of politician.
Forum	Shortened version of party name.
FVD	Shortened version of party name.
Geert	Name of politician.
Hans	Name of politician.
Hoekstra	Name of politician.
Raak	Name of politician.
Jesse	Name of politician.
Mark	Name of politician.
Kaag	Name of politician.
Lilian	Name of politician.
Lilianne	Name of politician.
Paul	Name of politician.
Ploumen	Name of politician.
Rob	Name of politician.
Ronald	Name of politician.
Rutte	Name of politician.
Sigrid	Name of politician.
Thierry	Name of politician.
Wilders	Name of politician.
luistertip	
kijktip	

ticket	
socialer	Part of the often-used phrase (by one party) "socialer leiderschap" (*more social leadership*.
leiderschap	Part of the often-used phrase "socialer leiderschap" (*more social leadership*).
studio	The name of an online show of one of the parties (Studio SGP).
uitzending	
vd	Short for "van de" (*of the*), only used by a party with only one writer. More stylistic than a conscious choice and thus removed.
nl	Short for "Netherlands", only used by a party with only one writer. More stylistic than a conscious choice and thus removed.
oa	Short for "onder andere" (*among other things*), only used by a party with only one writer. More stylistic than a conscious choice and thus removed.
app	Used by certain parties who use their WhatsApp to keep in contact with their followers.

Table A.1: Data-specific terms

toeslagenschandaal	klimaatakkoord
zorgverlener	AOW
kringlooplandbouw	klimaatcrisis
duurzaam	abortus
klimaatverandering	progressief
uitstoot	groen
green	europees
leraar	student
klimaat	schoon
vluchteling	grens
koopkracht	groninger
natuur	onderwijs
permanent	coronawet
oorlog	

Table A.2: Topic words

Appendix B

Party’s explanation of Twitter process

B.1 PvdA

April 24th, 2023 (through the official WhatsApp contact number of the online Team PvdA):
"Wij hebben een team van 9 mensen dat onze sociale media beheert. Een Tweet wordt vaak van tevoren bedacht, geschreven en ingepland - omdat we weten dat we op een bepaalde dag over een bepaald onderwerp willen delen. Bijvoorbeeld als er op een bepaalde dag over een wetsvoorstel van ons gestemd wordt.

Eén iemand uit ons team heeft dan ook de taak toegewezen gekregen om een eerste voorstel voor een berichttekst te schrijven. Daarna kijkt een ander paar ogen naar het tekstje. Want vaak kan het nog net wat beter of sterker. Bij het schrijven van berichten om te delen op sociale media gebruiken we zoveel mogelijk taal en woorden die passen bij onze kernboodschap.

Soms moet het wat sneller - als we op het nieuws willen inspelen of op andere politieke actoren willen reageren op Twitter. Dan stuurt iemand snel een voorstel voor een bericht, en wordt het ook snel door de rest van het team gecheckt en akkoord gegeven. Als het om politiek gevoelige dingen gaat, is het ook goed om te checken bij onze politici of hun woordvoerders.”

Translation:
"We have a team of 9 people who manage our social media. A Tweet is often planned, written, and scheduled in advance - because we know we want to share something on a specific topic on a certain day. For example, if there is a vote on a bill we support on a particular day.

One person in our team is assigned the task of writing the first draft of the message text. Then another set of eyes looks at the text. Because often it can be improved or made stronger. When writing messages to share on social media, we use language and words that fit our core message as much as possible.

Sometimes it needs to be faster - if we want to respond to the news or other political actors on Twitter. Then someone quickly sends a proposal for a message, which is also quickly checked and approved by the rest of the team. When it comes to politically sensitive matters, it is also good to check with our politicians or their spokespersons."

B.2 GroenLinks

May 1st, 2023 (through the official mail-address of GroenLinks):
"Ons Online Team schrijft de tweets. Dit wordt vaak door dezelfde persoon geschreven. Wel worden deze tweets afgestemd met de online woordvoerders van de Kamerleden en de Kamerleden zelf."

Translation:
"Our online team writes the tweets. This is often written by the same person. However, these tweets are coordinated with the online spokespersons of the Members of Parliament and the Members of Parliament themselves."

B.3 SP

May 5th, 2023 (through the official email address of the SP):
"Onze Kamerleden die twitteren, doen dat zelf, meestal op het gebied waar ze woordvoerder voor zijn. Ook Lilian Marijnissen schrijft haar eigen berichten en nieuwsbrieven. Algemene Social Media en mailadressen van de SP worden behandeld door de medewerkers van het communicatieteam en het secretariaat/publieksvoorlichting."

Translation:
"Our Members of Parliament who tweet, do so themselves, usually in their respective areas of spokespersonship. Lilian Marijnissen writes her own messages and newsletters as well. General social media accounts and email addresses of the SP are managed by the staff of the communication team and the secretariat/public information."

B.4 D66

May 24th, 2023 (through the official WhatsApp contact number of the online Team D66):
"Wij hebben bij D66 een communicatieteam waarbij een aantal mensen verantwoordelijk zijn voor onder andere het schrijven van tweets. Hierbij adviseren zij Kamerleden en schrijven zij ook tweets voor onze eigen kanalen. De manier waarop de tweet wordt geschreven heeft natuurlijk alles te maken met de doelgroep die we willen bereiken (zijn dit journalisten, burgers, andere kamerleden etc.). Het moet nieuwswaardig zijn of empatisch, het heeft in principe alles te maken met het doel van de tweet."

Translation:
"At D66, we have a communication team consisting of several individuals who are responsible for, among other things, writing tweets. They provide advice to Members of Parliament and also draft tweets for our own channels. The way a tweet is crafted depends on the target audience we want to reach (whether it's journalists, citizens, other Members of Parliament, etc.). It should be newsworthy or empathetic, ultimately the content of the tweet is aligned with its intended purpose."

B.5 SGP

June 5th, 2023 (through the official email address of the SGP):
"Ons landelijke SGP-account valt onder het beheer van meerdere communicatiemedewerkers. Zowel collega's die werken bij de fractie (in Den-Haag) als op het partijbureau (in Rotterdam) sturen zo nu en dan tweets via het account. Wij werken in principe met falcon, een onlineprogramma waarbij je je content in kunt plannen. Echter, niet alles valt natuurlijk te plannen, er zijn altijd adhoc gebeurtenissen waarop je in moet springen. Het ligt dan aan het onderwerp/gebeurtenis welk team iets oppakt en een Twitterbericht plaatst."

Translation:
"Our national SGP account is managed by multiple communication staff members. Both colleagues working at the parliamentary faction (in The Hague) and at the party headquarters (in Rotterdam) occasionally send tweets via the account. In principle, we use Falcon, an online program that allows the user to schedule content. However, not everything can be planned in advance, as there are always ad hoc events that require immediate response. Depending on the subject/event, it is up to the relevant team to address it and post a Twitter message."

B.6 ChristenUnie

June 13th, 2023 (through the official email address of the ChristenUnie):
"De tweets die wij versturen worden door meerdere mensen geschreven. Over het algemeen iemand uit ons social media team, maar dit kan ook een voorlichter of communicatieadviseur zijn. Hoe het proces verloopt hangt af van de zwaarte van het bericht. Als het een inhoudelijk bericht is dan wordt dit altijd kortgesloten met de desbetreffende voorlichter en/of Kamerlid. Voor wat meer laagdrempelige posts, zoals vacatures of aankondigingen voor een event is dat niet nodig en schrijft het social media team dit zelf."

Translation:
"The tweets we send out are written by multiple individuals. Generally, it is someone from our social media team, but it can also be a spokesperson or communication advisor. The process depends on the significance of the message. If it is a substantive message, it is always coordinated with the relevant spokesperson and/or Member of Parliament. For more straightforward posts such as job vacancies or event announcements, coordination may not be necessary, and the social media team handles them themselves."

Appendix C

Associations of participants with words

C.1 Populist words

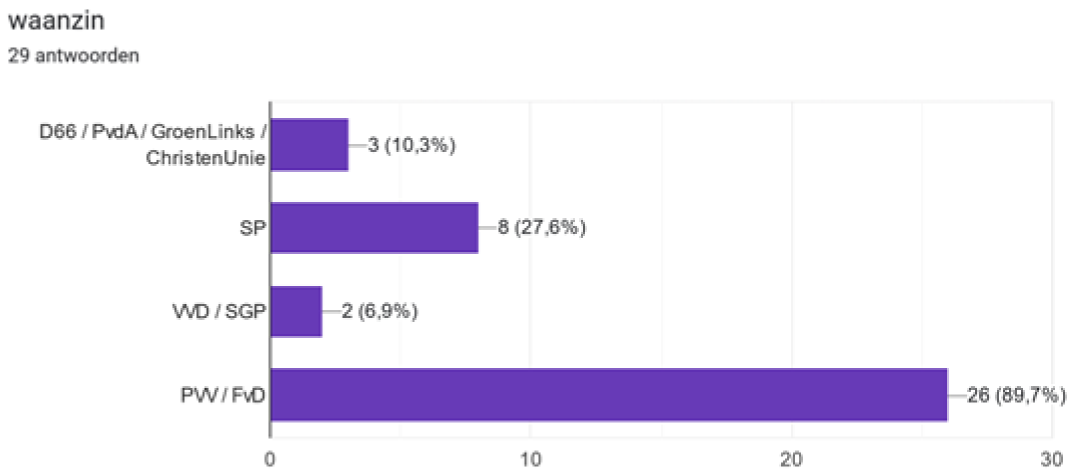

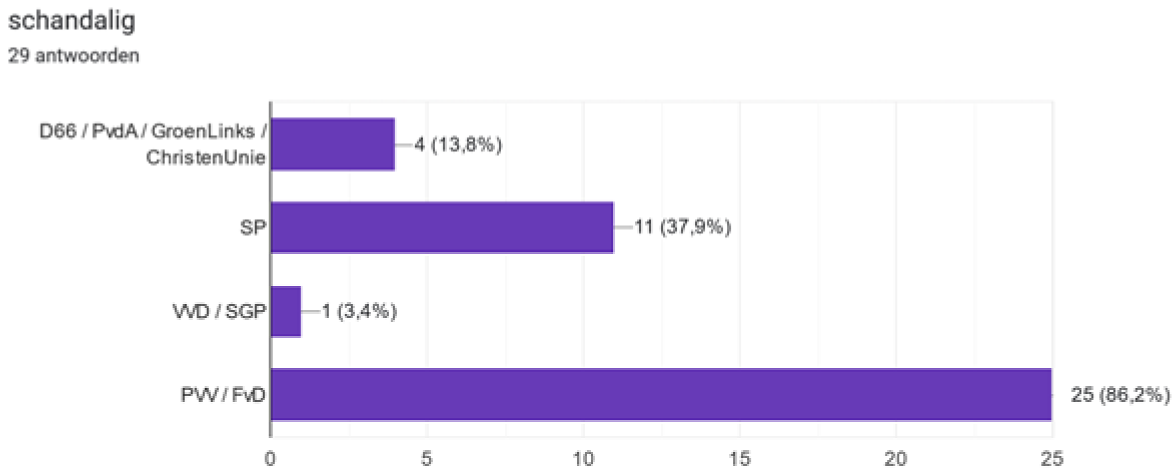

rechtvaardigheid

29 antwoorden

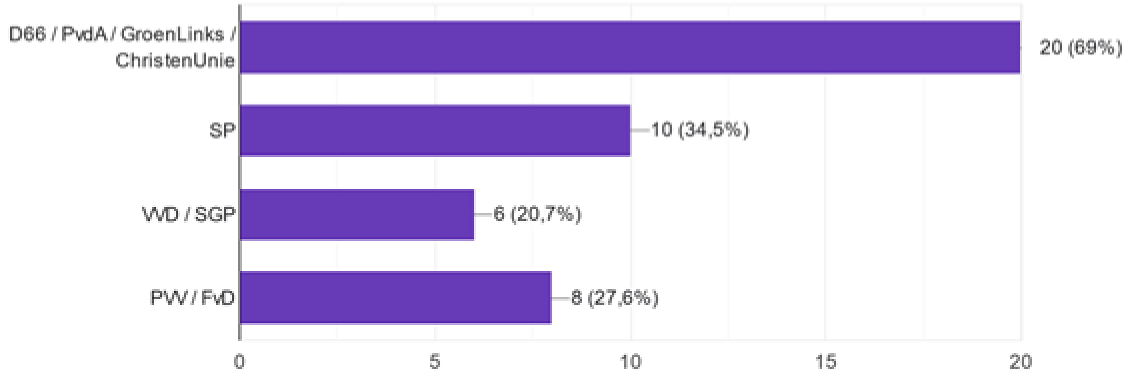

verzet

29 antwoorden

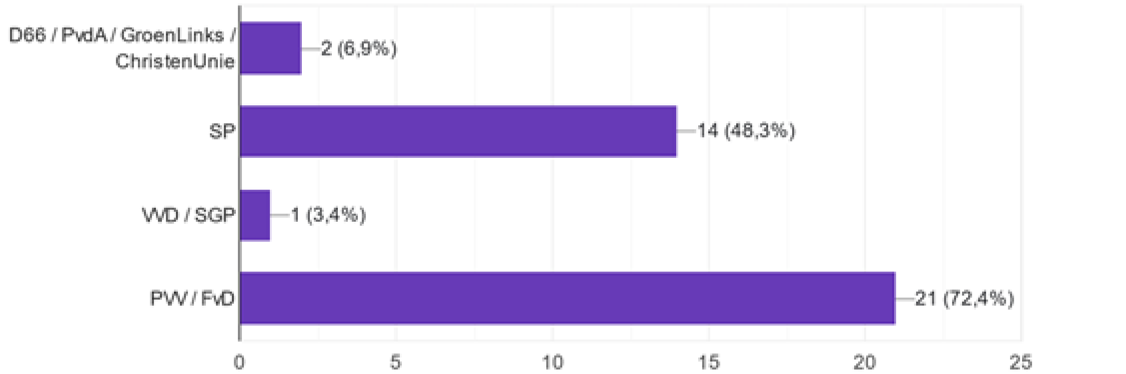

tuig

29 antwoorden

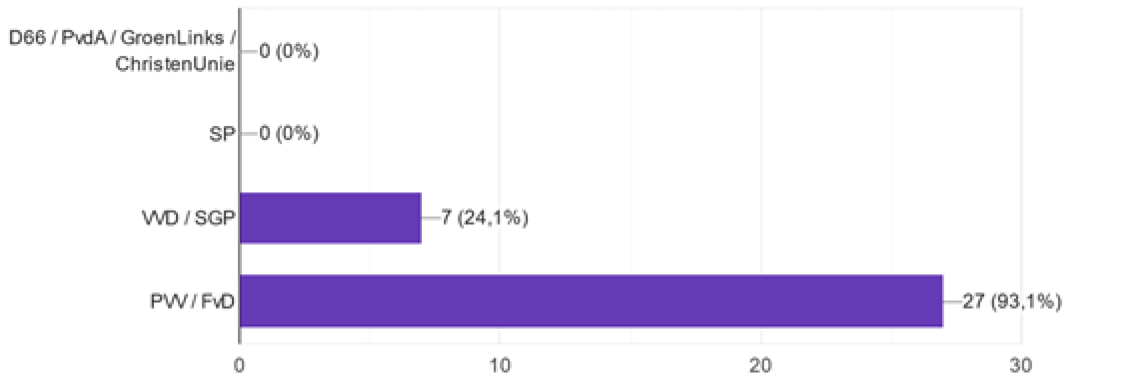

massaal

29 antwoorden

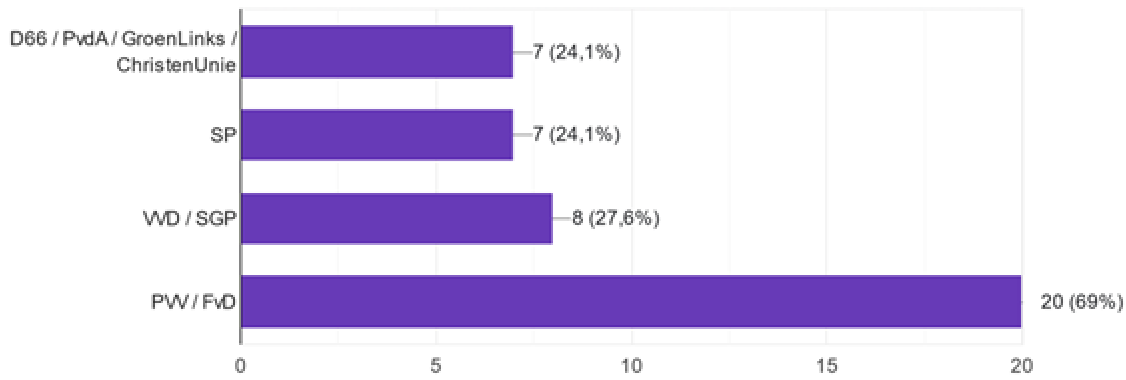

volk

29 antwoorden

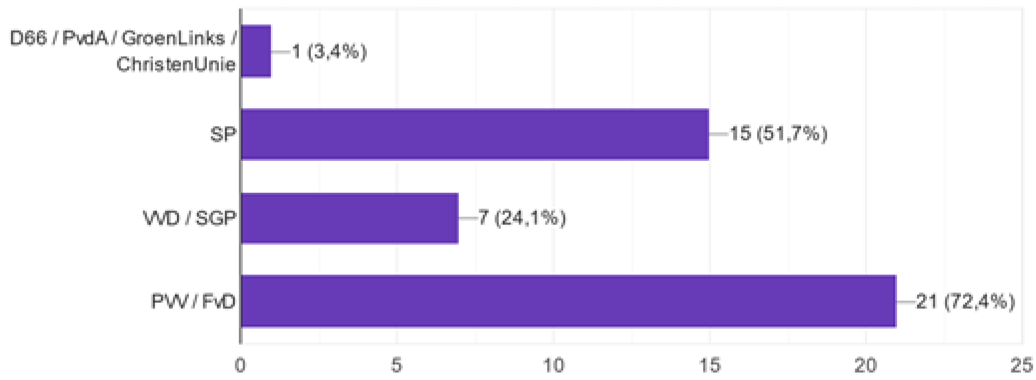

leugen

29 antwoorden

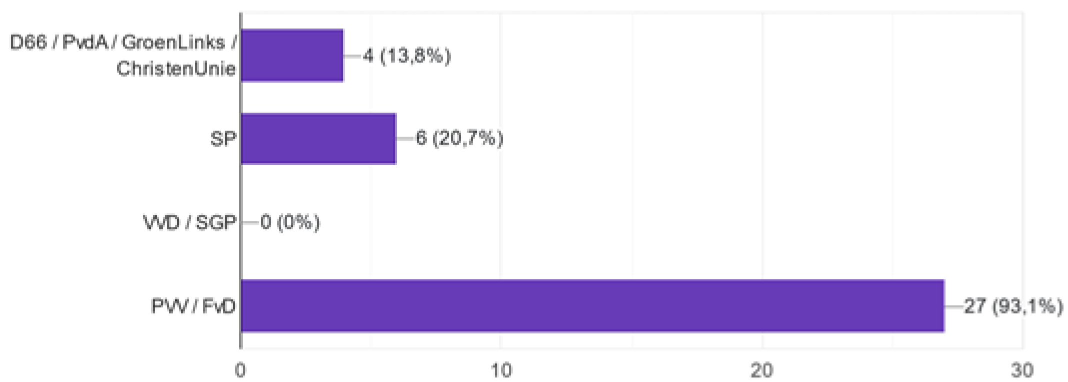

fantastisch

29 antwoorden

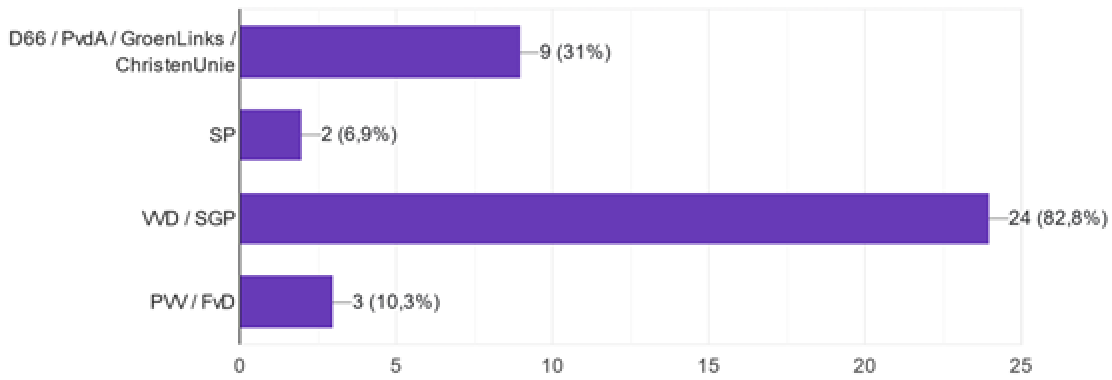

kamerdebat

29 antwoorden

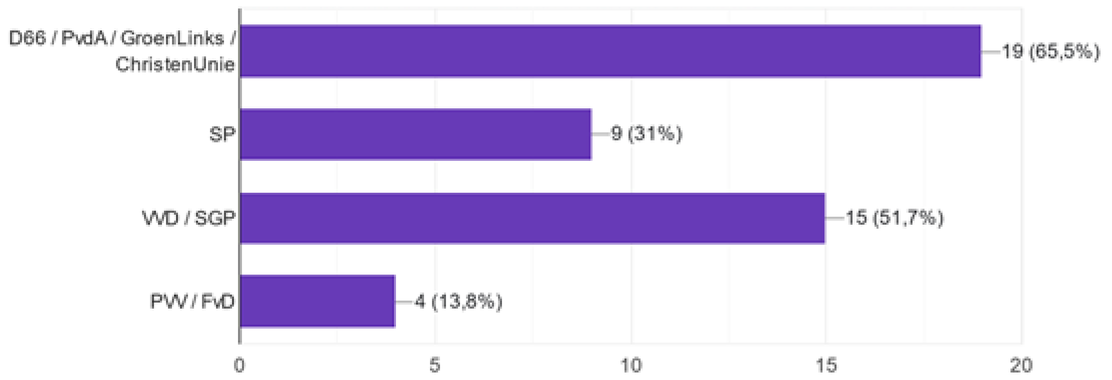

C.2 Non-populist words

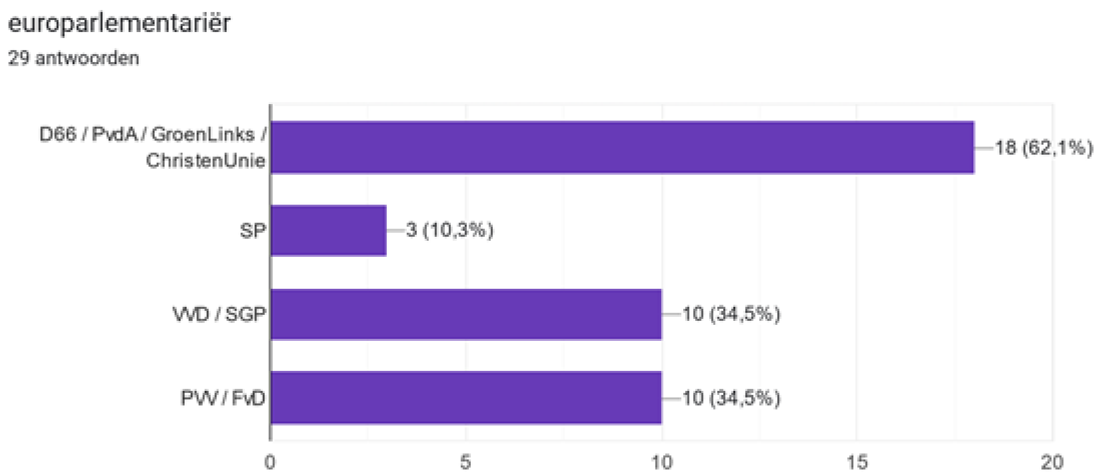

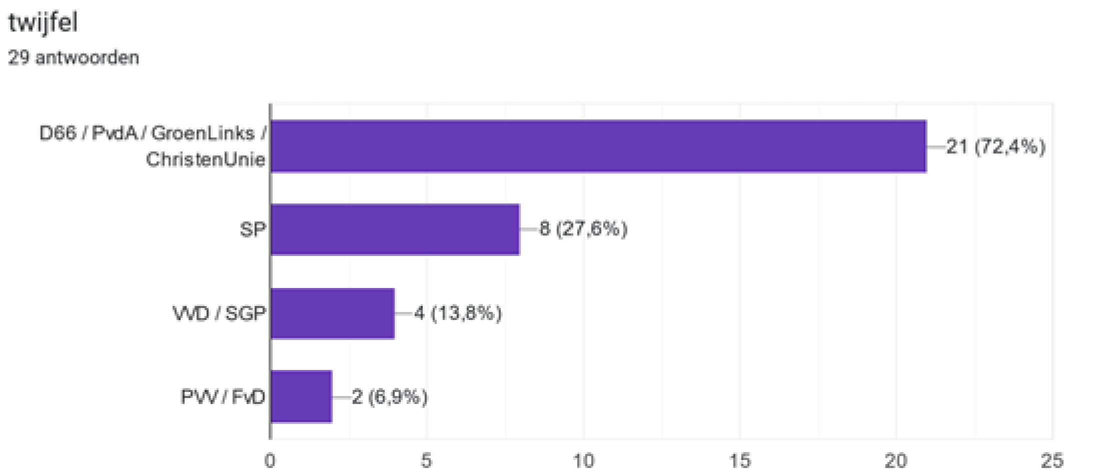

thema

29 antwoorden

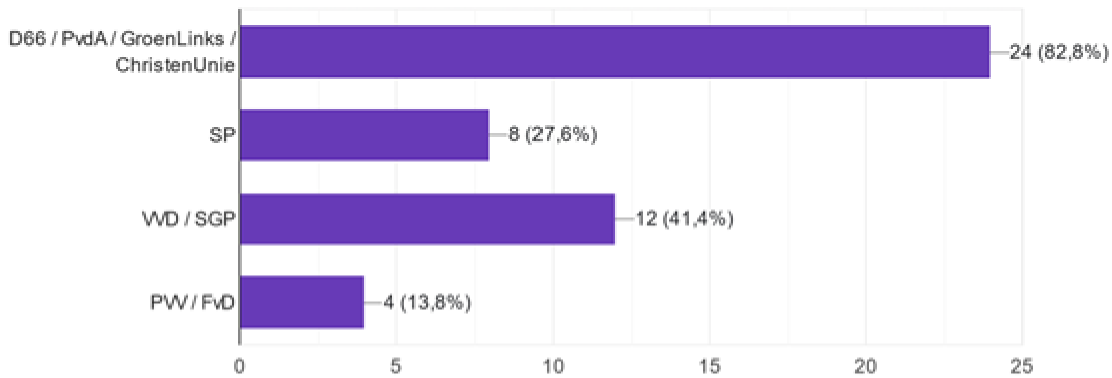

bijdrage

29 antwoorden

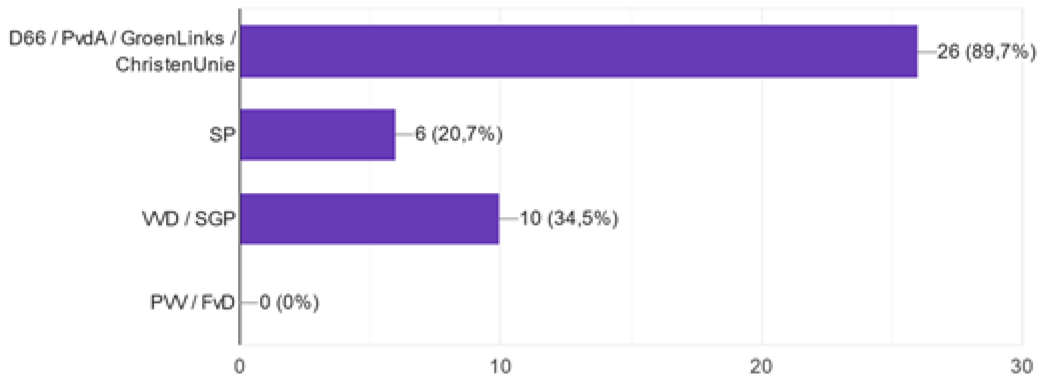

check

29 antwoorden

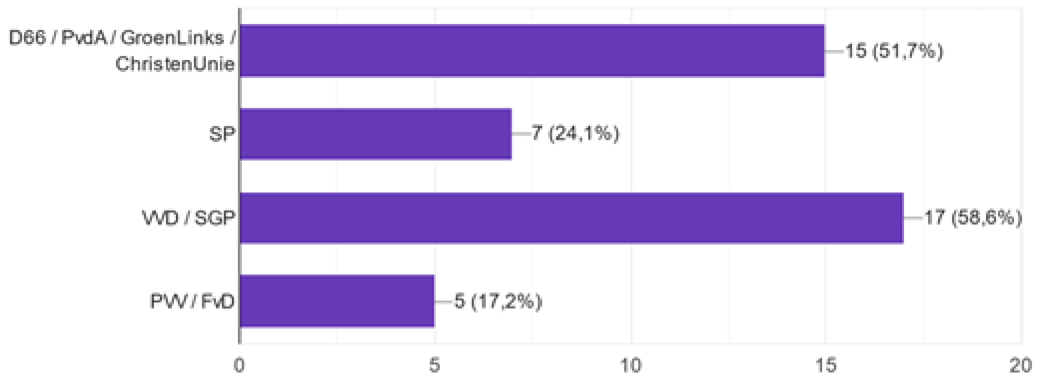

manifest

29 antwoorden

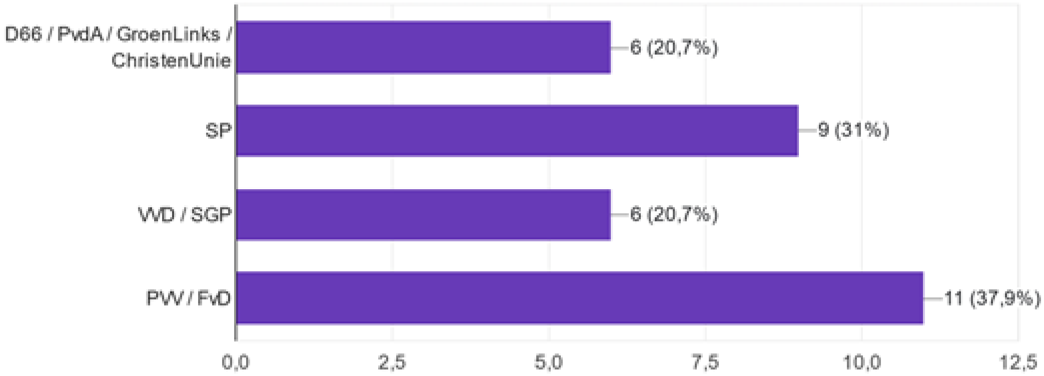

wetsvoorstel

29 antwoorden

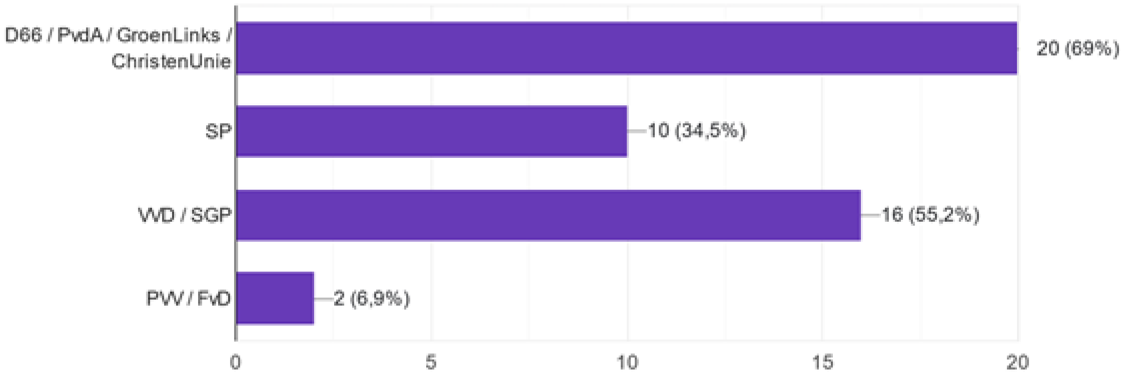

bijvoorbeeld

29 antwoorden

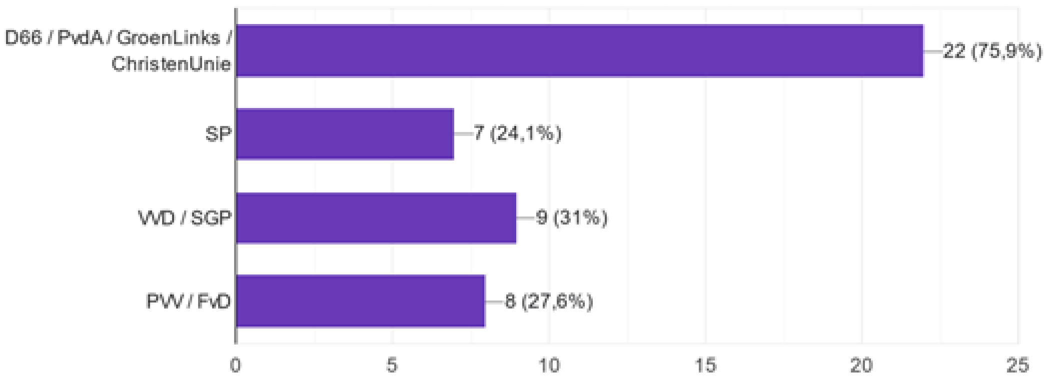

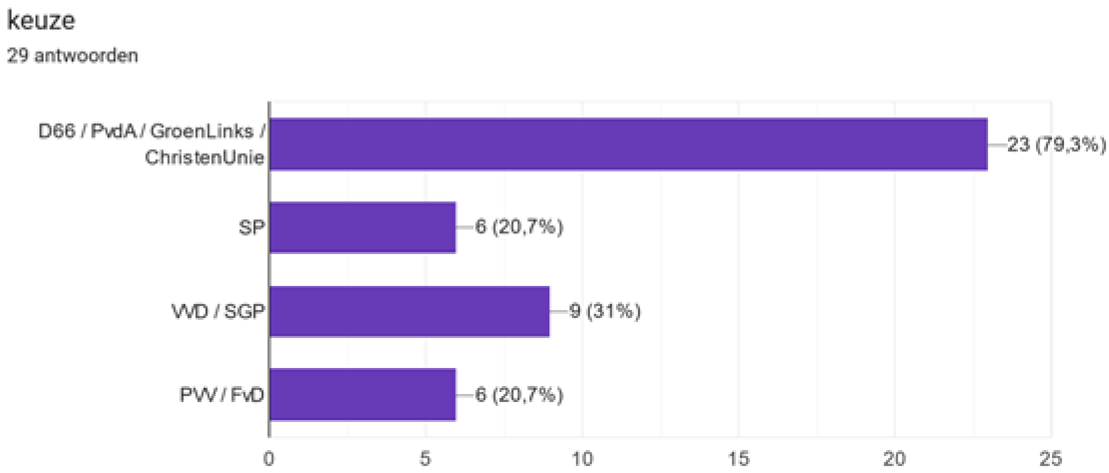
keuze
29 antwoorden
D66 / PvdA / GroenLinks / ChristenUnie
SP
VVD / SGP
PVV / FvD
23 (79,3%)
6 (20,7%)
9 (31%)
6 (20,7%)
0
5
10
15
20
25

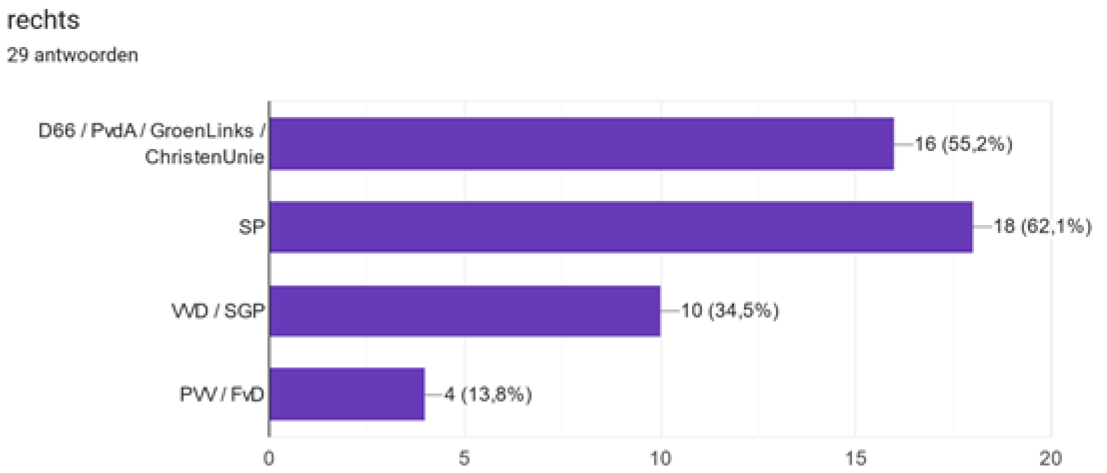
rechts
29 antwoorden
D66 / PvdA / GroenLinks / ChristenUnie
SP
VVD / SGP
PVV / FvD
16 (55,2%)
18 (62,1%)
10 (34,5%)
4 (13,8%)
0
5
10
15
20

Appendix D

Result tables

Class	**Precision**	**Recall**	**F1-score**
Non-populist	0.802	0.855	0.828
Populist	0.792	0.724	0.756
accuracy			0.798
macro avg	0.797	0.789	0.792
weighted avg	0.798	0.798	0.797

Table D.1: TF-IDF with all words as features.

Class	**Precision**	**Recall**	**F1-score**
Non-populist	0.786	0.835	0.810
Populist	0.765	0.703	0.732
accuracy			0.778
macro avg	0.776	0.769	0.771
weighted avg	0.777	0.778	0.776

Table D.2: TF-IDF without function words

Features	**Precision**	**Recall**	**F1-Score**
Content words	0.767	0.761	0.763
(+) articles	0	0	0
(+) prepositions	+0.001	+0.001	+0.001
(+) quantifiers	+0.003	+0.002	+0.003
(+) conjunctions	0	0	0
(+) pronouns	+0.006	+0.006	+0.007
(+) auxiliary verbs	+0.001	0	+0.001
(+) adverbs	+0.004	+0.003	+0.003
(+) modifiers	+0.002	+0.002	+0.002
(+) interjections	+0.001	0	+0.001

Table D.3: Precision, Recall, and F1-Score for TF-IDF, with different function word types added (macro averages). For all of these scores, topic words were removed.

Features	**Precision**	**Recall**	**F1-Score**
Content and function words	0.787	0.779	0.782
(-) articles	+0.001	+0.001	+0.001
(-) prepositions	-0.004	-0.003	-0.004
(-) quantifiers	-0.003	-0.003	-0.003
(-) conjunctions	+0.002	+0.002	+0.002
(-) pronouns	-0.008	-0.008	-0.008
(-) auxiliary verbs	+0.004	+0.004	+0.004
(-) adverbs	-0.002	-0.001	-0.001
(-) modifiers	0	+0.001	+0.001
(-) interjections	+0.001	+0.001	+0.001

Table D.4: Precision, Recall, and F1-Score for TF-IDF, with different function word types added (macro averages). For all of these scores, topic words were removed.

Class	**Precision**	**Recall**	**F1-score**
Non-populist	0.793	0.848	0.820
Populist	0.781	0.710	0.744
accuracy			0.788
macro avg	0.787	0.779	0.782
weighted avg	0.788	0.788	0.787

Table D.5: TF-IDF without topic words

Class	**Precision**	**Recall**	**F1-score**
Non-populist	0.780	0.828	0.803
Populist	0.754	0.695	0.723
accuracy			0.770
macro avg	0.767	0.761	0.763
weighted avg	0.769	0.770	0.769

Table D.6: TF-IDF without function and topic words

Populist	**Non-populist**
waanzin (*madness*)	europarlementariër (*MEP*)
schandalig (*outrageous*)	twijfel (*doubt*)
rechtvaardigheid (*justice*)	thema (*theme/topic*)
verzet (*resistance*)	bijdrage (*contribution*)
tuig (*scum*)	check (*check*)
massaal (*en masse*)	manifest (*manifest*)
volk (*people/nation*)	wetsvoorstel (*bill (law proposal)*)
leugen (*lie*)	bijvoorbeeld (*for example*)
fantastisch (*fantastic*)	keuze (*choice*)
kamerdebat (*debate in parliament*)	rechts (*right (wing)*)

Table D.7: Most important words to the model when classifying populist and non-populist tweets.

	precision	recall	f1-score
Coalition	0.606	0.371	0.460
Opposition	0.823	0.924	0.870
accuracy			0.791
macro avg	0.714	0.647	0.665
weighted avg	0.771	0.791	0.772

Table D.8: Precision, Recall, and F1-Score of word use as a feature when classifying opposition and coalition tweets. Function words and topic words have been removed.

Coalition	**Opposition**
kies (*choose*)	petitie (*petition*)
slag (*starting to work on*[1])	socialer (*more social*)
bezig (*working on*)	publiek (*public(ally)*)
stel (*ask/propose* [2])	verzet (*resistance*)
staatssecretaris (*Secretary of State*)	kamervragen (*questions in parliament*)
veilig (*safe*)	waarheid (*truth*)
vooruit (*forward*)	waanzin (*madness*)
kennis (*knowledge*)	beweging (*movement*)
veiligheid (*safety*)	kamerdebat (*debate in parliament*)
kamerlid (*member of parliament*)	tuig (*scum*)

Table D.9: Most important words to the model when classifying opposition and coalition party tweets.

	precision	recall	f1-score
Left	0.770	0.808	0.788
Right	0.741	0.694	0.717
accuracy			0.758
macro avg	0.755	0.751	0.753
weighted avg	0.757	0.758	0.757

Table D.10: Precision, Recall, and F1-Score of word use as a feature when classifying left- and right-wing party tweets. Function words and topic words have been removed.

Left	**Right**
ongelijkheid (*inequality*)	waanzin (*madness*)
socialer (*more social*)	kamervragen (*questions in parliament*)
rechtvaardigheid (*justice*)	verzet (*resistance*)
twijfel (*doubt*)	kamerdebat (*debate in parliament*)
minimumloon (*minimum wage*)	tuig (*scum*)
actie (*action/activity*)	volk (*people/nation*)
voorstel (*proposal*)	nederlander (*Dutchman*)
aanlezen (*read*[3])	ongelooflijk (*unbelievable*)
investering (*investment*)	zinloos (*pointless*)
check (*check*)	onbetaalbaar (*priceless/too expensive*)

Table D.11: Most important words to the model

	precision	recall	f1-score
Conservative	0.750	0.715	0.732
Progressive	0.757	0.789	0.773
accuracy			0.754
macro avg	0.754	0.752	0.752
weighted avg	0.754	0.754	0.754

Table D.12: Precision, Recall, and F1-Score of word use as a feature when classifying conservative and progressive party tweets. Function words and topic words have been removed.

Progressive	**Conservative**
eerlijker (*more fair*)	kamervragen (*questions in parliament*)
eerlijk (*fair*)	waanzin (*madness*)
socialer (*more social*)	verzet (*resistance*)
ongelijkheid (*inequality*)	volk (*people/nation*)
twijfel (*doubt*)	kamerdebat (*debate in parliament*)
minimumloon (*minimum wage*)	tuig (*scum*)
rechtvaardigheid (*justice*)	nederlander (*Dutchman*)
actie (*action/activity*)	ondernemer (*entrepreneur*)
voorstel (*proposal*)	ongelooflijk (*unbelievable*)
werkdruk (*work pressure*)	zinloos (*pointless*)

Table D.13: Most important words to the model

	precision	recall	f1-score
Non-populist	0.594	0.848	0.699
Populist	0.547	0.241	0.335
accuracy			0.585
macro avg	0.571	0.544	0.517
weighted avg	0.574	0.585	0.541

Table D.14: Precision, Recall, and F1-Score of an SVM-model trained with all readability features

	precision	recall	f1-score
Non-populist	0.595	0.848	0.699
Populist	0.550	0.244	0.337
accuracy			0.586
macro avg	0.572	0.546	0.518
weighted avg	0.575	0.586	0.543

Table D.15: Precision, Recall, and F1-Score of an SVM-model trained with readability features: Leesindex A, Flesch, and Flesch-Douma. w/s is left out.

	precision	recall	f1-score
Non-populist	0.594	0.848	0.699
Populist	0.546	0.240	0.333
accuracy			0.585
macro avg	0.570	0.544	0.516
weighted avg	0.573	0.585	0.540

Table D.16: Precision, Recall, and F1-Score of an SVM-model trained with readability features: w/s, Flesch, and Flesch-Douma. Leesindex A is left out.

	precision	recall	f1-score
Non-populist	0.594	0.848	0.699
Populist	0.546	0.240	0.333
accuracy			0.585
macro avg	0.570	0.544	0.516
weighted avg	0.573	0.585	0.541

Table D.17: Precision, Recall, and F1-Score of an SVM-model trained with readability features: w/s, Leesindex A, and Flesch-Douma. Flesch is left out.

	precision	recall	f1-score
Non-populist	0.595	0.847	0.699
Populist	0.548	0.244	0.337
accuracy			0.586
macro avg	0.571	0.545	0.518
weighted avg	0.574	0.586	0.542

Table D.18: Precision, Recall, and F1-Score of an SVM-model trained with readability features: w/s, Leesindex A, and Flesch. Flesch-Douma is left out.

	Populist		**Non-populist**	
	mean	*SD*	*mean*	*SD*
w/s	14.39	7.83	13.75	7.22
Leesindex A	63.86	23.83	71.26	23.9
Flesch	57.52	25.32	66.11	25.15
Flesch-Douma	70.86	23.06	78.68	22.9

Table D.19: Mean and standard deviation scores for the different readability features.

	Precision	**Recall**	**F1-score**
Coalition	0.000	0.000	0.000
Opposition	0.760	1.000	0.864
accuracy			0.760
macro avg	0.380	0.500	0.432
weighted avg	0.577	0.760	0.656

Table D.20: Precision, Recall, and F1-Score for readability

	Coalition		**Opposition**	
	mean	*SD*	*mean*	*SD*
w/s	11.87	5.72	14.74	7.87
Leesindex A	74.64	23.84	65.87	23.85
Flesch	68.31	26.10	60.41	25.10
Flesch-Douma	80.70	23.77	73.50	22.86

Table D.21: Mean and standard deviation scores for the different readability features.

	Precision	**Recall**	**F1-score**
Left	0.585	0.892	0.706
Right	0.592	0.199	0.298
accuracy			0.586
macro avg	0.588	0.545	0.502
weighted avg	0.588	0.586	0.526

Table D.22: Precision, Recall, and F1-Score for readability

	Left		**Right**	
	mean	*SD*	*mean*	*SD*
w/s	14.23	7.42	13.76	7.60
Leesindex A	69.73	23.68	65.85	24.57
Flesch	64.80	24.63	59.22	26.44
Flesch-Douma	77.49	22.43	72.41	24.08

Table D.23: Mean and standard deviation scores for the different readability features.

	Precision	**Recall**	**F1-score**
Conservative	0.592	0.310	0.407
Progressive	0.570	0.810	0.669
accuracy			0.575
macro avg	0.581	0.560	0.538
weighted avg	0.580	0.575	0.546

Table D.24: Precision, Recall, and F1-Score for readability

	Conservative		**Progressive**	
	mean	*SD*	*mean*	*SD*
w/s	13.88	7.63	14.15	7.38
Leesindex A	66.04	24.67	69.77	23.56
Flesch	59.57	26.58	64.78	24.44
Flesch-Douma	72.73	24.21	77.48	22.26

Table D.25: Mean and standard deviation scores for the different readability features.