A Natural Language Processing based text analysis of populist rhetoric in social media text messages

N.A. (Nizar) Hirzalla
MSc Artificial Intelligence, Vrije Universiteit, Amsterdam

KEYWORDS
Natural Language Processing
text mining
computationally detecting
populist rhetoric
social media

ABSTRACT
The political approach of populism is rising internationally while populist parties are gaining more followers. Populist parties tend to use social media as the primary method of spreading their ideas. Similarities have been detected in how these populist parties exactly spread the message on social media from a linguistic perspective, as these messages always contain a typical rhetoric and usage of words that can be described as populist rhetoric. Using definitions and characteristics of populism and populist rhetoric as coined by language and political experts an analysis based on Natural Language Processing (NLP) text analysis methods is done. The NLP text mining based methods of topic modelling and sentiment analysis are used for this purpose. These methods provide insights into the parties' rhetoric, the used vocabulary, what they deem the most important topics and what entities the parties perceive as good and bad. The Yellow Vest Moment, which has been a debating point regarding if it can be defined as populist movement, is used as an example case to seek if this type of analysis can classify this movement as populist. Two established Dutch populist parties, PVV and Identitair Verzet, are used as well for comparison.

1. Introduction

An international rise of populism has been observed in the last decade (Algan, Guriev, Papaioannou, & Passari, 2017). This is a trend that has been met by disapproval overall due to the negative connotations associated with populism, as well as been dubbed as a significant threat worldwide by several experts (Cox, 2017; Goodwin, 2011). Noticeable among populist parties is that they tend to use social media extensively to spread the core message. As social media is easily accessible, these parties potentially have a large outreach and as such can acquire new followers and spread their messages relatively easy (Engesser, Ernst, Esser, & Büchel, 2017). However, these types of messages have been found problematic as they usually antagonize a significant portion of the overall population (Gidron & Bonikowski, 2013). This can in turn lead to a divided nation or even the collapse of the democratic foundations of a country (Cox, 2017; Gidron & Bonikowski, 2013; Goodwin, 2011).

Rhetorical and linguistic patterns have been detected in how these messages are usually constructed by the populist writer. Such patterns manifest themselves, for instance, by repetitive usage of certain words in different messages or certain types of rhetoric that function repetitively as the core element of the message (Engesser et al., 2017; Jagers & Walgrave, 2007). With these patterns established, this could be used for further purposes such as detecting the rise of upcoming populist parties. This can be done through a computational analysis using different kinds of Natural Language Processing (NLP) techniques, the further implications being that new parties can computationally be detected using populist rhetoric and therefore can be approached as such. To illustrate this, an analysis based on a comparison between two popular and established populist parties in the Netherlands, PVV and Identitair Verzet (IV), and an upcoming party that gained traction recently, The Yellow Vest Movement (YVM), will be made. Geert Wilders' PVV ('Party for Freedom') and IV ('Identitarian Resistance') are both well-known right-wing populist parties that have a significant number of followers.

The YVM mainly combats economic inequality and has been named a populist movement but not officially yet seen as such due to it being new, making it an interesting example case (Hankla, 2018; Nossiter, 2018). The results of the analysis will determine if the YVM can potentially be classified as a populist movement through computational means and substantiate the populist label it received by the media.

However, the goal of this paper is not to define 'populism' or to classify parties as populist but rather to construct a prototype for a universal method to determine whether text messages match the criteria of populist rhetoric.

2. Related work

Literature from two domains are especially relevant to this paper. The first domain constitutes studies related to defining populism and populist rhetoric. The second domain relates to technical studies in which different NLP techniques are used.

2.1 Defining populism and populist rhetoric

There does not seem to be a universal consensus on the definition of populism (Mudde, 2013). However, a definition coined by Cas Mudde in 2004, based on a set of studies, has been dubbed as an influential and iconic definition due to its expressiveness and universal usage (in other studies but also among the general public) (Gidron & Bonikowski, 2013). Mudde defines populism as follows:

> "[populism is] a thin-centered ideology that considers society to be ultimately separated into two homogeneous and antagonistic groups, 'the pure people' versus 'the corrupt elite' and which argues that politics should be an expression of the volonté generalé (general will) of the people. "
>
> (Mudde, 2004,p.543)

Due to its influential nature, this definition will be used for this paper. Some intrinsic properties of populism can also be defined to get a better understanding of what typical populist rhetoric encompasses. 'Typical' patterns that emerge in populist rhetoric, as identified by political linguist Szilagyi (Szilagyi, 2017), are the following and are in accordance with a 2017 study (Engesser et al., 2017):

(i) Constantly referring and advocating to/for 'the people'
(ii) Usage of personal pronouns to connect with the audience, with the intended purpose to create a 'us versus them' mentality
(iii) Calling out and criticizing the elite
(iv) Repeated use of metonymies (using alternative words to describe what is factually meant)
(v) Emphasizing victimhood of the audience

These patterns do not necessarily always indicate whether a person is a populist or not. The major factor here is consistency: if these patterns form the central component of your rhetoric and are used consistently, then the person using this rhetoric can be identified as a populist (Szilagyi, 2017).

2.2 NLP techniques

NLP is a well-established domain within Artificial Intelligence and is concerned with processing and analyzing large amounts of natural language data, usually with the intent to enhance human-computer interaction or to computationally understand 'human language' (Collobert et al., 2011). Text mining is a technique often part of NLP-based analyses, as it can play an essential part in obtaining relevant information from large amounts of text, for example, through deriving patterns (such as the established patterns in populist-driven texts) (Kao, 2007).

Going more in depth, text mining can be used for many intents and purposes once a dataset or a corpus has been established. Such a purpose would be to detect how certain topics are perceived by the examined population, which is also known as a sentiment analysis. There are also NLP-techniques for summarizing purposes. Making a dictionary of a defined number of text posts with the intent to model topics to showcase the most discussed/important topics is one of such techniques (topic modelling).

For this paper, the NLP-techniques of text mining, topic modelling and a sentiment analysis will prove to be useful. These techniques can be performed to obtain relevant data from the text source, find the most important topics and usage of words in that data and how they are perceived sentimentally (Liu, Tang, Dong, Yao, & Zhou, 2016; Maier et al., 2018; Pak & Paroubek, 2010; Westerlund, Leminen, & Rajahonka, 2018). Therefore, this paper will seek if the usage of these techniques combined provides a good universal method for finding populist rhetoric in text messages.

3. Methods

The steps that will be taken for text mining, topic modelling and sentiment analysis will be thoroughly explained in this section. In addition, the steps for obtaining the data as well as pre-processing the data is elaborated upon.

3.1 Text mining

The data used for this paper consists of posts from the official PVV and IV Facebook-pages (approximately 300000 and 1000 followers). Due to YVM not having an official page, the three most active YVM Facebook pages (combined approximately 5000 followers) are instead used (‘Gele-Hesjes-NL’, ‘GeleHesjesVlaanderen’ and ‘DeGeleHesjes’). Upon visual inspection, the YVM based in Flanders (Belgium) did not have any significant differences with those based in the Netherlands and as such can be grouped as one.

For obtaining and text mining posts on Facebook, Netvizz (version 1.6)[1] is used. Netvizz is a text mining tool with multiple functions for analyzing text posts on Facebook (Rieder, 2013). In Netvizz, the 'Page posts' module is chosen as this focuses on retrieving posts. For all three parties, 999 posts are selected for retrieval (the maximum number that can be given as input for obtaining posts). For PVV and IV the last 999 posts were obtained this way. For YVM this netted 274 posts. This lower number of posts can be attributed to the fact that YVM is a recent movement and therefore logically will have a lower number of posts in comparison to the established parties. The obtained text was cleaned up and prepared for topic modelling by doing the following:

(i) All strings that solely contain a picture, video or a shared link are removed, as these cannot be interpreted for text analysis purposes.
(ii) A standard Dutch and English list of stop words[2] is used to filter stop words from the obtained text posts.
(iii) Specific stop words, not in the standard lists, used by the different parties are manually picked out and filtered out as well.
(iv) Single characters, concatenations of punctuation marks, smileys, and other insignificant fragments of text are filtered out as well.
(v) Strings that are empty as a result of the previous filters, are not included.

This produced the results as shown in Table 1. Table 1 shows the number of posts and the average length of a post from the three parties before cleaning up the data and after cleaning up the data.

1 https://tools.digitalmethods.net/netvizz/facebook/netvizz/index.php

2 Stop words are commonly used words that do not tell us much about a topic on their own (Rajaraman & Ullman, 2011). For example, words like: *which*, *here*, *else*, *over* and *the*.

Table 1

Number of text messages and the average length of a text message (in words) respectively before and after cleaning the different datasets.

Political party	Before clean-up	After clean-up
PVV	999 - 22	754 - 8
IV	999 - 26	892 - 13
Gele hesjes	274 - 58	272 - 27

3.2 Topic modelling

The process of topic modelling itself consists of multiple phases, which we will go through one by one in this section.

3.2.1 Lemmatisation

With the typical NLP-task of lemmatisation, all words are derived to their base form. This makes it possible to classify, and thus interpret, different forms of the same word as one. This gives us better results, as different forms of the same word will not be seen as different words. After preparing the texts, lemmatisation is done with Frog (version 0.15)[3], which makes it possible to lemmatise large quantities of text posts quickly. (Van den Bosch, Busser, Canisius, & Daelemans, 2007)

3.2.2 Building a viable corpus

After lemmatisation, a dictionary of the text is made and converted to a corpus based on the bag-of-words model. This model represents texts, which in the case of the used datasets for this paper ranged from a single sentence towards small monologues, as a multiset of words (an unordered collection of words where each word can appear a finite number of times). This multiset of words disregards grammar but instead focuses on the multiplicity and salience of the different words (Zhang, Jin, & Zhou, 2010).

3.2.3 Using latent Dirichlet allocation to model topics.

With Python module Gensim (version 3.7.1)[4] using latent Dirichlet allocation (LDA), topic modelling is performed (Rehurek & Sojka, 2010). LDA is a type of statistical classification that uses a generative model to classify similar data in a corpus (Blei, Ng, & Jordan, 2003). A generative model is based on the joint probability distribution that seeks if different variables can be classified in fitting any particular range or a set of values specified for that variable. In the case of LDA, it looks for connections between different words and seeks to explain why some parts of the data is similar by discriminating. It operates by viewing the input data as a mixture of various topics and whether these topics can be characterized by the words in the dataset. In the context of this paper, this means that it computes

3 https://languagemachines.github.io/frog/

4 https://pypi.org/project/gensim/

probabilities for topics that come forward prominently based on how certain words seem to relate to each other, which can then be linked to populist rhetoric. The purpose of this is to find the most important talking points of the party that published the texts.

To illustrate how LDA would work in practice, an example will be given. For this example the topic threshold will be set at 5. The case for this example will be a political party that frequently talks about immigration, climate change, leaving the EU, decreasing taxes and education on social media. LDA's generative model statistically classifies that terms related to these topics are consistently used in the same context and same messages, indicating a cluster.

The more the same terms occur in the same context in different text posts about the topics, the stronger the cluster will be and the stronger the probability will be of the most occurring terms in that cluster. Eventually this will produce an output of the mentioned five topics with the most occurring terms that define these topics.

3.2.4 Computing and visualizing the most relevant and salient terms

Using Python module pyLDAvis (version 2.1.2) (Mabey, 2018), the most relevant words and most salient terms within the topics are computed. pyLDAvis extracts information from a fitted LDA topic to create an interactive web-based visualization for easy interpretation of the topic modelling results. For calculating the relevancy of words within topics and saliency, in combination with the LDA results, the following formulas from a 2012 study (Chuang, Manning, & Heer, 2012) and a 2014 study (Sievert & Shirley, 2014), respectively, are used. These formulas are the following:

(i) *saliency(termw) = frequency(w) * [sumtp(t|w) * log(p(t|w)/p(t))]*

(Chuang et al., 2012)

(ii) *relevance(termw|topict) = λ * p(w|t) + (1 − λ) * p(w|t)/p(w)*

(Sievert & Shirley, 2014)

In these formulas *w* and *t* respectively stand for 'word' and 'topic'. Saliency is a measure for indicating how noticeable and important a term is, and in this case indicates how much it tells about the topic. Relevancy in this case is a weighted average of the probability of the word in the context of the given topic and the probability of the word in the context of the given topic normalized by the probability of the topic.

Furthermore, pyLDAvis shows how different topics from the LDA results are linked and as such whether there are similarities between the topics. The Python script for the procedures of (1) pre-processing text, (2) LDA topic modelling, and (3) visualizing the models has been published on Github[5].

5 https://github.com/NizarH/PopulismTopicModelling

3.3 Sentiment analysis

With sentiment estimating tool SentiStrength (version 2.3)[6], the sentiment analysis is performed. SentiStrength rates input data, in the form of sentences or small texts, on a scale from -4 (negative) to 4 (positive) using pre-defined data as a reference point (Thelwall, Buckley, Paltoglou, Cai, & Kappas, 2010), in this case a Dutch dataset of words[7]. The main list of words used for reference is a standard list of all Dutch words correlated with positive and negative sentiment. To get more accurate results, other lists are also included: 'booster words', negating words, slang words, question words, English words that get used often in Dutch texts, emoticons and derision terms. Booster words consist out of words that amplify a certain word, for instance the word 'very'. For example, in the case of 'that is very good' the positive score of *good* increases. Negating words do the opposite; if *good*, is preluded with 'not', the positive score becomes 0. Derision terms reduce negative scores slightly by looking at certain irony-associated words like *haha* or *lol*. With this reference data SentiStrength is applied on all of the cleaned datasets.

4. Results

For topic modelling, a threshold of five topics is chosen as this gives the best results. Fewer topics would cluster some terms too broadly, while more topics would return terms with too loose of a connection to the topic

4.1 Topic Modelling Results

The computed topics seem to range from promoting voting for the PVV in elections (topic 1, 5) to reiterating typical PVV talking points (topic 3, 4) (see Appendix A[8]). How these topics are related to each other can be seen in Figure 1.

The most salient terms show that there are some terms with a relatively high frequency: 'Nederland', 'PVV' and 'StemPVV' (translates to: 'vote for PVV') . The most salient terms seem to overlap with the most relevant terms in the five classified topics (see Appendix B).

Figure 1 shows no overlap between topics. Topics 1, 2, 3 and 5 are relatively close to each other (in the same hemisphere) while topic 4 is on the complete other side, indicating a more distinct topic in comparison to the other topics.

For IV, the topic modelling results show many similarities with PVV. There is no overlap between the IV topics and the spread of the topics is the same as with the PVV topics: topic 1, 2, 3, 5 are in the same hemisphere and close to each other while topic 4 is far-off, as can be seen in Figure 2.

6 http://sentistrength.wlv.ac.uk/index.htmlDownload

7 https://github.com/NizarH/PopulismTopicModelling/tree/master/SentiStrengthDutch/dutch

8 Find the online version of this article, including appendices, at linguujournal.nl.

Figure 1. Intertopic distance map. A representation of the relations and connectivity between the different LDA produced topics of the PVV dataset.

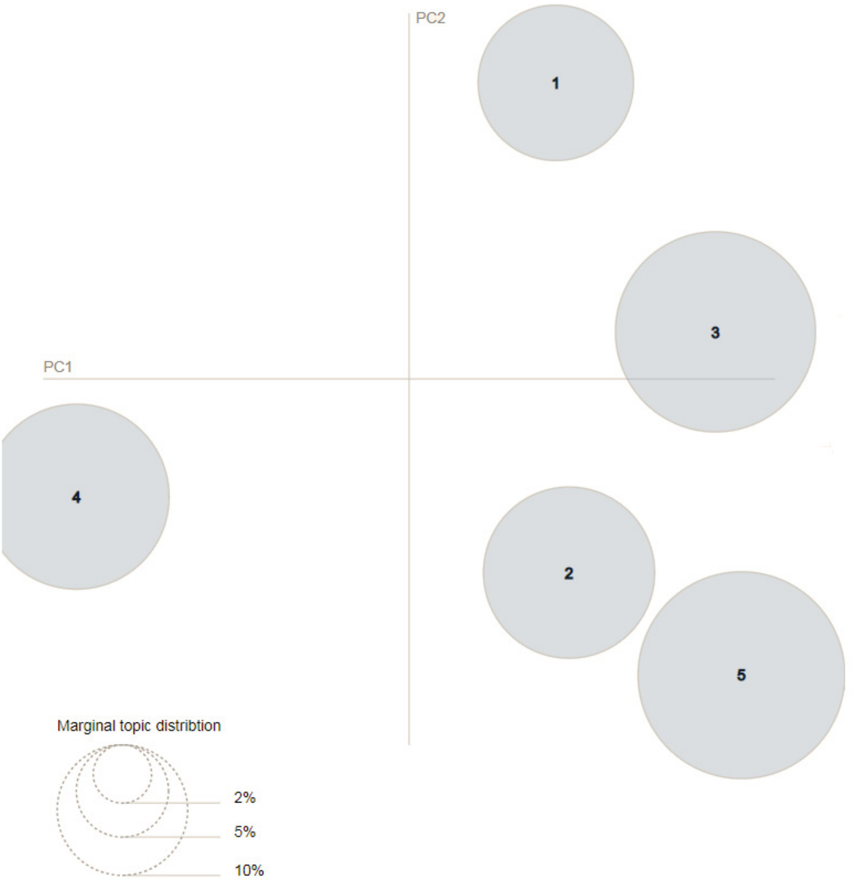

Figure 2. A representation of the relations and connectivity between the different LDA produced topics of the IV dataset.

As IV is an activist group, unlike political party PVV, there are no topics asking for votes. The topics all indicate to be centered thematically around political talking points. Terms such as ‘activist’ seem to indicate that some of these topics may also be centered around promoting activism (see Appendix A). For IV, the salient terms seem to have a higher frequency than the PVV salient terms, which can be explained due to a broader range of talking points (see Appendix B). There are some terms that stand out due to their very significant presence, such as 'Nederland', 'land', 'nieuw', 'groot' and 'verzet'. These terms translate, respectively, to 'the Netherlands', 'country', 'new', 'big' and 'resistance'. For the YVM, thematically, the topics all seem to promote some form of activism, but are also addressing the readers and reiterating their role as 'the people' (see Appendix A). The relations between these topics can be seen in Figure 3.

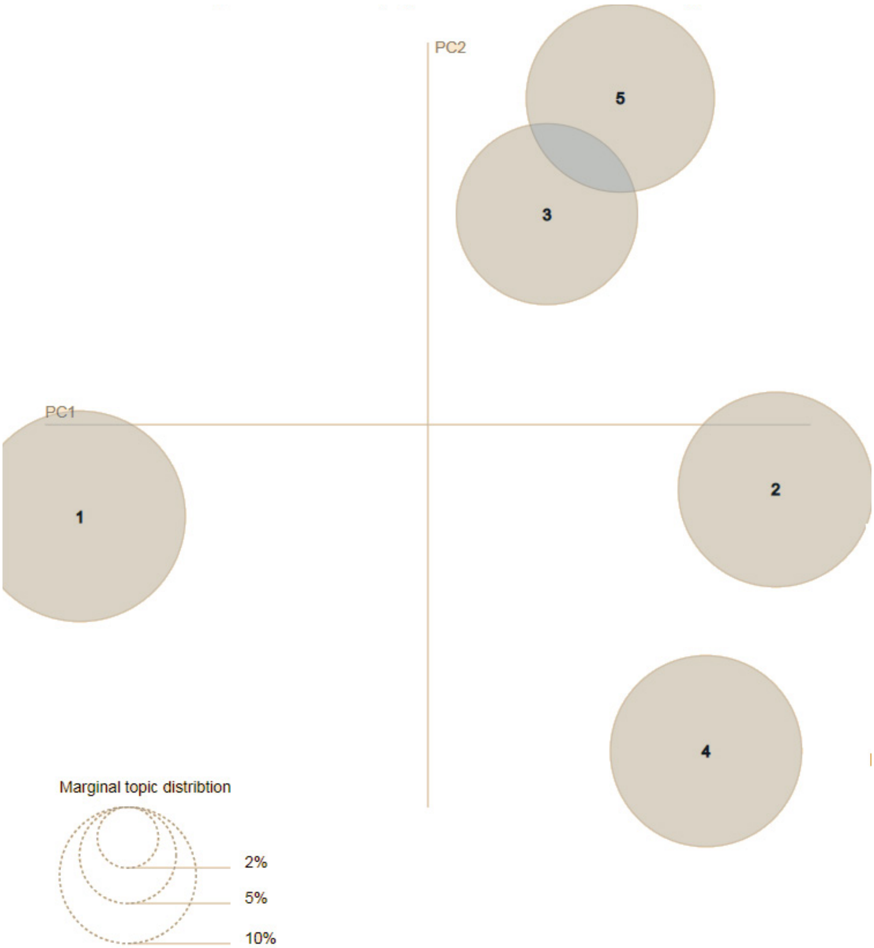

Figure 3. A representation of the relations and connectivity between the different LDA produced topics of the YVM dataset.

Figure 3 shows an overlap between topic 3 and 5. The structure of having four topics in the same hemisphere (relatively close to each other) with a distinct topic, as was also observed with the other two parties, seems to be returning with the YVM, this time topic 1 being the far-off, distinct topic. We can also see that most of the salient topics have a relatively low frequency on average, but there are also very noticeable high frequency salient terms such as 'mens', 'volk', 'geld', 'bank', 'alleen' and 'denken' (see Appendix B). These terms translate, respectively, to 'person', 'people', 'money', 'bank', 'alone', 'to think'.

4.2 Sentiment Analysis

For the sentiment analysis, terms that had the most positive scores and the terms with the most negative scores were considered and analysed. These were terms with a score ranging from 3 to 4 and -3 to -4, respectively. Scores from -1 to 1 were considered neutral and scores -2 or 2 were considered too marginal to give a good indication of how different terms are perceived. The strongly rated posts were also selected because they are the best interpretable and therefore give us the most insight. The results of the sentiment analysis for all three parties can be seen in Table 2.

Table 2 shows some similarities between the parties and how they perceive different topics. Terms such as 'Rutte', 'de overheid' ('the government'), 'kabinet' ('cabinet') are being perceived as a negative associated terms, for all three parties. Also when looking at positive associated terms there seems to be overlap with terms such as 'volk', 'wij', 'ons' (translates to 'the people', 'we' and 'us'). The PVV and IV seem to practically have the same perceptions thematically and content-wise, considering they share many similarities in both positive and negative associated terms.

5. Discussion

The results of topic modelling provide insight into the rhetoric used by the parties as well as their most discussed topics. These range from promoting activism/voting towards talking about topics that concern the parties. It becomes clear that social media is used for multiple purposes: (1) to promote (physical) activity in support of the party or the party’s ideology, (2) to inform the readers of topics that concern the party, and (3) to form a narrative in which an (emotional) reaction tends to be instigated from the readers.

Some of the produced topics were clear in what they exactly constitute, e.g. PVV topic 1, 4, IV topic 1, 5, and YVM topic 1 and 4. For example, a topic referencing to the ’good people’ and their supposed victimhood can be found in topic 4 of the YVM topics. Sometimes it was not clear what the topic exactly constitutes because the clustering of terms is too broadly, e.g. IV topic 2, 4 and YVM topic 3, 5. The intertopic distance charts showed us that the topics do in fact have much variance indicating a multitude of concerns for the parties. There was hardly any overlap with the different topics, except for one case in the YVM dataset. Some topics were clustered far away from others (topic 4 for PVV and IV, topic 1 of YVM), this is most likely caused due to the fact that these topics include combination of terms that differ significantly from the other four topics per the Intertopic Distance metric.

The sentiment analysis gained insights in populist rhetoric and the different parties’ perceptions on different topics. It helped providing some needed context for the topic modelling results. The sentiment analysis showcases a perceived ‘bad’ common foe for all three parties due to the sheer amount it is talked about and forms the backbone of the parties‘ rhetoric. This is also illustrated by a clear dichotomy in what is consistently perceived as negative versus what is consistently perceived as positive.

Table 2

The results of the sentiment analysis for the (1) PVV, (2) IV, and (3) The YVM (Gele Hesjes), showcasing the most positive and negative associated terms

Party	The most positive associated terms (at least a score of 3 or more)	The most negative associated terms (at least a score of -3 or less)
PVV	*Nederland*, *Nederlandse cultuur*, *Burgers* or *Volk*, *Ons* or *Wij*, *PVV* (PVV members but also people who vote PVV) Also a combination of people, groups or parties who speak out against the negative associated terms	*Rutte*, *VVD*, *Kabinet + Parlement*, *EU*, *Islam* (and also associations with Islamic beliefs, figures, culture or countries), *Terreuraanslag/ Terreur* in combination with Islam, *Geweld* or *Misdrijf* in combination with Islam/ muslims, morroccans, or migrants, *Gelukszoekers* or *Asielzoekers* (usually to refer to refugees or migrants)
Identitair Verzet	*Nederland*, *Nederlanders*, *Ons*, *Wij* or *Onze x* where *x* pertains to Dutch people, institutions, culture or achievements, people associated with Identitair Verzet (members, activists, protesters, supporters), *Afrikaners*	*Rutte*, *VVD*, Dutch left-wing parties, *Extreemlinks*, *EU*, *Europa*, *Islam* (and also associations with Islamic beliefs, figures, culture or countries), *Salafisten* or *Jihadisten*, *Cultuurmarxisme*, *Asielzoekers* or *Illegalen*, *Migranten*
Gele Hesjes	*Volk*, *Wij*, *Gele Hesjes*, *De normale burger* or *mensen*, *De arme burger* (and terms showcasing/emphasizing victimhood of 'the people'), terms associated with standing up for the 'victimized people', terms associated with improving the climate or *milieu* with collective behaviour change, terms associated with protesting or rebelling against the *rulers* or *elite*, terms associated with economic equality	*Rutte*, *De overheid*, *Banken*, *Verdeeldheid* or terms pertaining to economic inequality, *System*, *Europa*, *Belasting*, terms pertaining to making profit by 'the elite' at the expense of 'the people', *Klimaatmaatregelen* or terms pertaining to combat climat change with increasing ecological costs for 'the people'

For PVV and IV the negative perceived entities can be defined into two groups: The elite or 'the corrupt rulers' (VVD, Rutte, the government, EU) on one hand and an existential threat on the other hand (Islam, migrants). For the YVM, the same division can be found, with the former being Rutte, the government and Europe and the latter banks and economic inequality. They see these threats operating in real time and have a sense of being collectively oppressed and economically exploited by both the banks as well as the elite (Marijnissen, 2018).

The PVV and IV parties seem to be more aligned with nationalist and right-wing sentiments, while the YVM aligns more with socialist and left-wing sentiments (Otjes & Louwerse, 2013; van Elsas, Hakhverdian, & van der Brug, 2016). Content and main concerns of the parties aside, it becomes clear that all parties have the earlier defined 'us versus them' dichotomy forming their rhetoric. For instance, all parties constantly refer to 'the people' or other synonyms with the same effect. Additionally, emphasizing victimhood of 'the people' and the parties themselves happens constantly, most apparent in the YVM dataset. Often, personal pronouns such as 'we' or 'us' are used in order to connect with the audience. Furthermore, the usage of metonymies, such as PVV referring to all migrants and refugees as 'fortune seekers' is also a recurring theme. Lastly, calling out and criticising the elite seems to be the most used rhetoric as this is constantly done and overall seems to form the biggest priority for the parties. These outcomes combined seem to give a full match with the defined characteristics of populist rhetoric in the Related Work section, for all three parties.

5.1 Limitations of this paper and the used methods

While performing the different NLP-analyses, some issues and limitations became apparent. The first one being that the YVM dataset is significantly smaller than the PVV and IV dataset, possibly yielding less representative results. This was partially dealt with by using multiple Facebook pages to gain a larger YVM dataset.

From the topic modelling results, it was not always clear how to interpret the topics. Possibly, using larger datasets for all parties could cluster some terms with more significance which then could be interpreted more easily. Doing a sentiment analysis did in turn provide some needed context for the topic modelling results.

As for the sentiment analysis, there is a limitation in how the text messages are interpreted. The perception of what can be regarded as negative or positive has a subjective connotation. For instance, people with different characters can perceive the negativity or positively of a certain text messages differently, and as such the used simple linear scoring might not reflect reality adequately.

Additionally, SentiStrength has a hard time detecting sarcasm, potentially making some of the produced scores inaccurate. The mentioned 'derision terms' list of words as referenced in Section 3.3 can act as a simplified workaround for this, but this in itself poses problems. Some of these words might instead not be used for the purpose of irony, but then still get counted as such this is hardcoded into

SentiStrength. Properly detecting sarcasm or irony cannot be done at this moment, as detecting sarcasm is still a topic of research with difficulties on its own in NLP-research (González-Ibáñez, Muresan, & Wacholder, 2011). To somewhat counteract this, a randomly selected group of posts from all parties were manually inspected for sarcasm, although this way sarcasm was only sporadically detected. Furthermore, SentiStrength has difficulties seeing 'the bigger picture', for instance when a text post consists of multiple sentences with half of these sentences being positive and the other half being negative. This would net an overall neutral score, but this does not represent the content of the text post in a useful way for the purpose of this paper. The latter issue was dealt with by manually checking some of the longer text posts that received a neutral rating.

Furthermore, SentiStrength cannot detect specific jargon or 'invented' words of the populist parties, such as 'kansenparel' as used by IV, which is used to describe non-western migrants or refugees in a derogatory way. These types of words have to be added manually in the reference data, which is less efficient.

6. Conclusion and future work

Multiple NLP techniques have provided new insights in the YVM's use of rhetoric. This paper has shown that YVM adheres to the definition coined by Cas Mudde, engulfing what populism exactly constitutes (Mudde, 2004). It becomes clear from the sentiment analysis in combination with the topic modelling results that the 'us vs them' mentality where 'us' is defined as pure, good, victimized people and 'them' the corrupt, bad elite and/or an existential threat is very present in YVM's rhetoric. This dichotomy becomes especially clear when looking at Table 2, which illustrates the clear division between the positively perceived 'us' and related concepts and the negatively perceived 'them' and related concepts. Next to this, it also became clear that the YVM linguistic style of writing and rhetoric also consistently adheres to all five of the populist characteristics, as defined by language and political experts (Engesser et al., 2017; Szilagyi, 2017).

The same, on both fronts, was also true for the established populist parties, PVV and IV, as was expected. In comparison with these parties the YVM tends to be the same regarding using populist rhetoric and linguistic writing style, although there was difference found in the rhetoric content-wise. The YVM leans much more to left wing populism and socialism while PVV and IV adhere more towards right wing populism and nationalism.

In the greater context, this paper has also provided a prototype for a method to sufficiently find if new groups and/or parties can be computationally classified as populist. In our experiments, we created a corpus and applied automatic text analysis methods in the form of topic modelling and sentiment analysis to obtain insights into the linguistic characteristics of populist rhetoric.

For the future, this method could be improved by including a technique that accurately can detect sarcasm in text consistently. Sarcasm can twist the meaning of a text post significantly and as such should be considered when doing a sentiment analysis. Therefore, this paper recommends further research in the field of NLP regarding detecting sarcasm. Also, for future research this paper recommends to apply this method to more and bigger datasets. This can be applied to other upcoming parties to see if it can consistently classify parties as populist accurately. To add on to that, for future research this paper also recommends to seek if the method in this paper can also be accurately used for other languages (non-Dutch), other types of populist parties and for parties that do not post on social media but rather use more traditional ways of spreading the message. Finally, this paper also recommends to consider non-populist parties and to seek whether the method can confirm them as such.

Acknowledgement
I want to thank my supervisors Marieke van Erp and Victor de Boer for providing the means to carry out the research, their assistance and supervision. I also want to thank political experts Ernst van den Hemel and Sakina Loukili for their advice and feedback. ■

Find the online version of this article, including appendices, at linguujournal.nl.

Received April 2019; accepted August 2019.

References

Algan, Y., Guriev, S., Papaioannou, E., & Passari, E. (2017, November). *The European Trust Crisis and the Rise of Populism* (tech. rep. No. 12444). C.E.P.R. Discussion Papers. Retrieved from https://ideas.repec.org/p/cpr/ceprdp/12444.html

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet Allocation. J. *Mach.* Learn. Res. 3, 993–1022. Retrieved from http://dl.acm.org/citation.cfm?id=944919.944937

Chuang, J., Manning, C. D., & Heer, J. (2012). Termite: Visualization Techniques for Assessing Textual Topic Models. In *Advanced visual interfaces*. Retrieved from http://vis.stanford.edu/papers/termite

Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., & Kuksa, P. (2011). Natural language processing (almost) from scratch. *Journal of machine learning research, 12*(Aug), 2493–2537.

Cox, M. (2017, January). *The Rise of Populism and the Crisis of Globalisation: Brexit, Trump and Beyond*.

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

Gidron, N. & Bonikowski, B. (2013). Varieties of Populism: Literature Review and Research Agenda. Weatherhead Working Paper Series, No. 13-0004. Weatherhead Working Paper Series, No. 13-0004.

González-Ibáñez, R. I., Muresan, S., & Wacholder, N. (2011). Identifying Sarcasm in Twitter: A Closer Look. In *Acl*.

Goodwin, M. J. (2011, September). *Right response: Understanding and countering populist extremism*. Chatham House. London: The Royal Institute of International Affairs, Chatham House. Retrieved from https://kar.kent.ac.uk/54401/

Hankla, C. (2018, December). What French populists from the '50s can teach us about the 'yellow vests' roiling Paris today. Retrieved from http://archive.is/jBikW

Jagers, J. & Walgrave, S. (2007, May). Populism as political communication style: An empirical study of political parties' discourse in Belgium. *European Journal of Political Research, 46*(3), 319–345.

Kao, S., A; Poteet. (2007). *Natural Language Processing and Text Mining*. Springer International Publishing.

Liu, L., Tang, L., Dong, W., Yao, S., & Zhou, W. (2016, September). An overview of topic modeling and its current applications in bioinformatics. *SpringerPlus, 5*(1), 1608.

Mabey, B. (2018). pyLDAvis 2.1.2. Retrieved from https://pyldavis.readthedocs.io/en/latest/

Maier, D., Waldherr, A., Miltner, P., Wiedemann, G., Niekler, A., Keinert, A., Adam, S. (2018). Applying LDA Topic Modeling in Communication Research: Toward a Valid and Reliable Methodology. *Communication Methods and Measures, 12*(2-3), 93–118.

Marijnissen, H. (2018, December). De gele hesjes zijn het zat! Maar wat? Retrieved from https://www.trouw.nl/home/de-gele-hesjes-zijn-het-zat-maar-wat-~a184703a/

Mudde, C. (2004, September). The Populist Zeitgeist. *Government and Opposition, 39*(4), 541–563.

Mudde, C. (2013). Are populists friends or foes of constitutionalism? The Foundation for Law, Justice and Society.

Nossiter, A. (2018, December). How France's 'Yellow Vests' Differ From Populist Movements Elsewhere. Paris. Retrieved from https://www.nytimes.com/2018/12/05/world/europe/yellow-vests-france.html

Otjes, S. & Louwerse, T. (2013, November). Populists in Parliament: Comparing Left-Wing and Right-Wing Populism in the Netherlands. *Political Studies, 63*(1), 60–79.

Pak, A. & Paroubek, P. (2010). Twitter as a corpus for sentiment analysis and opinion mining. In *Lrec* (Vol. 10, 2010, pp. 1320–1326).

Rajaraman, A. & Ullman, J. D. (2011). Data Mining. In *Mining of massive datasets* (pp. 1–17). Cambridge University Press.

Rehurek, R. & Sojka, P. (2010). Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the lrec 2010 workshop on new challenges for nlp frameworks* (pp. 45–50). Valletta, Malta: ELRA.

Rieder, B. (2013). Studying Facebook via Data Extraction: The Netvizz Application. In *Proceedings of the 5th annual acm web science conference* (pp. 346–355). WebSci '13. New York, NY, USA: ACM.

Sievert, C. & Shirley, K. (2014). LDAvis: A method for visualizing and interpreting topics. In *Proceedings of the workshop on interactive language learning, visualization, and interfaces* (pp. 63–70). Association for Computational Linguistics.

Szilagyi, A. (2017). How to talk like a populist. VICE News. Retrieved from https://news.vice.com/en_us/article/gyd98y/how-to-talk-like-a-populist

Thelwall, M., Buckley, K., Paltoglou, G., Cai, D., & Kappas, A. (2010, December). Sentiment Strength Detection in Short Informal Text. *J. Am. Soc. Inf. Sci. Technol. 61*(12), 2544–2558.

van den Bosch, A., Busser, B., Canisius, S., & Daelemans, W. (2007, January). *An efficient memory-based morphosyntactic tagger and parser for Dutch.*

van Elsas, E. J., Hakhverdian, A., & van der Brug, W. (2016, November). United against a common foe? The nature and origins of Euroscepticism among left-wing and right-wing citizens. *West European Politics, 39*(6), 1181–1204.

Westerlund, M., Leminen, S., & Rajahonka, M. (2018). A Topic Modelling Analysis of Living Labs Research. *Technology Innovation Management Review, 8*, 40–51.

Zhang, Y., Jin, R., & Zhou, Z.-H. (2010, December). *Understanding bag-of-words model: A statistical framework.*