

PopAut: An Annotated Corpus for Populism Detection in Austrian News Comments

Ahmadou Wagne, Julia Neidhardt, Thomas E. Kolb

Christian Doppler Laboratory for Recommender Systems, Department of Informatics, TU Wien
Favoritenstraße 9-11, 1040 Vienna, Austria
{ahmadou.wagne, julia.neidhardt, thomas.kolb}@tuwien.ac.at

Abstract

Populism is a phenomenon that is noticeably present in the political landscape of various countries over the past decades. While populism expressed by politicians has been thoroughly examined in the literature, populism expressed by citizens is still underresearched, especially when it comes to its automated detection in text. This work presents the PopAut corpus, which is the first annotated corpus of news comments for populism in the German language. It features 1,200 comments collected between 2019-2021 that are annotated for populist motives anti-elitism, people-centrism and people-sovereignty. Following the definition of Cas Mudde, populism is seen as a thin ideology. This work shows that annotators reach a high agreement when labeling news comments for these motives. The data set is collected to serve as the basis for automated populism detection using machine-learning methods. By using transformer-based models, we can outperform existing dictionaries tailored for automated populism detection in German social media content. Therefore our work provides a rich resource for future work on the classification of populist user comments in the German language.

Keywords: populism detection, machine-learning, language resources

1. Introduction

Populism, a complex and contested concept, has become a significant phenomenon in various political contexts worldwide. Political parties sharing populist views, such as the Austrian FPÖ, German AfD, Argentinian PJ, or French RN, have garnered substantial attention and electoral success, emphasizing the need for a common definition of populism that transcends geographical and political boundaries. While different definitions, including *ideational* (Mudde, 2004; Wirth et al., 2016; Mudde and Rovira Kaltwasser, 2017), *political-strategic* (Barr, 2009; Jansen, 2011; Laclau, 2005) or *stylistic* (Krämer, 2014; Canovan, 1999; Bracciale and Martella, 2017), exist, a core set of motives underlies the essence of populism, with Mudde’s definition of a “thin ideology” serving as a fundamental reference point (Mudde, 2004).

Prior research in this area has mainly focused on analyzing populist content expressed by politicians or political parties (Jagers and Walgrave, 2007; Rooduijn and Pauwels, 2011; Hawkins et al., 2019), often manually or through dictionary-based methods. Few attempts have employed machine learning (ML) algorithms and these have also focused on manifestos and speeches (Ulinskaite and Pukelis, 2021; Hawkins and Silva, 2016; Di Cocco and Monechi, 2022; Klamm et al., 2023). With the rise of social media, there is a growing recognition that citizens, particularly on these platforms, express populist ideas through their messages, thereby reproducing ideas originating from politicians and the media.

However, the analysis of user comments has received limited attention. A notable ML-based study of user comments by Cabot et al. focused on a “Us vs. Them” mentality involving different out-groups (Cabot et al., 2021), a motive not considered characteristic in most works.

Therefore, a lack of annotated data for the training of machine learning models for detecting populism in the German language is evident. One of the only validated and promising solutions for large-scale content analysis is a dictionary built by Gründl (Gründl, 2022) that was tailored to detect populism in the social media content of German politicians. While this approach gave highly valid results, it is limited in flexibility, as it is optimized for high precision and suffers from general problems with dictionary-based approaches, as it relies on certain words to be present, and can not deal with negation or incorporate context. Additionally, Klamm et al. provided a German and English cross-lingual dataset for investigating the mentioning of “the elites” or “the people” in parliamentary debates using ML (Klamm et al., 2023), but not populist motives specifically.

This work wants to address this observed gap and provides the first corpus of news comments annotated for populist rhetoric called PopAut. The corpus should serve as a resource for ML models to enable large-scale analysis of the expression of populism by citizens in the online sphere. The code for this project is available on Github¹ includ-

¹<https://github.com/ahmadouw/COV-Populism-Standard>

ing a reference to the data set. In experiments, we demonstrate the performance of transformer-based methods for text classification compared to existing dictionary-based methods and other ML baselines. The PopAut corpus was initially collected to measure the populist tendencies expressed under news articles and investigate the relation between populist rhetoric and the topic of COVID-19 in the news, which is done in a follow-up work. Therefore it consists of news comments under articles of Austrian daily newspaper *Der Standard*² from 2019 (as a reference time frame before the pandemic) until November 2021.

In the following, we introduce the related work including definitions of populism and available resources. Section 3 describes the creation of the PopAut corpus in detail. In Section 4 the classification experiments are presented, while Sections 5 and 6 discuss the findings and reflect on limitations and potentials for future work.

2. Related Work

To treat populism as a classification problem in the context of ML, it is important to find a definition that is both widely agreed on and operable to be identified at the text level. Other than “thick” political ideologies such as socialism, populism primarily focuses on a narrow set of issues and simplistic solutions. In contrast, “thick” ideologies are comprehensive political beliefs that address a wide range of social issues and have well-developed theories and programs. On the other hand, populism can coexist with and borrow elements from other ideologies (Kaltwasser, 2012) and is not specific to any kind of political direction.

In the early 2000s we saw a rise in populism studies that did not reach beyond observing it on a national scale (Rooduijn, 2019), while later many scholars aimed for a more comparative research of populism (Mudde, 2004; Roberts, 2006; Weyland, 2001). In the following years, there was a growing consensus about the essentials of populism and most of the scholars agree on populism including anti-elitism and people-centrism (Rooduijn, 2019).

Some researchers see populism as a *political strategy* expressed through discourse to mobilize the masses (Laclau, 2005; Abts and Rummens, 2007; Weyland, 2001). Another stream defines it as a *communication style* (Krämer, 2014; Canovan, 1999; Bracciale and Martella, 2017; Moffitt, 2016). Here specific stylistic features are assigned to populist messages that include colloquial and emotional language, toxicity/incivility or simplification (Bracciale and Martella, 2017) or dramatization, polarization, moralization, etc. (Bos et al., 2011).

Mudde (Mudde, 2017) raises critique to this perspective, as he states that not every populist message includes those stylistic elements and they are not sufficiently distinct from other forms of discourse.

However, when observing populism on the text level, we want to know what is being said. In cases where this is of interest, the most widely used definition is based on the definition by Mudde, who defines it as:

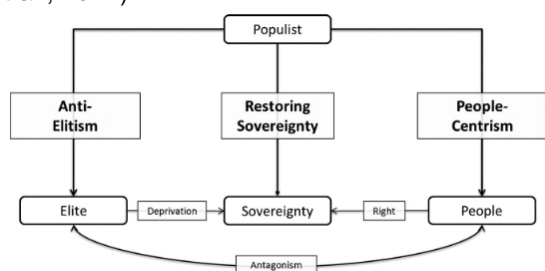
“An ideology that considers society to be ultimately separated into two homogeneous and antagonistic groups, ‘the pure people’ versus ‘the corrupt elite’, and which argues that politics should be an expression of the volonté générale (general will) of the people.” (Mudde, 2004)

This minimal definition sums up the characteristics that are shared by all populist actors, disregarding their political beliefs. The core notion of this ideology is that in the current system, the people lack power and feel the need to be represented more in shaping society, laws, etc. This representation is suppressed by a certain group of elites or by representatives of the elites. Elites hereby can be political elites, financial elites, cultural elites or media elites (Rooduijn, 2019). The people are seen as a homogeneous group that shares a variety of good common attributes, opinions or beliefs, which are praised by populists. Some definitions propose more out-groups than elites, like immigrants (Jagers and Walgrave, 2007; Cabot et al., 2021), to form a “thick” populism, but this often leads to fluent borders to ideologies like nativism. A populist puts themselves into the position of speaking for the people and expressing their needs that are denied by elites. This is a very simplistic view of societies and denies pluralism. However, this view can also empower people who are actually not represented by political elites in an unjust way and serve to correct flaws in democratic systems (Kaltwasser, 2012) but is often seen as a threat to liberal democracy (Canovan, 1999; Mudde and Cristóbal, 2012; Mazzoleni, 2008).

From Mudde’s definition, we can derive three motives that build the main framework for the detection of populism in this work. In Figure 1, we can see a conceptual model based on this definition, which includes anti-elitism, restoring sovereignty (called people-sovereignty in this work) and people-centrism and the interplay of the different actors. Wirth et al. further operationalized the motives and described their features broken down into key messages, which are tied to an underlying ideology, to facilitate a more fine-grained distinction (Wirth et al., 2016). These key messages are detectable in texts and are therefore considered an operable

²<https://www.derstandard.at/>

Figure 1: Conceptual model of thin populism (Ernst et al., 2017)



framework for this work. These three dimensions of populism are moreover used in various other works on the matter (Schulz et al., 2018; Ernst et al., 2019; Wirz, 2018; Rooduijn, 2014).

As we want to observe populism in news comments, we focus on the citizen as a populist actor. Most of the previous research that deals with analyzing populism in texts, focused on politicians or political parties as populist actors (e.g. (Hawkins et al., 2019; Ernst et al., 2019; Rooduijn and Pauwels, 2011)). However recent work has dealt with people expressing populism, especially in messages posted online (Engesser et al., 2017a; Krämer, 2017). In the context of news, Esser et al. describe “populist citizen journalism” (Esser et al., 2017) and criticize that some, although rather neutral themselves, news outlets open space for reader’s populist comments and Hameleers et al. (Hameleers et al., 2016) claim that populism is especially prevalent in the online forums of tabloid newspapers. In general, the online space and social media gave citizens the means to reach a broad audience by sharing their opinions and worldviews, which offers this position as a populist actor. Galpin and Trenz call this phenomenon “participatory populism” (Galpin and Trenz, 2019). Hameleers (Hameleers, 2019) studied populist messages online and found group dynamics that form a unified identity that empowers the collective, which expresses its discontent towards the system. Fernández-García and Salgado discovered similar patterns in an analysis of online comments in Portugal and Spain (Fernández-García and Salgado, 2020) and added that of the core motives, the most frequently found one was anti-elitism.

The manual coding of populism in text began with qualitative studies, such as the analysis of Belgian political broadcasts (Jagers and Walgrave, 2007) and the examination of Hugo Chávez’s populism (Hawkins, 2009). Other studies extended to coding smaller documents like manifestos, statements, and social media posts (Rooduijn, 2014; Rooduijn and Pauwels, 2011; Pauwels and Rooduijn, 2015; Aslanidis, 2018; Müller et al., 2017; Manucci and Weber, 2017; Engesser et al., 2017a) or employ-

ing (expert) surveys (Polk et al., 2017; Stavrakakis et al., 2017). While providing high-quality results, manual methods are expensive.

The use of computational power to automatically classify populist text aims at quantitatively analyzing large amounts of documents. There are dictionary-based methods by Rooduijn and Pauwels (Rooduijn and Pauwels, 2011) and Gründl (Gründl, 2022), which we use as a baseline, as they provide resources for German text especially. Gründl furthermore uses the same definition as this work and opts to identify anti-elitism, people-centrism and people-sovereignty. Other dictionary-based methods can be found in the publications of Bonikowski and Gidron (Bonikowski and Gidron, 2015; Bart Bonikowski, 2016).

In 2016 Hawkins and Castanho Silva applied ML algorithms to classify party manifestos and political speeches (Wirth et al., 2016). They found elastic net regression (Zou and Hastie, 2005) and logit boosting (Friedman et al., 2000) performed best on their data, however, they were only good at detecting non-populist documents and rather poor at detecting populist ones. In the following years, some annotated resources were published (Rooduijn et al., 2019; Hawkins et al., 2019; Klamm et al., 2023), but they either contain party manifestos and speeches or populism scores for parties. Subsequently, there is a lack of (openly available) annotated data to train ML models. Further attempts (Di Cocco and Monechi, 2022; Dai, 2019; Dai and Kustov, 2022) achieved good results using a random forest classifier (Breiman, 2001). Few attempts to use transformer models can be found that again classify manifestos or speeches (Ulin-skaite and Pukelis, 2021; Klamm et al., 2023).

To the best of our knowledge, there is no available publication that uses ML to classify populist user comments or social media data according to the ideational definition. Cabot et al. manually labeled English Reddit³-comments and used a RoBERTa-based model to classify them (Cabot et al., 2021). However, they focus on an “Us vs. Them” mentality and inspect the data for various non-elite outgroups like immigrants, liberals or Muslims. This is mostly not considered a part of populism in general, but rather seen as an expression of the underlying political attitude (Rooduijn, 2019).

3. Corpus Development

In this section, we describe the data sampling process, the conducted annotation study and its evaluation along with the observed metrics.

³<https://www.reddit.com/>

3.1. Data Collection

Our data set was kindly provided by Austrian daily newspaper *Der Standard* and features data from 1999 until November 2021 originating from their news forum. *Der Standard* is one of the two most read quality newspapers in Austria and considers itself as liberal and politically neutral, but is classified as left-liberal by press review *eurotopics*⁴. As tabloid news attract more populist comments and content and Blassnig et al. state that populist news articles lead to more populist user comments (Blassnig et al., 2019), the choice of a rather neutral source mitigates the effect of the reproduction of a populist statement found in the source article. To leave a comment, a user must register with an email and a pseudonym, which provides a means to maintain anonymity among users.

In regards to comparing populism during the pandemic, three different samples are drawn. The first contains only comments from articles that deal with the topic of COVID-19, the second contains comments from articles that were posted during the COVID-19 crisis (starting with the first COVID-19-related article in January 2020), but deal with all other topics. The last part contains comments that were posted in 2019, the year before the outbreak of COVID-19, as a reference. Table 1 shows the statistics of all three samples considering the comments, articles and users found in them.

After examining the data, we draw a sample that is

Table 1: Article and user stats from all samples

	Reference	COVID-19	Non-COVID
Comments	9,683,153	14,391,704	16,885,895
Articles	49,421	14,135	68,642
Users	55,780	64,401	82,354

labeled in an annotation study. For that, we clean the data by first filtering out all short comments containing ten or fewer words. To remove overtly long comments, we choose to remove all comments outside the 99% quantile of the whole data, which affects posts with more than 110 words.

A general problem in ML is having an imbalanced sample. Out of the whole body of comments, we consider populism to be a rare phenomenon and therefore use the dictionary of Gründl (Gründl, 2022) to pre-select potentially populist comments. To achieve this, we apply the dictionary to all samples and assess the number of dictionary terms present in each comment. To calculate the matching terms, the R-package provided by Gründl⁵ is used, which implements pattern matching with regular expressions to capture various forms of the

⁴<https://www.eurotopics.net/de/148488/der-standard>

⁵<https://github.com/jogruue/popdictR>

dictionary terms. The dictionary analysis confirms the expected rarity of comments containing populism, as across all samples a maximum of 3.7% include at least one term. Within those, there are only a few comments that feature six or more populist words, which makes it possible to manually observe them. There is one outlier with a score of 60 that repeats the word “Schande” (German word for “shame”) 60 times, which is a problem for most comments with a high score. Consequently, all posts with excessive repetition are removed.

Considering the available resources for the annotation study, we draw a sample of the size of 1,200 that entails 400 comments from each of the samples to account for potential differences in populist wordings based on the context. To address the issue of class imbalance, we opt to randomly select half of the sample exclusively from comments that were identified as populist by the dictionary. With the assumption that more populist words result in a higher likelihood of populism, a sample from all comments with a score greater than or equal to two is drawn. The other half is a completely random sample. Subsequently, we draw a sample of 200 random comments and a sample of 200 dictionary-annotated comments from each of the three samples to end up with the annotation corpus of 1,200 comments. The random sample is drawn to account for the bias of the dictionary. Due to the infrequency of populist comments, it is probable that the randomly selected sample will comprise only a limited number of such comments. This must be considered when determining the recall of the dictionary during evaluation.

The final sample’s populism scores can be found

Table 2: Absolute distribution of dictionary scores for the annotation sample

Score	Number of comments
0	587
1	12
2	551
3	46
4	4

in Table 2. There are 587 comments with a score of zero, which indicates that the random sample only includes 13 potentially populist comments with a dictionary score of one or more. The annotation study is performed on a randomly shuffled version of this prepared sample.

3.2. Annotation Study

To create the PopAut corpus, a labeling study involving multiple coders is conducted to thoroughly examine each comment for any instances of populist content. The goal is to annotate the com-

ments for the three populist motives presented earlier. There are scholars who argue that populism is a non-compensatory multidimensional phenomenon, which means that it requires all motives to be present (Wuttke et al., 2020). However, the majority of research reviewed in Section 2 sees it as sufficient if at least one of the motives is present, which is also reasonable given the brevity of the analyzed texts. Rather than reflecting the ideology as a whole, single statements mostly refer to certain aspects of it (Engesser et al., 2017b; Ernst et al., 2017; Esser et al., 2017). Thus, the setting entails a multiple-choice selection from the three motives of anti-elitism, people-centrism, and people-sovereignty, along with a “none” option. Hawkins showed that inter-rater reliability is high for political speeches and party manifestos and often a single coder is sufficient (Hawkins and Silva, 2016). Given the substantially distinct nature of our data, we aim to augment the validity of our annotation by involving external coders. Five participants (including one of the authors) with various academic backgrounds are invited to participate in the study. The goal is that each comment gets annotated by exactly three different people. Due to the resource availability, the data is split into three parts of 400 comments each. While two annotators review all 1,200 comments, the remaining three participants are assigned a subsample of 400 comments for annotation.

The study is carried out using the open-source tool LimeSurvey⁶ on a web server. Before launching the survey it is very important to establish precise annotation guidelines so that everyone has a shared concept of what populism is and how the motives can be defined. The main information is presented in text form, but every participant also receives a personal briefing in advance, where they get information about the data usage and a statement that they could withdraw from the study at any given point if they decide to. Additionally, they receive a few questions to collect demographic data and ask about their understanding of the content, because of the peculiarities in Austrian German compared to Standard German. In the following, the established procedure of the annotation study is explained step by step.

The first page the participants see during the personal briefing is the welcome page where they are presented with the most important information about the study. The instructor stresses that the goal is not to focus on detecting populism in the comments, but to detect comments that actually contain populist messages. This is done to emphasize the fact that a considerable proportion of comments are evidently not populist, and the objective is not to force any identification of populism

in them. It is furthermore explicitly mentioned that it is a multiple-choice task and comments can and will often include multiple motives.

Next, the guidelines on how to classify the content are displayed, which follow the work Ernst et al. (Ernst et al., 2019), who already conducted an annotation study on social media content, using the same definitions. The ideational definition of Mudde (Mudde, 2004) is given as the central reference, along with the publication. According to that, the three motives are introduced and the reference to the relevant publication is included so that the participants can get further information if desired. The following pages include an explanation of the motives, key messages and example comments as seen in Appendix 9.2. To facilitate the process of labeling, this information is available to them at any time in a text document, to not having to revisit the introduction if any uncertainties arise. In addition, we showcase a pre-selected example comment from the *Der Standard* data for each motive, providing an impression of how content related to that motive could look. The presentation of the motives comes with explanations in three dimensions: the populist key message that should be identified, along with the underlying ideology and a description of how it can manifest in a populist text. For anti-elitism, we added further information on who or what can be seen as an elite and for people-centrism we added a note that the German translation for people is “Volk” in this context, to make sure that people are seen as a group that includes the countries population.

After this introduction, the participants are presented with one comment after another. To further check if the participants have a shared perception of clearly populist comments, gold label comments are added. This means that for every batch of 400 comments, three comments are manually created following the definitions of the motives. The participants’ annotations for those comments are extracted later and a 100% agreement is expected. Each batch features one comment of every motive, resulting in the nine golden samples highlighted in Appendix 9.1, which are inserted at a random position.

3.3. Measurements

For the experiments, we create a binary label called *populism*, which is assigned if a participant finds any of the motives, as the models are trained for a binary classification task. We report the distribution and the agreement of all motives separately because those could be of interest for further work, but the criterion, which is important for the experiments is the agreement on the *populism* label. We measure the agreement of two or more coders by using Krippendorff’s α (Hayes and Krippendorff,

⁶<https://www.limesurvey.org/de/>

2007). α is an inter-rater reliability measure used to assess the agreement of multiple coders on different types of data (nominal in our case). Its basic assumptions rely on the concept of observed and expected agreement. Observed agreement is the proportion of actual agreement of the coders on the same class, while expected agreement is the proportion of agreement by chance. With observed disagreement as D_o and expected disagreement as D_e , Krippendorff's α is calculated as:

$$\alpha = 1 - \frac{D_o}{D_e}$$

This is applied to incorporate the probability of random disagreement based on the positive class. With this approach, α has a number of advantages over other inter-rater reliability measures. It is able to handle any number of coders, deals with different sizes of data sets and uses nominal weights to measure the importance of different categories, which is beneficial for our imbalanced sample. Hawkins and Castanho Silva (Hawkins and Silva, 2016) observed an agreement of $\alpha = 0.87$, which is high for such a complex phenomenon. When dealing with short, user-generated text, we do not expect such a high value. Thiele measured an agreement of $\alpha = 0.81$ for populist Facebook comments, which is very satisfactory. We additionally calculate the agreement on our self-created gold label comments separately and report those numbers to assess the accuracy of our annotation guidelines. After this, we create the final labels by majority vote, which means that a comment is populist if at least two participants detect any populist motive and subsequently agree that the comment features populist content, as the boundaries of the motives are sometimes fluent and can still be subjective.

3.4. Annotation Evaluation

In the following, the annotating author is called *P0*, the first participant, who also annotated all 1,200 comments, is called *P1* and the other participants, who annotated 400 comments each are called *P2-4*. We report demographic data collectively and not separately for every participant to retain anonymity. Concerning their native language, three participants indicate that they speak Standard German as their mother tongue, whereas one participant reports Austrian German as their native language. One of the participants has never lived in Austria, while the others state that they have lived there for two, six and 25 years and they lived in Vienna or Upper Austria for the most part. The participants across a range of academic domains, including education, veterinary medicine, and international business administration, all report a Master's or Diploma degree as their highest educational level.

Table 3: Distribution of labels for the first batch of 400 comments in the annotation study

	P0	P1	P2
Anti-elitism	90	102	76
People-centrism	18	13	32
People-sovereignty	6	10	6
None	302	288	308
Populism	98	112	92

Table 4: Distribution of labels for the second batch of 400 comments in the annotation study

	P0	P1	P3
Anti-elitism	95	100	109
People-centrism	19	16	22
People-sovereignty	8	11	7
None	296	290	284
Populism	104	110	116

The demographic questions reveal that the participants have a strong academic background and exhibit proficiency in the German language, along with a familiarity with the cultural context of Austria for most of them.

Due to the fact that three participants only labeled a batch of 400 comments each, we report the distribution of the given labels separately for those batches in Tables 3, 4 and 5.

The tables provide an overview of all potential responses and the corresponding populism label assigned upon the detection of at least one motive. The prevalence of *anti-elitism* is prominent among comments categorized as populist by all annotators. Additionally, the results indicate a strong class imbalance, which occurs despite the pre-selection of 50% of the sample with the populism dictionary.

The agreement of all participants in Table 6 shows satisfactory results for our populist label. An α value of 0.79 comes close to the results of Thiele (Thiele, 2022). Notably, the highest agreement can be observed for anti-elitism. In contrast, agreement on people-centrism is lower, potentially due to the subjective boundaries between people-centrism and people-sovereignty, as well as the multiple-choice format of the task.

Table 5: Distribution of labels for the third batch of 400 comments in the annotation study

	P0	P1	P4
Anti-elitism	87	97	104
People-centrism	14	10	14
People-sovereignty	9	8	9
None	304	299	286
Populism	96	101	114

Table 6: Krippendorff’s α for all participants of the annotation study across all labels

	Agreement
Anti-elitism	0.79
People-centrism	0.54
People-sovereignty	0.72
Populism	0.79

Table 7: Label distribution of the whole sample and the training, validation and test split

	Full	Training	Validation	Test
Populist	297	198	50	49
None	903	602	350	351

The high agreement on anti-elitism suggests that participants share a common understanding of who qualifies as an elite. Based on the high agreement, we use a majority vote and report the resulting label distributions for the PopAut corpus and sub-samples used in the experiments in Table 7. As expected, the results confirm a high class imbalance, with only approximately 24.8% of comments classified as populist. There are 30 comments marked as populist in the random half of the sample, which were not detected by the dictionary. The published PopAut corpus contains all comments in plain text, without further information about the author. Furthermore, it does not include the usernames of other authors. The corpus is published along with all participants’ vote on populist motives, to enable further investigation and use within the community, along with the final populist label. Lastly, we evaluate the agreement on the manually created gold standard samples placed in the annotation study. Here we can observe perfect agreement on each sample (i.e. $\alpha = 1$), suggesting that the participants have a shared perception of evidently populist comments.

4. Experiments

In the experimental section of this work, the validation of the annotated corpus is conducted by employing a series of classification techniques. To establish a robust baseline, existing dictionary-based methods and basic ML models are applied. These serve as reference points against which the performance of a BERT-based (Devlin et al., 2018) solution is compared, to showcase that the PopAut corpus is suitable for a classification task of populist user comments using transformer models. For the ML models, the data is split into 800 training, 200 validation and 200 test comments. The first baseline is the Gründl-dictionary, which is applied to the test data, labeling a comment as

populist if it contains at least one dictionary term. The second dictionary baseline is the German version of the dictionary of Rooduijn and Pauwels (Rooduijn and Pauwels, 2011), which is however only tailored towards anti-elitist content. Next, the best-performing model setting of Hawkins and Castanho Silva (Hawkins and Silva, 2016) is used by employing elastic net regression. A few adjustments are made as we deal with user-generated content and the original approach was made for political texts. As a first step, URLs within comments are removed, as they hold no semantic information. Further, in the original paper, they create a document-term-matrix with a cut-off of high and low-frequency terms, which is not done here because of the small document size. The following baselines used by Hawkins and Castanho Silva are added: a logistic-regression (LR), a support-vector-machine (SVM) and a random forest (RF) classifier implemented using the standard parameters given by the scikit-learn framework (Pedregosa et al., 2011).

For the transformer-based solution, we experiment with different models available in the Hugging Face transformers library⁷ and pre-processing steps and report the best setup. Therefore the “bert-base-german-cased”⁸ model is used along with the recommended parameter settings reported by Devlin et al. (Devlin et al., 2018). To reduce possible noise contained in news comments, HTML tags, non-ASCII characters, digits, single-letter words and multiple white spaces are removed. Because training can be unstable, especially for small sample sizes (Zhang et al., 2020), we report the average score over five different training runs with five epochs each. The model is optimized for F1-score using the validation set and early stopping.

Table 8 shows the results of all experiments regarding accuracy, precision, recall and F1-score. Scores are calculated for the minority class of populist comments. However, in our study, the recall needs to be interpreted carefully, because we use the dictionary to create parts of the annotated sample. Therefore, the recall score achieved by the dictionary was expected to generally be high. Subsequently, the goal is to reach a substantially higher precision using a transformer-based model, while also retaining a high recall, to outperform the dictionary-based method in regards to F1.

Indeed the Gründl-dictionary has an expected high recall of 0.88 and only accounts for seven false negatives. Its low precision of 0.44 and 0.45 is not surprising, as the results of the annotation study already showed that there are only less than 300 samples where the participants actually

⁷<https://huggingface.co/>

⁸<https://huggingface.co/bert-base-german-cased>

Table 8: Performance of highlighted text classification techniques (best performance for each score is written in bold font).

	Gründl	R&P	EN	LR	SVM	RF	BERT
Accuracy	0.69	0.77	0.77	0.81	0.74	0.77	0.80
Precision	0.44	0.53	0.54	0.92	0.46	0.58	0.57
Recall	0.88	0.51	0.39	0.22	0.33	0.22	0.79
F1	0.59	0.52	0.45	0.36	0.38	0.32	0.66

agreed on the populist label of the dictionary. R&P references the dictionary by Rooduijn and Pauwels (Rooduijn and Pauwels, 2011). It only has a relatively low recall of 0.51 as it is solely designed to capture anti-elitism and is comprised of 20 dictionary terms only. It only outperforms the Gründl-dictionary in regards to accuracy and precision. When looking at the baseline ML classifiers, we can see that the training size is insufficient to perform satisfactorily for those algorithms because they only find a small number of actual populist comments. However, they tend to be more precise in the few populist predictions they make. With our BERT-model, we reach the goal of improving on the precision of the dictionary-based method by 12%, while keeping a recall of 79%, which results in an improved F1-score by 7%.

The experiments with the annotated sample could meet our expectations of outperforming the dictionary in terms of precision and subsequently F1-score. This is an interesting finding, as it was achieved with only 800 samples for BERT fine-tuning. The transformer architecture strongly outperforms the other ML methods that serve as a baseline and can better learn the underlying patterns of the complex phenomenon of populism. The improvement in precision can be seen as an improvement in detecting non-populist comments and subsequently, the ability to distinguish them better from actual populist comments. Especially when compared to the performance of human annotators who reach an agreement of 0.79 on the task.

5. Discussion

Populism detection is still a research field that is at an early stage, which is why data is lacking for example, but our findings suggest high potential in the use of the ML methods when annotated data is available. In this section, we discuss the results of our study.

Overall the annotation of the PopAut corpus provides promising results for its further use within the research community. The annotators could reach a high agreement on whether a comment contains populist motives or not. Regarding the distribution of the populism motives in the PopAut corpus, it is

noticeable that anti-elitism is a lot more prevalent than all other motives. We furthermore observe the highest agreement on anti-elitism. This could be an indicator that the citizen's populism rather focuses on blaming politicians or experts, especially during the COVID-19 pandemic they were held responsible for the countermeasures that drastically changed social life. Therefore more fine-grained annotations for the future could entail labels for mentioned elites similar to the work of Klamm et al. (Klamm et al., 2023). Additionally, upon qualitative examination of the data, a notable prevalence of in-group versus out-group dynamics emerges. Users frequently classify individuals into those who adhere to and trust governmental measures versus those who do not, as well as vaccinated versus unvaccinated individuals. This motive is not represented in the minimalistic definition and might be a characteristic of the "pandemic populism" indicated in the work of Boberg et al. (Boberg et al., 2020). After the annotation study, we demonstrated that we can outperform existing methods for automated populism detection with the use of pre-trained BERT-based models fine-tuned on our human-annotated data. We could especially reduce the number of false positives, which gives us more precise results and helps to capture the phenomenon more accurately, which is desirable for large-scale analysis. However, with a precision of 57%, the amount of false positives is still reasonably high.

6. Limitations & Future Work

Despite the rich potential of the provided dataset, we want to address the limitations of our annotation study. Due to the time-consuming nature of manual data annotation, the sample size of 1,200 annotated is relatively small, as seen for the ML baselines. However, fine-tuning of large pre-trained models could still achieve good results with the amount of training data, which is promising regarding the brevity of the text and the complexity of populism.

Another factor to be considered is the bias of the dictionary that was necessary to be used to sample half of the corpus, because of the rarity of populist comments among the entirety. Further bias is introduced by only incorporating data between 2019-2021. Consequently, this makes it an especially valuable resource to examine the prevalence of populist rhetoric during the COVID-19 pandemic in Austria. Considering the scale of this work, the choice of *Der Standard* as a media source that is not considered to have a populist orientation gives a good overview, but incorporating other news sources might give a more diverse representation of public opinion. Geographically the work is limited to Austria, which

comes with peculiarities in the language used, but the annotation framework and definitions used aim at being used to compare populism agnostic to geographic or political backgrounds, which means that it can be used to extend the corpus in the future, by incorporating German data for example. Further potentials for increasing the breadth of the dataset's coverage include exploring semi-supervised text annotation methodologies and leveraging generative language models to annotate or classify data. Finally the advancements in the area of large language models additionally promise enhanced capabilities for the classification task.

7. Acknowledgement

The financial support by the Austrian Federal Ministry of Labour and Economy, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association is gratefully acknowledged.

8. Bibliographical References

- Koen Abts and Stefan Rummens. 2007. [Populism versus democracy](#). *Political Studies*, 55(2):405–424.
- Paris Aslanidis. 2018. [Measuring populist discourse with semantic text analysis: an application on grassroots populist mobilization \[quality & quantity\]](#). *Quality and Quantity*, 53:1241–1263.
- Robert R. Barr. 2009. [Populists, outsiders and anti-establishment politics](#). *Party Politics*, 15(1):29–48.
- Noam Gidron and Bart Bonikowski. 2016. [Populism in legislative discourse : Evidence from the european parliament , 1999-2004](#).
- Sina Blassnig, Sven Engesser, Nicole Ernst, and Frank Esser. 2019. [Hitting a nerve: Populist news articles lead to more frequent and more populist reader comments](#). *Political Communication*, 36(4):629–651.
- Svenja Boberg, Thorsten Quandt, Tim Schatto-Eckrodt, and Lena Frischlich. 2020. [Pandemic populism: Facebook pages of alternative news media and the corona crisis – a computational content analysis](#).
- Bart Bonikowski and Noam Gidron. 2015. [The Populist Style in American Politics: Presidential Campaign Discourse, 1952–1996](#). *Social Forces*, 94(4):1593–1621.
- Linda Bos, Wouter van der Brug, and Claes de Vreese. 2011. [How the media shape perceptions of right-wing populist leaders](#). *Political Communication*, 28(2):182–206.
- Roberta Bracciale and Antonio Martella. 2017. [Define the populist political communication style: the case of italian political leaders on twitter](#). *Information, Communication & Society*, 20(9):1310–1329.
- L Breiman. 2001. [Random forests](#). *Machine Learning*, 45:5–32.
- Pere-Lluís Huguet Cabot, David Abadi, Agneta H. Fischer, and Ekaterina Shutova. 2021. [Us vs. them: A dataset of populist attitudes, news bias and emotions](#). In *EACL*.
- Margaret Canovan. 1999. [Trust the people! populism and the two faces of democracy](#). *Political Studies*, 47(1):2–16.
- Yaoyao Dai. 2019. [Measuring populism in context: A supervised approach with word embedding models](#). *Manifesto Corpus Conference, Berlin*.
- Yaoyao Dai and Alexander Kustov. 2022. [When do politicians use populist rhetoric? populism as a campaign gamble](#). *Political Communication*, 39(3):383–404.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).
- Jessica Di Cocco and Bernardo Monechi. 2022. [How populist are parties? measuring degrees of populism in party manifestos using supervised machine learning](#). *Political Analysis*, 30(3):311–327.
- Sven Engesser, Nicole Ernst, Frank Esser, and Florin Büchel. 2017a. [Populism and social media: how politicians spread a fragmented ideology](#). *Information, Communication & Society*, 20(8):1109–1126.
- Sven Engesser, Nayla Fawzi, and Anders Olof Larsson. 2017b. [Populist online communication: introduction to the special issue](#). *Information, Communication & Society*, 20(9):1279–1292.
- Nicole Ernst, Sina Blassnig, Sven Engesser, Florin Büchel, and Frank Esser. 2019. [Populists prefer social media over talk shows: An analysis of populist messages and stylistic elements across six countries](#). *Social Media + Society*, 5(1):2056305118823358.

- Nicole Ernst, Sven Engesser, Florin Büchel, Sina Blassnig, and Frank Esser. 2017. [Extreme parties and populism: an analysis of facebook and twitter across six countries](#). *Information, Communication & Society*, 20:1–18.
- Frank Esser, Agnieszka Stępińska, and David Nicolas Hopmann. 2017. [Populism and the media: cross-national findings and perspectives](#). In T. Aalberg, F. Esser, C. Reinemann, J. Strömbäck & C. d. Vreese (Eds.), *Populist Political Communication in Europe*. Routledge, 365-380.
- Belén Fernández-García and Susana Salgado. 2020. [Populism by the people: An analysis of online comments in portugal and spain*](#). pages 210–219.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. 2000. [Additive logistic regression: a statistical view of boosting \(With discussion and a rejoinder by the authors\)](#). *The Annals of Statistics*, 28(2):337 – 407.
- Charlotte Galpin and Hans-Jörg Trenz. 2019. [Participatory populism: Online discussion forums on mainstream news sites during the 2014 european parliament election](#). *Journalism Practice*, 13(7):781–798.
- Johann Gründl. 2022. [Populist ideas on social media: A dictionary-based measurement of populist communication](#). *New Media & Society*, 24(6):1481–1499.
- Michael Hameleers. 2019. [The Populism of Online Communities: Constructing the Boundary Between “Blameless” People and “Culpable” Others](#). *Communication, Culture and Critique*, 12(1):147–165.
- Michael Hameleers, Linda Bos, and Claes de Vreese. 2016. [The Netherlands: A Heartland Full of Insights into Populist Communication](#).
- K. A. Hawkins, R. Aguilar, B. Castanho Silva, E. K. Jenne, B. Kocijan, and C. Rovira Kaltwasser. 2019. [Measuring populist discourse: The global populism database](#). *EPSA*.
- Kirk A. Hawkins. 2009. [Is chávez populist?: Measuring populist discourse in comparative perspective](#). *Comparative Political Studies*, 42(8):1040–1067.
- Kirk A. Hawkins and Bruno Castanho Silva. 2016. [A head-to-head comparison of human-based and automated text analysis for measuring populism in 27 countries*](#).
- Andrew F. Hayes and Klaus Krippendorff. 2007. [Answering the call for a standard reliability measure for coding data](#). *Communication Methods and Measures*, 1(1):77–89.
- Jan Jagers and Stefaan Walgrave. 2007. [Populism as political communication style: An empirical study of political parties’ discourse in belgium](#). *European Journal of Political Research*, 46(3):319–345.
- Robert S. Jansen. 2011. [Populist mobilization: A new theoretical approach to populism](#). *Sociological Theory*, 29(2):75–96.
- Cristóbal Rovira Kaltwasser. 2012. [The ambivalence of populism: threat and corrective for democracy](#). *Democratization*, 19(2):184–208.
- Christopher Klamm, Ines Rehbein, and Simone Paolo Ponzetto. 2023. [Our kind of people? detecting populist references in political debates](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1227–1243, Dubrovnik, Croatia. Association for Computational Linguistics.
- Benjamin Krämer. 2014. [Media Populism: A Conceptual Clarification and Some Theses on its Effects](#). *Communication Theory*, 24(1):42–60.
- Benjamin Krämer. 2017. [Populist online practices: the function of the internet in right-wing populism](#). *Information, Communication & Society*, 20(9):1293–1309.
- Ernesto Laclau. 2005. [On populist reason / Ernesto Laclau](#). Verso New York.
- Luca Manucci and Edward Weber. 2017. [Why the big picture matters: Political and media populism in western europe since the 1970s](#). *Swiss Political Science Review*, 23.
- Gianpietro Mazzoleni. 2008. [Populism and the Media](#), pages 49–64. Palgrave Macmillan UK, London.
- Benjamin Moffitt. 2016. [The Global Rise of Populism: Performance, Political Style, and Representation](#), hardcover edition. Stanford University Press.
- Cas Mudde. 2004. [The populist zeitgeist](#). *Government and Opposition*, 39(4):541–563.
- Cas Mudde. 2017. [Populism: An Ideational Approach](#). In *The Oxford Handbook of Populism*. Oxford University Press.
- Cas Mudde and Kaltwasser Cristóbal, Rovira, editors. 2012. [Populism in Europe and the Americas: Threat or Corrective for Democracy?](#), hardcover edition. Cambridge University Press.

- Cas Mudde and Cristóbal Rovira Kaltwasser. 2017. *Populism: A Very Short Introduction*. Oxford University Press.
- Philipp Müller, Christian Schemer, Martin Wettstein, Anne Schulz, Dominique S. Wirz, Sven Engesser, and Werner Wirth. 2017. *The Polarizing Impact of News Coverage on Populist Attitudes in the Public: Evidence From a Panel Study in Four European Democracies*. *Journal of Communication*, 67(6):968–992.
- Teun Pauwels and Matthijs Rooduijn. 2015. *Populism in Belgium in times of crisis: Intensification of Discourse, Decline in Electoral Support*. ECPR Press, Colchester, UK.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Jonathan Polk, Jan Rovny, Ryan Bakker, Erica Edwards, Liesbet Hooghe, Seth Jolly, Jelle Koedam, Filip Kostelka, Gary Marks, Gijs Schumacher, Marco Steenbergen, Milada Vachudova, and Marko Zilovic. 2017. *Explaining the salience of anti-elitism and reducing political corruption for political parties in europe with the 2014 chapel hill expert survey data*. *Research & Politics*, 4(1):2053168016686915.
- Kenneth M. Roberts. 2006. *Populism, political conflict, and grass-roots organization in latin america*. *Comparative Politics*, 38(2):127–148.
- M. Rooduijn, S. Van Kessel, C. Froio, A. Pirro, S. De Lange, D. Halikiopoulou, P. Lewis, C. Mudde, and P. Taggart. 2019. *The populist: An overview of populist, far right, far left and eurosceptic parties in europe*.
- Matthijs Rooduijn. 2014. *The nucleus of populism: In search of the lowest common denominator*. *Government and Opposition*, 49(4):573–599.
- Matthijs Rooduijn. 2019. *State of the field: How to study populism and adjacent topics? a plea for both more and less focus*. *European Journal of Political Research*.
- Matthijs Rooduijn and Teun Pauwels. 2011. *Measuring populism: Comparing two methods of content analysis*. *West European Politics*, 34(6):1272–1283.
- Anne Schulz, Philipp Müller, Christian Schemer, Dominique Wirz, Martin Wettstein, and Werner Wirth. 2018. *Measuring populist attitudes on three dimensions*. *International Journal of Public Opinion Research*, 30:316–326.
- Yannis Stavrakakis, Ioannis Andreadis, and Giorgos Katsambekis. 2017. *A new populism index at work: identifying populist candidates and parties in the contemporary greek context*. *European Politics and Society*, 18(4):446–464.
- Daniel Thiele. 2022. *Pandemic populism? how covid-19 triggered populist facebook user comments in germany and austria*. *Politics and Governance*, 10(1):185–196.
- Jogile Ulinskaite and Lukas Pukelis. 2021. *Identifying populist paragraphs in text: A machine-learning approach*. *CoRR*, abs/2106.03161.
- Kurt Weyland. 2001. *Clarifying a contested concept: Populism in the study of latin american politics*. *Comparative Politics*, 34(1):1–22.
- Werner Wirth, Frank Esser, Martin Wettstein, Sven Engesser, Dominique S. Wirz, Anne Schulz, Nicole Ernst, Florin Büchel, Daniele Caramani, Luca Manucci, Marco R. Steenbergen, Laurent Bernhard, Edward Weber, Regula Hänggeli, Caroline Dalmus, Christian Schemer, and Philipp Müller. 2016. *The appeal of populist ideas, strategies and styles: A theoretical model and research design for analyzing populist political communication*. (88).
- Dominique Stefanie Wirz. 2018. *Persuasion through emotion? an experimental test of the emotion-eliciting nature of populist communication*. *International Journal of Communication*, 12:25.
- Alexander Wuttke, CHRISTIAN SCHIMPF, and Harald Schoen. 2020. *When the whole is greater than the sum of its parts: On the conceptualization and measurement of populist attitudes and other multidimensional constructs*. *American Political Science Review*, 114:1–19.
- Tianyi Zhang, Felix Wu, Arzoo Katiyar, Kilian Q. Weinberger, and Yoav Artzi. 2020. *Revisiting few-sample bert fine-tuning*.
- Hui Zou and Trevor Hastie. 2005. *Regularization and variable selection via the elastic net*. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 67(2):301–320.

9. Appendix

9.1. Manually Created Gold Standard Comments

Following are the comments that are manually created and randomly inserted in the annotation study to validate the annotation guidelines along with an English translation:

- “Die da oben mit ihrer Coronapolitik sind doch an allem Schuld!” (anti-elitism)
- “Those folks up there with their COVID policies are to blame for everything!” (anti-elitism)
- “Das ist doch eh alles nur noch gesteuert von den finanziellen Eliten, die mit den Problemen von uns Normalos nichts zu tun haben.” (anti-elitism)
- “It’s all just being controlled by the financial elites who have nothing to do with the issues us ordinaries, anyway.” (anti-elitism)
- “Auf das was die Medien berichten, kann man sich als Bürger nicht mehr verlassen, die plappern doch nurnoch die Politiker nach” (anti-elitism)
- “As a citizen, you can’t rely on what the media reports anymore, they’re just parroting the politicians” (anti-elitism)
- “Die letzten Wochen haben doch ganz klar die Meinung des Volkes gezeigt. Ich gebe hier nur die allgemeine Stimmung wieder.” (people-centrism)
- “The past few weeks have clearly shown the people’s opinion. I’m just echoing the general mood here.” (people-centrism)
- “Diese Entwicklung ist definitiv der Bevölkerung zuzuschreiben, ich bin stolz auf alle Österreicher” (people-centrism)
- “This development is definitely thanks to the citizens, I’m proud of all Austrians” (people-centrism)
- “Ich denke, ich spreche hier für alle Österreicher, wenn ich sage, dass dieses Gesetz keiner braucht” (people-centrism)
- “I think I speak for all Austrians when I say that nobody needs this law.” (people-centrism)
- “Nach den letzten gescheiterten Versuchen, sollte die Bevölkerung nun doch selbst entscheiden dürfen, was sie will. Abstimmung über Lockdowns jetzt!” (people-sovereignty)

- “After the recent failed attempts, the population should be allowed to decide for themselves now what they want. Vote over lockdowns now!” (people-sovereignty)
- “Das ist doch alles keine Demokratie mehr, wenn der gemeine Bürger nichts mehr zu sagen hat.” (people-sovereignty)
- “This isn’t even democracy anymore when the ordinary citizen has nothing to say anymore.” (people-sovereignty)
- “Also in Zukunft sehe ich schwarz für dieses Land, wenn die Meinung der Bürger nicht stärker ins Gewicht fällt..” (people-sovereignty)
- “So, in the future, I see dark times ahead for this country if the citizens’ opinion don’t carry more weight..” (people-sovereignty)

9.2. Annotation Guidelines

Figures 2, 3 and 4 show the operationalization of populist motives and key motives as outlined in (Ernst et al., 2019). The participants of the annotation study were given these instructions along with an example comment for every motive. Rough English translations of the respective examples are:

- “Unfortunately, the elite has forced the redistribution of wealth from the bottom to the top, leading us into these precarious work situations today. Politicians can change this right away. Why don’t they? Because they follow the commands and desires of these elites and no longer serve the people. It’s time for a system change.” (Figure 2)
- “The people do not want to vaccinate, the people prefer to wear mask and test! The government can do whatever it wants, the people do what the people want! And votes Sincerely ...” (Figure 3)
- “Even though unlikely, I hope that my lines will find their way to you, my esteemed PEOPLE representatives. Even if improbable, I offer you a personal meeting, in which you can hear a voice from exactly this PEOPLE, which finances your “common sense outbursts”. For we are all aware that the quiet voices of your people can only be heard in election year. Even if this personal meeting is highly unlikely, I would be happy to be taught why YOUR common sense is better than ours;” (Figure 4)

Figure 2: Description of the motive anti-elitism in the annotation study, including an example comment

Anti-Elitism:

Here it is necessary to describe what an 'elite' is first. Elites in the context of populism are seen as an actor or group defined by power, to morally separate them from the people. Groups that are considered elites here are: economic elites holding financial power, political elites, media elite and cultural elites.

Populist key message	Underlying ideology	Description
Discrediting the elite	Elites are corrupt.	Elites are accused of being malevolent, criminal, lazy, stupid, extremist, racist, undemocratic, and so on. The elite are called names and denied morality, charisma, credibility, intelligence, competence, consistency, and so on.
Blaming the elite	Elites are harmful.	Elites are described as a threat/burden, responsible for negative developments/situations, or as having committed mistakes or crimes. Elites are described as not being a source of enrichment or responsible for positive developments/situations.
Detaching the elite from the people	Elites do not represent the people.	Elites are described as not belonging to the people, not being close to the people, not knowing the people, not speaking for the people, not caring for the people, or not performing everyday actions.

Example comment containing this motive:

Leider hat die Elite die Umverteilung von unten nach oben forciert, wodurch wir heute in diese präkeren Arbeitsverhältnisse getrieben wurden. Die Politiker könnten dies sofort wieder ändern.

Warum sie das nicht tun? Weil Sie die Befehle und Wünsche dieser Eliten ausführen und nicht mehr dem Volk dienen.

Es ist Zeit für einen Systemwechsel!

Figure 3: Description of the motive people-centrism in the annotation study, including an example comment

People-Centrism:

(Note: People here would be translated with 'Volk' in German)

Populist key message	Underlying ideology	Description
Stressing the people's virtues	The people are virtuous.	The people are bestowed with morality, charisma, credibility, intelligence, competence, consistency, and so on. The people are exempt from being malevolent, criminal, lazy, stupid, extremist, racist, undemocratic, and so on.
Praising the people's achievements	The people are beneficial.	The people are described as being enriched or responsible for a positive development/situation. The people are described as not being a threat/burden, not being responsible for negative developments/situations, nor as having committed mistakes or crimes.
Stating a monolithic people	The people are homogeneous.	People are described as sharing common feelings, desires, or opinions.
Demonstrating closeness to the people	The populist represents the people.	The speaker describes himself as belonging to the people, being close to the people, knowing the people, speaking for the people, caring for the people, agreeing with the people, or performing everyday actions. The speaker claims to represent or embody the people.

Example comment containing this motive:

Das Volk will nicht impfen, das Volk will lieber Maske tragen und testen !

Da kann die Regierung am Kopf stehen und mit den Beinen Fliegen fangen, das Volk macht, was das Volk will !

Und wählt Mit freundlichen Grüßen ...

Figure 4: Description of the motive people-sovereignty in the annotation study, including an example comment

People-Sovereignty:

Populist key message	Underlying ideology	Description
Demanding popular sovereignty	The people are the ultimate sovereign.	The speaker argues for general institutional reforms to grant the people more power by introducing direct-democratic elements or increasing political participation. The speaker argues in favor of granting more power to the people within the context of a specific issue (e.g., election, immigration, security).
Denying elite sovereignty	The elites deprive the people of their sovereignty.	The speaker argues in favor of granting less power to elites within the context of a specific issue (e.g., election, immigration, security).

Example comment containing this motive (also contains people-centrism):

Wenn auch unwahrscheinlich, so hoffe ich doch, dass meine Zeilen den Weg zu Euch, meine geschätzten VOLKsvertreter finden. Wenn auch unwahrscheinlich, so biete ich Euch ein persönliches Treffen an, indem Ihr eine Stimme aus genau diesem VOLK, welches Eure „Hausverstand-Ergüsse“ finanziert, hören könnt. Denn es ist uns allen bewusst, dass die leisen Stimmen Eures Volkes nur im Wahljahr zu hören sind. Wenn auch dieses persönliche Treffen höchst unwahrscheinlich ist, so ließe ich mich gerne belehren, warum EUER Hausverstand besser ist als Unserer;