Proceedings of the 13th Conference on Language Resources and Evaluation (LREC 2022), pages 5849–5858
Marseille, 20-25 June 2022

Who's in, who's out? Predicting the Inclusiveness or Exclusiveness of Personal Pronouns in Parliamentary Debates

Ines Rehbein, Josef Ruppenhofer
Mannheim University, Leibniz-Institut der Deutschen Sprache, Mannheim
rehbein@uni-mannheim.de, ruppenhofer@ids-mannheim.de

Abstract

This paper presents a compositional annotation scheme to capture the clusivity properties of personal pronouns in context, that is their ability to construct and manage in-groups and out-groups by including/excluding the audience and/or non-speech act participants in reference to groups that also include the speaker. We apply and test our schema on pronoun instances in speeches taken from the German parliament. The speeches cover a time period from 2017-2021 and comprise manual annotations for 3,126 sentences. We achieve high inter-annotator agreement for our new schema, with a Cohen's κ in the range of 89.7-93.2 and a percentage agreement of $> 96\%$. Our exploratory analysis of in/exclusive pronoun use in the parliamentary setting provides some face validity for our new schema. Finally, we present baseline experiments for automatically predicting clusivity in political debates, with promising results for many referential constellations, yielding an overall 84.9% micro F1 for all pronouns.

Keywords: clusivity, pronouns, political text analysis

1. Introduction

This paper investigates the use of personal pronouns in political discourse. Personal pronouns are a central rhetorical device for building group identities and for *othering* third parties (Dam, 2015). Tyrkkö (2016) considers personal pronouns as "one of the primary linguistic features used by political speakers to manage their audiences' perceptions of in-groups and out-groups". This makes them an important device for populist rhetoric where the speaker evokes a dichotomous view of society, *us-versus-them* (see, e.g., Mudde (2004; Mudde and Kaltwasser (2017)).

Studies on the interface between corpus linguistics, pragmatics, discourse studies and political science have presented empirical investigations of the use of personal pronouns in political discourse. One example is Tyrkkö (2016) who presents a diachronic investigation of the use of pronouns in political speeches and notes a shift in the early 20th century where many politicians changed their self-centric style and started using more instances of the inclusive plural forms.

One of the drawbacks of previous work, however, is that the presented empirical findings are usually based on word frequencies (Vuković, 2012; Tyrkkö, 2016; Alavidze, 2017), and only few studies have tried to systematically investigate this topic in more detail, i.e., by trying to measure the agreement between human annotators for disambiguating the referents of personal pronouns in political speeches, or by presenting large-scale studies of the use of personal pronouns beyond word frequencies.

More linguistically oriented studies, on the other hand, stress the capacity of 1st person personal pronouns (1PPL) to distinguish between *inclusive* and *exclusive* uses, depending on whether an addressee (or addressees) are included in or excluded from the set of referents which also contains the speaker (Filimonova, 2005). Cysouw (2002), for example, distinguishes 3 senses of *we*: (i) `minimal inclusive` uses that refer to speaker and hearer only (1.1); (ii) `augmented inclusive` uses that add a third party to the set (1.2); and (iii) `exclusive` uses of *we* that exclude the addressee (1.3).

Example 1.1. Shall we talk on the phone tomorrow?

Example 1.2. Kim will arrive at 11. Shall we go to lunch then?' [all three of us]

Example 1.3. 'We're going to the movies. What are your plans?'

According to Bickel and Nichols (2005), inclusive-exclusive oppositions are common in the Americas, near-universal in Australia, common in eastern Asia, rare in the rest of Eurasia, and fairly rare in Africa.[1] Languages that have an inclusive-exclusive opposition differ in the exact distinctions that are made. For instance, some languages have a special pronoun for the inclusive, but the marking of the exclusive is identical to 'I'. Other languages not only differentiate inclusive and exclusive but also within each category differentiate the case 'exactly two' from the case 'more than two'. We refer the reader to the discussion in Bickel and Nichols (2005) and WALS online for details.

Unsurprisingly, clusivity is not directly coded in Standard German, the language of the political texts that

[1] WALS Online `https://wals.info/feature/39A#3/23.08/98.53;lastaccessed:2022-04-24` also shows the distinction to be grammaticalized largely in the Circum-Pacific linguistic area.

we study, nor in English.[2] We argue that the missing grammaticalisation allows speakers to exploit the inherent ambiguity of referents of plural personal pronouns (PPL) for rhetorical purposes.

In previous work (Rehbein et al., 2021), we carried out an annotation study in which we classified instances of 1st person plural pronouns *wir, uns, unser* 'we, us, our' in German parliamentary debates, assigning instances to one of 9 different categories such as GOVERNMENT, PARLIAMENT, PARTY, COUNTRY etc. We showed that inter-annotator agreement varies systematically, with some classes being easier to discriminate while others are inherently ambiguous, thus providing evidence for the scattergun effect described by Allen (2007) below.[3]

> "Shifting identity through pronoun choice and using pronouns with ambiguous referents enables politicians to appeal to diverse audiences which helps broaden their ability to persuade the audience to their point of view. It is a scattergun effect —shoot broadly enough and you'll hit something". (Allen, 2007:12)

In this work, we extend our previous efforts by augmenting our corpus with annotations that capture the clusivity properties of instances of 1st, 2nd and 3rd person plural pronouns. We present a general, compositional annotation scheme that can be applied to texts from any domain. Our annotations provide the basis for investigations of the strategic use of personal pronouns in political text. Our main focus is on investigating how different political actors make use of inclusive versus exclusive uses of pronouns in context, in order to build group identity and for *othering* third parties.

Towards that end, the paper starts with an overview of related work (Section 2). Then we present our annotation scheme and report results from our annotation study (Section 3.2). In Section 4.2, we provide baselines for the automatic prediction of clusivity features in political text. In Section 5, we conclude and outline future work.

2. Related Work

We now present a short overview of related work, focussing on the use of pronouns in political discourse.

Pronouns in parliamentary text Research on the interface between corpus linguistics, pragmatics, discourse studies and political science has studied the use of personal pronouns in political text, using qualitative and empirical methods. Tyrkkö (2016) presents a diachronic study of the use of personal pronouns in political speeches over two centuries, showing shifts from a self-centric style (marked by frequent use of *I*) towards the more inclusive use of 1PPL forms in the 1920s, which the author ties to the emergence of broadcast media. The study does not disambiguate 1PPL forms but counts all of them as inclusive.

Íñigo-Mora (2004) studies the use of *we* in 5 Question Time Sessions of the British parliament, where MPs ask questions of government ministers. She distinguishes what she calls exclusive, inclusive, generic and parliamentary uses of *we* and examines their distribution across different combinations of interactants (opposition MP to member of government; member of government to opposition MP; member of government and supportive MP (in either direction)).[4] The frequency distribution is interpreted along two dimensions: (i) power and distance and (ii) identity, community and persuasion. Among the findings is that exclusive uses of *we* constitue the most common type overall, accounting for 53.4% of all tokens. Exclusive *we* is at its most dominant in interactions from government supporting MPs to opposition MPs (76.1%) while it is hardly ever used in questions from opposition MPs to a member of government, which is taken to reflect the power dynamics. Inclusive uses of *we* were found to be much rarer overall, making up 14.5% of all tokens. None of these are uttered by opposition members speaking to members of government, while three quarters are produced between government supporting MPs and members of government, expressing shared identity. Opposition MPs mostly use generic and parliamentary *we*, thus affiliating themselves with the parliament as a distinct branch of government and the country at large, likely because that is where persuasion is most likely to succeed. It is unclear to what extent these results carry over to the plenary setting.

Non-parliamentary political discourse Studies of 1PPL have also targeted other types of interactions. Bull and Fetzer (2006) analyze the use of *you* and *we* in tv interviews with British politicians that were broadcast during the 1997 and 2001 British general elections and just before the war with Iraq in 2003. The focus of the study was on question-response sequences in which politicians make use of pronominal shifts as a means of equivocation to effect shifts of accountability and responsibility. Proctor and Su (2011) examine the use of *we* by four (vice-)presidential candidates in debates and interviews around the time of the 2008 US election. The study focuses on which groups are the referents of *we* and which entities are picked out by possessive NPs

[2] See, however, Simon (2005) who presents some evidence for a grammaticalisation of clusivity for 2nd person plural pronouns in Bavarian, a dialect spoken in Southern Germany. In that case, the relevant 2nd person forms are argued to grammaticalize not only reference to a set of addressees but also necessarily to non-speech-act-participants.

[3] Percentage agreement for the more frequent classes was in the range of 86-92% while for the less frequent classes, IAA varied between 58.8–66.7%. Cohen's κ for all 1PPL annotations was quite high with 0.82.

[4] There is no generally agreed-upon terminology used to distinguish uses of *we*, either in general or in the political or parliamentary context. For Inigo-Mora the generic *we* refers to "a kind of patriotic "we" that embraces all British people". In the terminology of Quirk et al. (1985) this would be called a collective use. In our previous work (Rehbein et al., 2021), the uses at issue would be labeled "COUNTRY".

S	A	O	Example	Reference to
+	+	+	**We** must stand together.	the country/all citizens of this country
+	+	-	**We** must approve the budget.	the parliament
+	-	+	As Free Democrats, **we** ...	all members of the party
+	-	-	**I** am very grateful to you.	the speaker (sg only)
-	+	+	The YPG fights alongside **you** against the IS.	The German government & the US
-	+	-	I am very grateful to **you**.	specific MPs (sg or pl)
-	-	+	**They** are fighting against the IS.	The YPG
-	-	-	200 years ago, **we** have founded this nation.	generic uses (here: excluding everybody alive)
NONE			That's a difficult question but we must answer **it**.	the question (German: *sie*, 3.ps.sg.fem)
NONE			Those are difficult questions and we must solve **them**.	the questions (German: *sie* 3.ps.pl.fem)

Table 1: Annotation scheme for *clusivity*, capturing inclusive/exclusive uses of personal pronouns (**S**: Speaker, **A**: Addressee, **O**: Other).

of the form *our N*, considering the results in light of the candidates' political stature and targeted office as well as the differences between debate and interview settings. Wieczorek (2009) studies a broader notion of clusivity on 2008 campaign speeches by Barack Obama. Clusivity in that work is defined along spatial, temporal and ideological axes relative to the speaker as a deictic center, and items other than personal pronouns can code clusivity. In this paper, we retain the focus on personal pronouns in political text but extend previous work by presenting a compositional, domain-independent annotation schema for capturing inclusive and exclusive uses of personal pronouns. We apply our schema to parliamentary debates and report empirical results from an annotation study.

3. Annotation Study

We extracted a dataset of parliamentary debates from the German Bundestag, covering a time period from Oct 24, 2017 to May 19, 2021.[5] The corpus includes speeches by 777 different speakers, with over 330,000 sentences and over 16,5 mio tokens.

For our previous work on disambiguating 1st person plural pronouns (1PPL) (Rehbein et al., 2021), we randomly sampled a subset of the data, with roughly the same number of speeches/tokens for each party (see Table 4). This resulted in a testset with 36 speeches by different speakers (52,027 tokens) where all instances of 1st person plural pronouns (*wir, uns, unser / we, us, our*) were disambiguated manually by two coders.

In this work, we use the same corpus and add annotations for 1st, 2nd and 3rd person singular and plural pronouns in the data, using our compositional clusivity features schema. We describe our annotations in more detail below (see Section 3.1).

3.1. Annotating clusivity in parliamentary debates

We encode clusivity features for all personal pronouns in our data. The coders' task consisted in annotating three features for each pronoun in context, namely whether the pronoun includes/excludes reference to

(i) the **S**peaker ($S^{+/-}$)

(ii) the **A**ddressee ($A^{+/-}$)

(iii) **O**ther parties ($O^{+/-}$).

Table 1 provides examples for each feature combination. Our compositional schema results in 8 complex labels ($S^{\pm}A^{\pm}O^{\pm}$). We introduce a simplex 9th label, NONE, for pronoun referents that are neither persons nor institutions/organisations and can thus not fill the slot of an addressee (see the examples in Table 1).

Annotation procedure The annotators, two computational linguists,[6] were presented with the speech texts where all instances of pronouns have been highlighted for annotation. We used INCEpTION (Klie et al., 2018) as annotaton tool (see Figure 1). Unambiguous instances of pronouns that can only take one particular clusivity feature combination were excluded from the manual coding (and IAA computation) and have been annotated using a rule-based approach. This includes the pronoun forms *ich (I)* and its morphological variants (*mich/mir/mein/... (me/me/my/...)* which always has the clusivity feature set $S^{+}A^{-}O^{-}$. Other singular forms are ambiguous. For instance, forms of *du (you)* can take a generic reading (Example 3.1, contextual meaning: *there is/there exists*) and 2 of the 4 instances in our corpus are generic uses of *you*.[7]

Example 3.1. ... because otherwise people will say, "**You** only have a few countries to go to."

The more frequent form *er (he)* and its morphological variants, however, are easy to disambiguate for humans and have thus been annotated by one coder only. Please note that the German pronoun *sie* is ambiguous between a) 3.ps.sg.fem (she/it), b) the honorific address for 2.ps.sg (you), in which use it is normally capitalized

[5]The data is available in XML format from `https://www.bundestag.de/services/opendata`.

[6]The authors of this paper.

[7]We consider those instances to be of class $S^{+}A^{+}O^{+}$ as their existential meaning includes everybody.

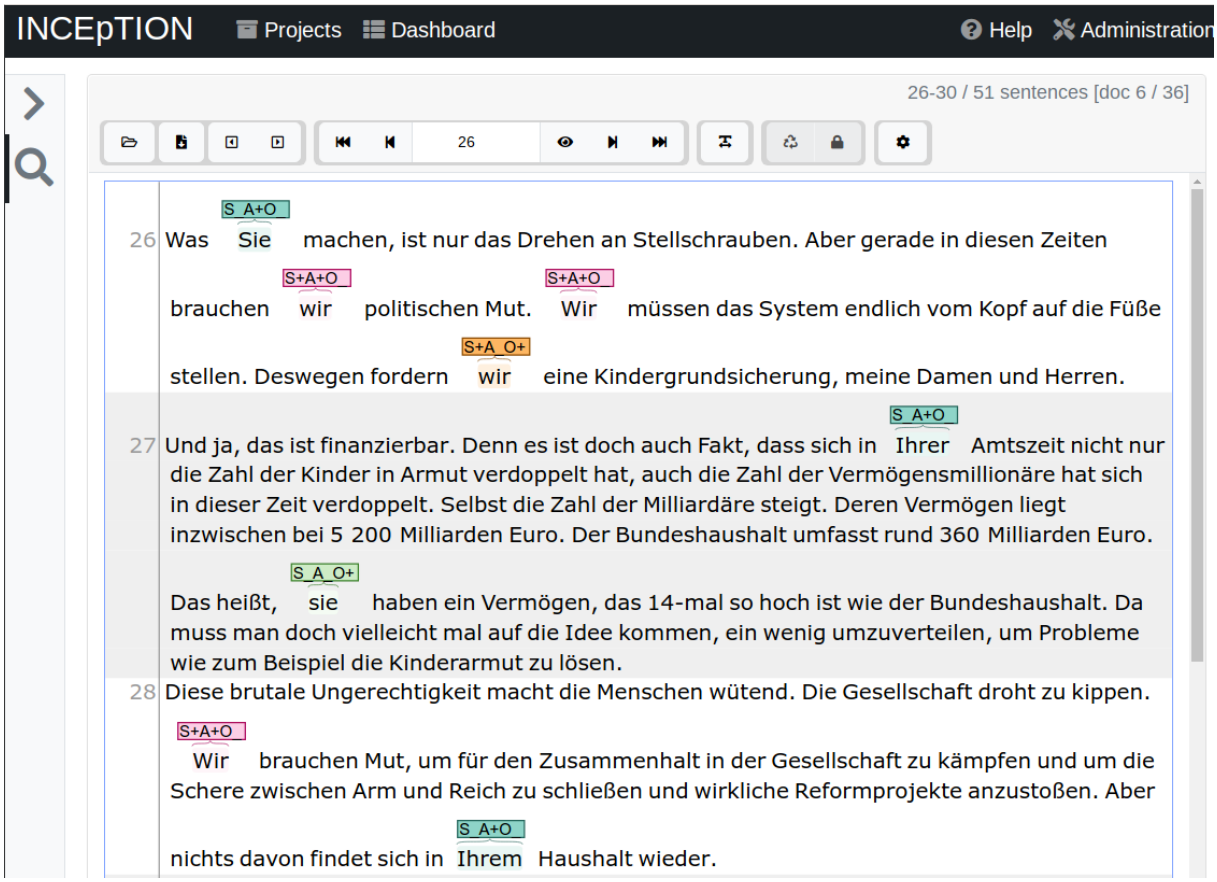

Figure 1: Screenshot of the annotation process and tool (INCEpTION).

26	What you are doing is just making minor adjustments. But it is precisely in these times that we need political courage. We must finally turn the system upside down. That's why we're calling for a basic child allowance, ladies and gentlemen.
27	And yes, it can be financed. After all, it is also a fact that not only the number of children living in poverty has doubled during your term of office, but the number of millionaires has also doubled during this time. Even the number of billionaires is rising. Their assets now total 5,200 billion euros. The federal budget is around 360 billion euros. That means they have assets that are 14 times as high as the federal budget. So maybe we should think about redistributing a little in order to solve problems such as child poverty.
28	This brutal injustice makes people angry. Society is threatening to topple over. We need courage to fight for cohesion in society and to close the gap between rich and poor and initiate real reform projects. But none of this is reflected in your budget.

Table 2: English glosses for the debate shown in the screenshot above (Figure 1).

in standard German orthography, and c) 3.ps.pl (they). Due to its homonymy, *sie (she)* has been included in the annotation of the more ambiguous plural forms.

IAA for clusivity features We report Cohen's κ and percentage agreement for two annotators on a subset of our corpus, consisting of 13 speeches with 18,947 tokens.[8] Table 3 shows IAA for the label combinations from our compositional clusivity schema. We report results separately for 1. and 2./3.ps.pl pronouns where we see a division of labour, i.e., 1PPL include the speaker per definition while 2/3PPL always exclude the speaker. Agreement for both subsets is very high with a percentage agreement >96% (κ: 93.2 for 1PPL and 89.7 for 2/3PPL).

As expected, uses of PPLs that neither include the speaker, the addresse nor any other party are implausible ($S^{-}A^{-}O^{-}$) and do not occur in our corpus, and the same goes for uses of plural pronouns that refer to the speaker only ($S^{+}A^{-}O^{-}$).

[8] Unambiguous singular pronoun forms have been excluded from the evaluation.

	1PPL	freq (%)	2/3PPL	freq (%)
$S^+A^+O^+$	97.8	58.2	0	0
$S^+A^+O^-$	98.8	24.2	0	0
$S^+A^-O^+$	89.6	17.5	0	0
$S^+A^-O^-$	0	0	0	0
$S^-A^+O^+$	0	0	100.0	0.3
$S^-A^+O^-$	0	0	98.4	76.5
$S^-A^-O^+$	0	0	93.1	17.3
$S^-A^-O^-$	0	0	0	0
NONE	0	0	87.2	5.9
Total (%)	96.1	100.0	96.0	100.0

Table 3: IAA (F1) and support (percentage of annotated instances) for the double-annotated sample (734 instances) of 1st, 2nd and 3rd person plural pronouns.

Instances that include addressee and others but not the speaker were rare in our data (see Example 3.2). Those cases, however, are not implausible and we expect to find more examples in a larger corpus.

Example 3.2. Sie [die YPG] kämpfen an der Seite der US-Amerikaner, und **Sie** [die deutsche Regierung] unterstützen das.
They [the YPG] are fighting alongside the U.S. and **you** [the German government] support that.

Adding singular pronoun senses Encouraged by the high inter-annotator agreement, we completed the annotation of the corpus, using a single annotator. The annotator also marked the missing clusivity features for singular personal and posessive pronouns (*I/my/me, you/your/you, she/her/her, he/his/him*). We exclude the neutral pronoun *it* as it is rarely used to refer to political actors. As described above, annotations for 1st person singular (1PSG) pronouns are not ambiguous and have been assigned automatically. Clusivity features for 2nd and 3rd PSG pronouns have been resolved by the human coder(s).
After the annotation was completed, the two annotators discussed and resolved all disagreements in the double-annotated subset to create a ground truth dataset that we use as evaluation data in our experiments (Section 4.2).

Limitations As we only had access to a restricted pool of two annotators, we were not able to systematically control for cultural or political background. This is suboptimal as annotators with different backgrounds might interpret and resolve underspecified references in different ways. For example, an annotator with positive stance towards a liberal immigration policy might be biased to interpret a (hypothetical) statement by the German chancellor "We welcome all refugees" as a reference to WE, THE COUNTRY while a person with opposite stance is more likely to resolve this as a reference to WE, THE GOVERNMENT. It would be interesting to investigate this issue in more depth and we would expect to see a systematic correlation between clusivity interpretation and political background.

A second limitation is the treatment of only one language in one political setting. In principle, our annotation scheme should apply to debates in all languages that have not grammaticalized an inclusive-exclusive distinction and whose 1PPL pronouns therefore need to be interpreted in context for this feature. However, we would not expect all languages to behave the same: language specificities are likely to lead at least to different frequency distributions. For example, inspection of how instances of German *wir* are translated into English in the context of the European parliament suggests that quite a few generic instances are unexpressed in the translation because the English version presents a state of affairs as a bald fact (*There is a crisis*) rather than something that everybody is experiencing (*We have a crisis*). On a more practical level, our pronoun-focused approach would have to be reframed in terms of 1PPL verbal forms for languages such as Italian where pro-drop is widely available (see Example 3.3 below).

Example 3.3. *Parliamo* di problemi creati in ventitrent'anni di mala gestione politica. (ParlaMint-IT_2019-12-11-LEG18-Sed-172.u9)
'We are talking about problems created in twenty-thirty years of political mismanagement.'

The specifics of the political institutions involved are likely also relevant because different types of parliamentary interactions – for instance, plenary debates versus Prime Minister's Question Time – are likely to feature different distributions of inclusive and exclusive references, involving different in- and out-groups. From that perspective, it could be worthwhile to, for instance, extend the annotation of German data to debates in the Austrian and Swiss parliaments.

3.2. Data exploration

We now present an exploration of our annotated data set where we look out for systematic patterns in the use of pronouns in parliamentary debates.

Inter-party differences in the use of pronouns Table 4 shows the distribution of 1st, 2nd and 3rd person pronouns for speakers from different parties. Most speakers use far more 1st person than 2nd person pronouns and more 2nd person than 3rd person pronouns. Members of the government (consisting at the time of data collection of the centre-right Christian democratic union parties CDU/CSU and the social democratic SPD) show a higher than average ratio of pronouns in general (see Table 4) which can be traced back to a higher use of 1st person pronouns and, in particular, 1PPL (not shown here). At the same time, their use of 2nd person pronouns is much lower than the one for members of the opposition. Intuitively, this makes sense as members of the government often talk about what they have planned or achieved (thus a high use of 1st person pronouns) while members of the opposition focus on criticising or attacking the government policies which is reflected in a higher use of 2nd person pronouns.

Party	# Speakers	# Sent	# Tok	1.ps %	2.ps %	3.ps %	freq	per 1000
CDU/CSU	5	590	10,674	75.6	16.1	8.3	640	60.0
SPD	4	462	7,438	74.6	12.0	13.3	465	62.5
FDP	7	491	7,358	62.7	25.4	11.8	397	53.9
GRUENE	5	451	7,457	58.1	31.1	10.8	418	56.0
LINKE	6	586	9,310	51.4	31.6	17.0	424	45.5
AfD	8	546	8,993	46.7	35.3	18.0	490	54.5
unaffiliated	1	15	797	23.8	69.0	7.1	42	52.7
TOTAL/AVG	36	3,126	52,027	62.3	25.4	12.3	2,876	55.0

Table 4: Distribution of 1st, 2nd and 3rd person pronouns per party and number of annotated pronouns (unaffiliated: speaker was a member of the Alternative for Germany (AfD) until 2017).

	S-A+O-	S+A-O-	S+A+O+	S+A+O-	S+A-O+	S-A-O+
CDU	15.8	23.3	25.9	8.4	18.4	6.4
SPD	12.0	21.9	24.5	13.3	14.8	9.5
FDP	25.2	20.9	19.9	11.1	10.8	9.6
GRUENE	32.1	25.8	15.5	13.9	3.3	5.7
LINKE	31.4	20.7	18.6	8.0	4.0	14.6
AfD	35.3	17.8	12.6	9.4	7.1	12.6
AVG	25.3	21.7	19.5	10.7	9.7	9.7

Table 5: Distribution of clusivity features across parties (%).

Inter-party differences regarding clusivity We now focus on the more fine-grained patterns of inclusive and exclusive meanings of pronouns for the different parties (Table 5). Please note that the numbers do not add up to 100% as we did not consider the NONE instances in the analysis. As before, we can see crucial differences between members of the government and the opposition. Members of the government show

- a more frequent use of $S^+A^+O^+$ (we, the country; generic uses) and $S^+A^-O^+$ (we, the government)
- a less frequent use of $S^-A^+O^-$ (exclusive address of specific persons/groups/parties)

We also observe an interesting overuse of the label $S^-A^-O^+$ for members of the left-wing party *Die Linke* and of the extreme right, *AfD* (mostly 3rd person references to specific persons or instances of the people; see Examples 3.4 and 3.5.

Example 3.4. Wenn Menschen **ihren** Job verlieren ...
When people lose **their** jobs ... (*AfD*)

Example 3.5. dass Arbeitslose im Jobcenter schikaniert werden, als würden **sie** um Almosen betteln
that the unemployed are harassed at the employment agency as if **they** were begging for alms (*Die Linke*)

The higher ratio of references to instances of the people could be interpreted as an indicator for *people-centrism*, often described in the literature as one of the crucial features of populism (see, e.g., Mudde and Kaltwasser (2017; Van Leeuwen (2019; Wirth et al. (2016)).

4. Experiments

In this section, we explore how well we can predict the clusivity of personal pronouns in parliamentary debates.

4.1. Exerimental setup

In order to investigate the use of personal pronouns as a rhetorical device on a larger scale, e.g., from a diachronic perspective, we need to be able to automatically predict the inherent inclusiveness or exclusiveness with sufficient accuracy. We now use our new data set to test how well we can predict clusivity features for pronouns in political debates.
To increase to interpretability of our models, we start with a conventional feature-based approach. We first focus on the question what type of information might help to solve the task.

Baseline Our baseline model is a linear SVM[9] that uses bag-of-words tf-idf features with a context size of 20 tokens to the left and 20 tokens to the right of the pronoun. Our input units are paragraphs from the speeches and we extract one instance for each pronoun that we want to disambiguate. Please note that we do not extract context features across paragraphs.
We train the classifier in a 6-fold cross-validation setting where we divide the different speeches into 6 folds so that a) each fold has roughly the same number of tokens and b) that the speeches from different parties are fairly

[9]We also experimented with other classifiers, i.e., SGD, randomForest, Ridge regression, random forests and decision trees but got best results for the LinearSVC classifier from scikit-learn: `https://scikit-learn.org`.

evenly distributed over the different folds. We do not split up the speeches to make sure that each test set includes only instances from speakers that have not been seen during training. This might make it harder for the classifier but also provides a more realistic test bed and makes it harder for the classifier to pick up on speaker-specific or topic-specific features.

Feature extraction To determine the optimal maximum number of features for each fold, we use χ^2 to select the highest-scoring features on a development set that we randomly extract from the training data for this fold.

For all models, we encode the normalised form of the pronoun as an additional feature.[10] As we have seen a strong correlation between party membership and the distribution of clusivity features in our data (see Subsection 3.2 and Table 4), we hypothesize that meta-information on party affiliation will provide useful information for the classifier. To test this, we encode two additional features, (i) party membership (PARTY) and (ii) whether an instance was produced by a member of the government or of the opposition (IS_GOV). The first feature can take 7 values, i.e., the party names, while our second feature is binary (1|0).

Another feature we test encodes the dependency subtree of the pronoun we want to disambiguate (Figure 2).

German text:
Da hätten Sie Herrn Erdogan klare Kante geben und sagen müssen : Wir , und zwar nicht nur **wir** Deutsche , sondern auch die europäischen NATO-Nationen [...]

English translation:
'You should have taken a firm stand and told Mr. Erdogan quite clearly: We, and not only **we** Germans, but also the European NATO nations [...]'

Dependency-based features:

	pron-form	**GF**	**head**	*head POS*	**child nodes**
DE:	wir	NK	deutschen	NOUN	–
EN:	*we*	NK	*Germans*	NOUN	–

Figure 2: Example for extraction of dependency features (NK: noun kernel).

We use spaCy v3.2.0 (Honnibal and Montani, 2017) with the `de_core_news_sm` model for all preprocessing steps, including tokenisation, lemmatisation and dependency parsing. We extract the normalised pronoun form for each instance, its grammatical function, its head lemma form und POS tag and the lowercased word forms for the head's child nodes, if present, and concatenate this information. The motivation for this feature is that we want to help the classifier to focus on the more relevant bits of information in the context.

We determine the best hyperparameters for each fold, based on a grid search, again extracting a random development set from the training data for this fold.[11]

BERT-based model We compare the results from our feature-based classification with a transformer-based text classifier (Devlin et al., 2019), initialised with pretrained contextual embeddings (`bert-base-german-cased`) provided by the Huggingface library (Wolf et al., 2020). Input to the model are the text segments where the respective pronouns are marked with a preceding underscore. We fine-tune the model on the same folds that we used in our previous experiments and perform model selection, based on eary stopping. Specifically, we use 10% of the training data as a validation set and stop training when the evaluation loss on the validation set stops decreasing by at least 0.01 for 2 consecutive evaluations. Given that our training set is rather small, we decided on a small patience value to prevent overfitting.

We then used the model that gave best results on the development set to predict clusivity features for the test set. Due to the small size of our data, we refrain from further hyperparameter tuning but stick to the default parameters.[12]

4.2. Classification results

Baseline Table 6 shows results (micro F1) for our baseline SVM, for different feature combinations. We can see that the dependency-based features alone outperform the tf-idf features by more than 2% while the combination of both yields a further small improvement, up to 78% F1. Surprisingly, the addition of meta-information on party membership and participation in government fails to give further improvements.

While the overall score of $> 78\%$ seems satisfying for such a difficult task, we can see that the performance for different classes varies considerably. Two of the label combinations obtain near perfect scores: $S^+A^-O^-$ and $S^-A^+O^-$. This is to be expected as the first class, $S^+A^-O^-$, mostly consists of the 1st person pronoun *ich (I)* which is not ambiguous. The high score for the label $S^-A^+O^-$ can be explained by the many instances of the 2nd person honorific address *Sie* (you). While the pronoun form is ambiguous between a) 3.ps.sg.fem (she/it), b) the honorific address for 2.ps.sg (you) and c) 3.ps.pl (they), the majority of instances of *Sie* (95%) are, in fact, instances of label $S^-A^+O^-$.

[10] We use the lowercased form and trace morphological variants (e.g., *unsere, unsre, unser, unseres*) back to a normalised form (*unser*).

[11] The hyperparameters we use for tuning are: loss: hinge, squared_hinge; penalty: l1, l2; multi_class: ovr, crammer_singer; max_iter: 500, 1000, 1500.

[12] We use the Simpletransformers, an easy-to-use wrapper for the Huggingface transformers library.

		Baseline: SVM classifier						*BERT*			
Class	**# gold**	**tfidf**	**dep**	**tfidf+dep**	**+party**	**+gov**	**all**	**prec**	**rec**	**F1**	**(stdev)**
S+A+O+	569	62.1	65.8	66.3	64.2	66.4	64.0	73.9	77.8	75.9	(0.12)
S+A+O-	301	37.4	45.0	51.1	49.0	50.3	49.7	63.0	50.8	56.2	(0.55)
S+A-O+	298	34.1	41.9	43.1	46.6	40.8	46.6	61.2	66.3	63.6	(0.06)
S+A-O-	617	99.7	99.8	99.7	99.6	99.7	99.6	100.0	99.8	99.9	(0.00)
S-A+O+	1	0	0	0	0	0	0	0	0	0	(0.00)
S-A+O-	726	96.7	96.6	96.5	96.5	96.5	96.5	97.6	98.6	98.1	(0.00)
S-A-O+	273	82.1	82.8	82.2	82.1	81.9	82.4	91.0	89.3	90.2	(0.25)
NONE	91	62.1	40.3	44.9	45.3	42.6	47.5	77.7	80.6	79.1	(1.36)
Total	2,876	75.6	77.9	78.5	78.0	78.2	78.1			84.9	

Table 6: Results (micro F1) for a 6-fold cross-validation with linear SVM (# gold: no. of ground truth instances per class) and results (micro **prec**ision, **rec**all, **F1**) averaged over 3 runs (with standard deviation in parentheses), for a 6-fold cross-validation with BERT. The highlighted columns show best baseline and BERT results (both micro F1).

	wir				**sie**				**uns**			
	# gold	P	R	F1	# gold	P	R	F1	# gold	P	R	F1
S+A+O+	368	69.5	72.0	70.8	–	–	–	–	70	67.9	75.7	71.6
S+A+O-	246	63.8	53.7	58.3	–	–	–	–	47	55.8	51.1	53.3
S+A-O+	225	60.8	67.6	64.0	1	0	0	0	48	65.9	60.4	63.0
S-A+O-	–	–	–	–	502	97.8	99.6	98.7	–	–	–	–
S-A-O+	–	–	–	–	110	92.6	80.0	85.8	–	–	–	–
NONE	–	–	–	–	53	75.0	84.9	79.6	–	–	–	–
Total	839	**Acc = 65.4**			666	**Acc = 95.0**			165	**Acc = 64.2**		

Table 7: Results (micro **P**recision, **R**ecall, **F1**) for individual pronoun forms for a 6-fold cross-validation with BERT.

Prediction with BERT We now compare our feature-based SVM to a state-of-the-art pretrained language model (Devlin et al., 2019). We report averaged results over 3 runs, with standard deviation in round brackets (Table 6).

The transformer-based model outperforms our feature-based models significantly, with an F1 close to 85%. The improvements are distributed over the different classes: for the inclusive label, $S^+A^+O^+$, we see an increase of 10% F1; for uses that include the speaker and others but exclude the addressee, $S^+A^-O^+$, results improve by around 20%; and for the identification of non-relevant instances, NONE, BERT outperforms our baseline by >30% F1.

Table 7 presents results for the three most frequent pronouns in our data: *wir (we), sie (they/you/she)* and *uns (us)*. Again, we can see that most of the ambiguity concerns 1PPL. This is also confirmed by the confusion matrix (see Appendix) which reveals that our models mostly struggle with the speaker-inclusive labels (i.e., 1PPL). Many of the misclassified examples are inherently ambiguous, which is reflected in the lower human agreement for 1PPL pronouns (also see Rehbein et al. (2021)). The examples below (translated into English) illustrate the three most frequent error types, showing the inherent ambiguity between references to the parliament ($S^+A^+O^-$), to the whole country ($S^+A^+O^+$) or to one specific party ($S^+A^-O^+$).

Example 4.1. But I also think that **we** must accept our international responsibility and therefore come to terms with the issue of the deployment of the Bundeswehr in Afghanistan. GOLD: S+A+O-; PRED: S+A+O+

Example 4.2. Yes , **we** lament that in the past we were one of the main destination countries of illegal migration. GOLD: S+A+O+; PRED: S+A-O+

Example 4.3. It will be **our** task to revive freedom, responsibility and competition as principles of our social and economic order. GOLD: S+A-O+; PRED: S+A+O+

5. Conclusions

In the paper, we presented a compositional schema for annotating the clusivity properties of personal pronouns in context. We applied our schema to German parliamentary debates and showed that human annotators can achieve high IAA for this task. In addition, our data exploration presented some evidence for the face validity of our annotations. We then presented first experiments towards the automatic prediction of clusivity in political text. While the overall results were promising, a closer look revealed that especially 1PPL pronouns pose a challenge for automatic disambiguation. This is in line with previous insights from qualitative and empirical studies on the interface between corpus linguistics, pragmatics, discourse studies and political science. In future work, we plan to improve the classifier to be able to apply it to large-scale studies of the use of personal pronouns beyond word frequencies.

6. Acknowledgements

This work was supported in part by the SFB 884 on the Political Economy of Reforms at the University of Mannheim (projects B6 and C4), funded by the German Research Foundation (DFG).

7. Bibliographical References

Alavidze, M. (2017). The use of pronouns in political discourse. *International Journal of Arts & Sciences*, 9(4):349–356.

Allen, W. (2007). Australian political discourse: Pronominal choice in campaign speeches. In *Selected Papers from the 2006 Conference of the Australian Linguistic Society, ed. by Mushin Ilana and Mary Laughren*.

Bickel, B. and Nichols, J. (2005). Inclusive/exclusive as person vs. number categories worldwide. In E. Filimonova, editor, *Clusivity: Typology and Case Studies of Inclusive-exclusive Distinction*, chapter 2, pages 49–72. John Benjamins Publishing Company.

Bull, P. and Fetzer, A. (2006). Who are we and who are you? the strategic use of forms of address in political interviews. *Text and Talk*, 26, 01.

Cysouw, M. (2002). The impact of an inclusive/exclusive opposition on the paradigmatic structure of person marking. In *Pronouns: Grammar and Representation*, pages 41–62. John Benjamins Publishing.

Dam, L. (2015). The functionality of personal pronouns in constructions of communities. *Globe: A Journal of Language, Culture and Communication*, 31-42(1).

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.

Filimonova, E. (2005). *Clusivity: Typology and Case Studies of Inclusive-exclusive Distinction*. Typological studies in language. J. Benjamins Pub.

Honnibal, M. and Montani, I. (2017). spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.

Íñigo-Mora, I. (2004). On the use of the personal pronoun we in communities. *Journal of Language and Politics*, 3(1):27–52.

Klie, J.-C., Bugert, M., Boullosa, B., de Castilho, R. E., and Gurevych, I. (2018). The inception platform: Machine-assisted and knowledge-oriented interactive annotation. In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 5–9. Association for Computational Linguistics, June.

Mudde, C. and Kaltwasser, C. R. (2017). *Populism: A very short introduction*. Oxford, UK: Oxford University Press.

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

Proctor, K. and Su, L. I.-W. (2011). The 1st person plural in political discourse – American politicians in interviews and in a debate. *Journal of Pragmatics*, 43(13):3251–3266.

Quirk, R., Greenbaum, S., Leech, G., and Svartvik, J. (1985). *A Comprehensive Grammar of the English Language*. Longman, London.

Rehbein, I., Ruppenhofer, J., and Bernauer, J. (2021). Who is we? disambiguating the referents of first person plural pronouns in parliamentary debates. In *Proceedings of the 17th Conference on Natural Language Processing (KONVENS 2021)*, pages 147–158, Düsseldorf, Germany, 6–9 Sep. KONVENS 2021 Organizers.

Simon, H. J. (2005). Only you? philological investigations into the alleged inclusive-exclusive-distinction in the second person plural. In Elena Filimonova, editor, *Clusivity. Typology and case studies of the inclusive-exclusive distinction*. Benjamins, Amsterdam/Philadelphia.

Tyrkkö, J. (2016). Looking for rhetorical thresholds: Pronoun frequencies in political speeches. *Studies in Variation, Contacts and Change in English*, 17.

Van Leeuwen, M. (2019). Measuring people-centrism in populist political discourse: A linguistic approach. In Jan Zienkowski et al., editors, *Imagining the Peoples of Europe, Populist discourses across the political spectrum*, pages 315–340. Amsterdam/Philadelphia: John Benjamins.

Vuković, M. (2012). Positioning in pre-prepared and spontaneous parliamentary discourse: Choice of person in the parliament of montenegro. *Discourse & Society*, 23(2):184–202.

Wieczorek, A. (2009). This is to say you're either in or out: Some remarks on clusivity. *Critical Approaches to Discourse Analysis across Disciplines*, 3(2). 118-129.

Wirth, W., Esser, F., Wettstein, M., Engesser, S., Wirz, D., Schulz, A., Ernst, N., Büchel, F., Caramani, D., Manucci, L., Steenbergen, M. R., Bernhard, L., Weber, E., Hänggli, R., Dalmus, C., Schemer, C., and Müller, P. (2016). The appeal of populist ideas, strategies, and styles: A theoretical model and research design for analyzing populist political communication. Technical report, University of Zurich, Zurich.

Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., Drame, M., Lhoest, Q., and Rush, A. (2020). Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical*

Methods in Natural Language Processing: System Demonstrations, pages 38–45, Online, October. Association for Computational Linguistics.

8. Appendix

Gold \ Pred	NONE	S+A+O+	S+A+O-	S+A-O+	S+A-O-	S-A+O+	S-A+O-	S-A-O+
NONE	74	0	0	0	0	0	5	12
S+A+O+	1	441	56	69	0	0	2	0
S+A+O-	0	91	156	54	0	0	0	0
S+A-O+	0	61	38	197	0	0	0	2
S+A-O-	0	0	0	0	616	0	0	1
S-A+O+	0	0	0	0	0	0	1	0
S-A+O-	1	1	0	0	0	0	717	8
S-A-O+	19	1	0	0	0	0	10	243

Table 8: Confusion matrix for the prediction of clusivity features for pronouns in parliamentary debates (BERT).