Between Cows and Capitalism: Measuring the Abstractness of Historical Parliamentary Speeches

Ruben Ros[1]

[1]*Centre for Contemporary and Digital History, Luxembourg University*

Abstract

This paper proposes a method for estimating the abstractness of Dutch historical parliamentary speeches. The degree to which language is abstract or concrete has gained significant attention in the field of cognitive science and linguistics. The paper uses proposed computational approaches to abstractness in these fields to model and explore variation in historical parliamentary data in the context of a larger project on technocratic rhetoric. By first scoring individual terms based on the vector averaging of an annotated set of abstract and concrete terms the abstractness of speech paragraphs is estimated. The paper shows that this information captures differences in rhetorical style that remain difficult to identify with established text mining methods. Abstractness sheds light not only on what is said in parliament, but how it is said, hereby bridging semantic and stylistic analysis in digital history.

Keywords

abstractness, political rhetoric, parliamentary history, conceptual history

1. Introduction

Parliamentary debate features both abstract and concrete forms of language. Abstract ideological rhetoric alternates with practical deliberation over detailed policy measures, and jargon-saturated legislative language appears alongside concrete calls to action by parliamentary representatives. This dimension of abstractness - defined as the opposite of concreteness - ties in with the so-called "scientization" of democracy and parliament [1, 2, 3]. In the eyes of for example Jürgen Habermas,the growing role of the sciences in politics leads to the replacement of democratic deliberation with technocratic calculation [4]. Parliamentary debate becomes a form of technical discussion under the pressure of waning ideologies, increasingly complex legislation and the growing dominance of expert knowledge and institutions [5].

There are several indicators that point to abstractness as a linguistic expression of these historical processes of scientization and technocratization. Technocratic rhetoric is considered in the literature as drawing heavily on abstract (bureaucratic and managerial) language and historical studies point at the rise of more technical debates that are saturated with legislative jargon [6, 7, 8]. At the same time, technocrats themselves often self-identify as "men of practice", shying away from (abstract) ideological politics, and instead focusing on (concrete) depoliticized technical management [9, 10].

Digital Parliamentary Data in Action (DiPaDA 2022) workshop, Uppsala, Sweden, March 15, 2022.

ruben.ros@uni.lu (R. Ros)

0000-0002-5303-2861 (R. Ros)

This paper describes a method for measuring abstractness in parliamentary debate. Drawing on work in computational linguistics, it models abstractness in parliamentary speech paragraphs and tries to use these abstractness signals to find forms of language variation that map onto the issue of technocratization and technocraticness.

Abstraction has recently gained significant traction in the fields of psychology and cognitive science. [11, 12]. Understood as a core component of human cognition, it is generally studied through language. Our capacity for abstract reasoning reflects in the concepts we use. Consequentially, the topic of abstraction has been taken up in (computational) linguistics, where several contributions have tried to estimate the abstractness of individual words alongside other metrics, such as age of acquisition and imagery [13, 14]. This, in turn, has led to applications of abstractness-oriented work in sociolinguistics. In this field, lexical abstractness has been used in research into linguistic intergroup bias. Borden et al. [15], for example, show an increase in abstract language in times of crisis, while Dragojevic et al. [16] use abstractness to demonstrate intersubjectivity in news coverage. These examples show that the linguistic study of abstractness is relevant to other fields that use language difference as the basis for humanistic research. Variation in lexical abstractness on the level of groups, genres or time periods can be used as a basis for studying history, as the recent work in the field of Computational Literary Studies has shown[17].

The paper puts forward a method for estimating parliamentary speech abstractness. It shows how abstractness can be estimated on the word level using word embeddings and an annotated set of terms. The paper explains the way in which word-level estimates are subsequently used to estimate parliamentary speech abstractness. It also features a brief analytical part that concentrates on (explanations for) speaker-level differences, functioning as an example of the type of analysis that could be performed with the method outlined in the paper. Lastly, several shortcomings and venues for further research are discussed.

2. Data

The paper uses a corpus of Dutch (Lower House) parliamentary debates from the period between 1917 and 1986 [18]. The abstractness estimation method is performed on the full corpus. The analytical section focuses on the period between 1917-1923. In total, the data contains 194415 paragraphs, a total of 13.979.075 tokens and 194.415 types. The text data was preprocessed by lowercasing tokens, removing tokens with multiple digits and removing stop words. For the topic modelling, a lemmatized version of the text data was used.

The data contains rich metadata that connect speech paragraphs to speakers, parties, roles (government minister, MP or guest speaker) and dates. This enables the investigation of various potential explanations for abstract language.

The quality of Optical Character Recognition, often a restraining factor in text analysis, is hard to measure, since no accuracy scores are available. Nevertheless, the significant improvement observed in the text that has been subjected to OCR post-correction using PICCL (compared to the original digitized proceedings) indicates a relatively high quality [19]. Type-token ratio's, used as a crude measure for OCR quality, suggest significant improvement over the course of the century [20].

3. Method

3.1. Estimating Word Abstractness

Previous work in computational linguistics uses manually scored vocabularies to create abstractness classification models [21, 22, 23]. The lexical context of known abstract and concrete terms is used to predict the abstractness of unseen terms, based on the assumption that abstract and concrete terms have similar contexts [24, 25]. The resulting word-level scores are then used to score larger units of text by means of counting (an extended list of) paradigmatic abstract and concrete terms in a text, or averaging the word scores in the text [17].

The weakness of this method is the labour-intensive process of creating annotated data. However, several annotated datasets are available. This paper relies on an extensive list of word concreteness scores produced in a study by Brysbaert et al. (hereafter BBT) [26]. This data contains concreteness ratings for over 30.000 Dutch words rated through the *Amazon Mechanical Turk* platform. Compared to other sets, the BBT set covers a high number of words, is scored by a multiple annotators is and one of the few that focuses on Dutch and does not require translation [27, 28]. According to Brysbaert et al., annotators were asked to rate a concept's concreteness on a five point scale based on the extent to which a word can be experienced directly through one of the senses.

We use the BBT data to score individual terms in a series of steps. First, we train word embedding models on the parliamentary proceedings. We use the Skipgram Negative Sampling (SGNS) variant of the Word2Vec algorithm implemented in the popular Gensim module in Python [29]. Based on the work by Rabinovich et al. [25], we use a relatively small window of three (n-grams) in the training. Moreover, words with a frequency of lower than one hundred are excluded. We are aware that this potentially excludes terms of interest, but in light of the volatility of low-frequency terms in the embedding models, we deem this choice justified. We train models on both the full period and six-year time slices to take historical language variation into account. Because the slices contain relatively little data in early time periods and individual word vectors tend to vary in different models trained on the same data, we use bootstrapped training and take the average word vectors over ten models for every term in all subsequent steps [30].

As shown in earlier studies, the word embedding models allow the construction of an average abstractness vector [17]. This is done by averaging the vectors of a list of known abstract and concrete terms. For this study, we use only adjectives and nouns. Experiments show that including other word forms, such as verbs and conjunctions, hinder the effectivity and explainability of our method. Another important factor is the number of seed terms used for training. It is not clear from the outset how many of the top abstract and concrete terms from the BBT-data should be used. Figure 1 shows the correlation between annotated and estimated scores plotted against different numbers of seed terms. This pattern shows that the correlation rapidly increases when taking a higher number of terms, but eventually decreases when we take too many. After around 480 terms, the correlation stabilizes, meaning that adding more seed terms to the average vector does not amount to a higher correlation between estimated and annotated abstractness scores. Based on this distribution, we take 480 abstract and concrete nouns and adjectives (meaning 220 concrete nouns, 220 concrete adjectives, 220 abstract nouns

and 220 abstract adjectives) as the basis for the averaging.

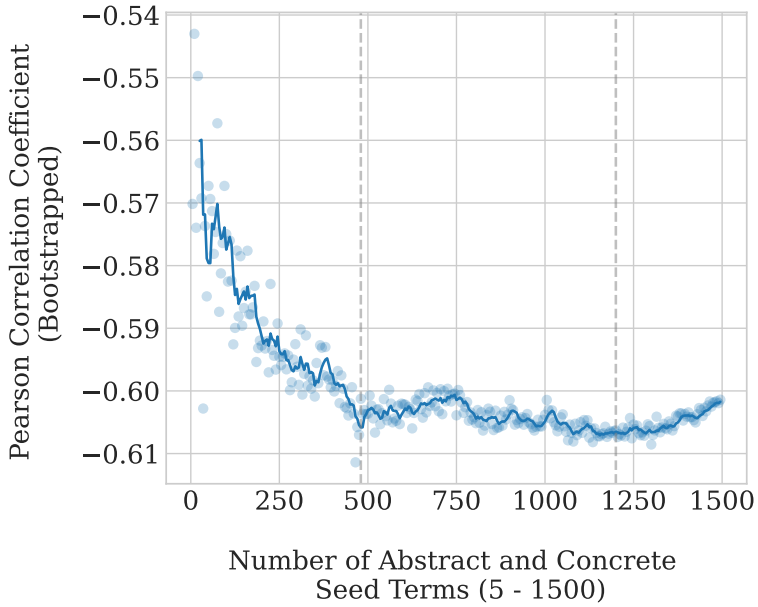

Figure 1: Bootstrapped Pearson's correlation coefficients against increasing number of seed terms. Vertical lines show stabilization at around 480 seed terms and a decline in correlation coefficients after around 1200 terms.

To estimate term abstractness, we calculate the mean vectors of the abstract and concrete terms and subtracted the latter from the former. The cosine similarity of a word vector to this average abstractness vector is then used as a measure for abstractness. Figure 2 shows the individual (estimated) word abstractness scores (derived from the bootstrapped models trained on the full period) plotted against the BBT-annotations. Since this means that abstractness is offsetted against concreteness, the correlation is negative. The figure shows a relatively broad scatter and a moderate correlation of around -0.6. This means that there remains a group of words that is annotated as concrete, while estimated as abstract and vice versa. Looking at these words points at the issue of sense ambiguity. For example, in Table 1, that shows the top most concrete and abstract estimated terms in the period between 1918 and 1923, the verb "stroken" ("match") appears as a highly abstract term. In the BBT-annotations, the term is likely to have been understood as a highly concrete noun ("strips"). However, in light of the observed limited number of these ambiguous terms, we believe that this problem does not fundamentally hinder our method.

3.2. Estimating Text Abstractness

Based on the word scores generated with the seed term lists and vector averaging we calculated abstractness scores for parliamentary speech paragraphs. Looking at overall speech abstractness proved unworkable, since abstract and concrete words cancel each other out, leading to little variation. We considered scoring texts on the level of sentences. However, the quality of the

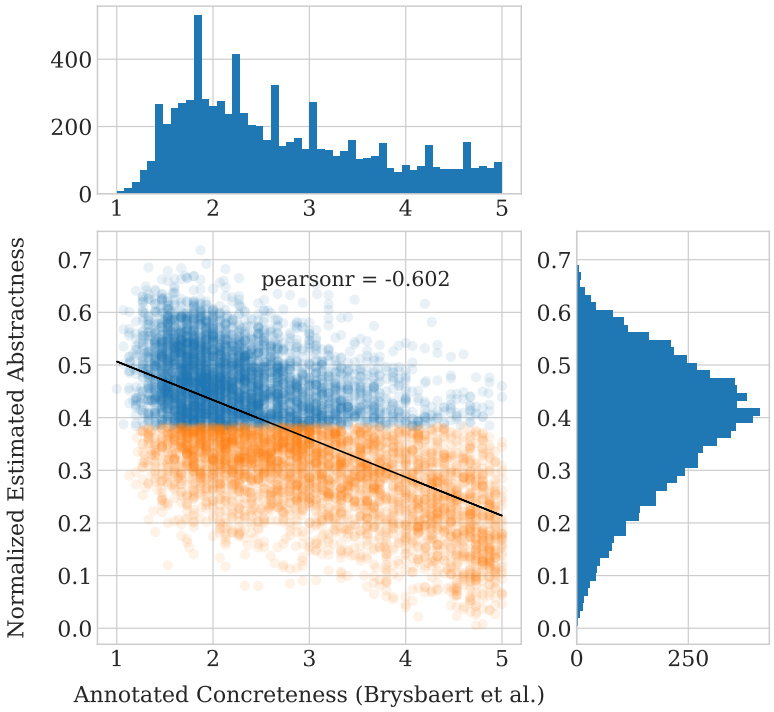

Figure 2: Scatterplot of all words in the model trained on the full period (1917-1986) that shows the correlation between estimates and annotations (Pearson's R = -0.602) and the distributions of the estimated abstractness scores and annotated concreteness scores.

	Abstract	Concrete
1	code of conduct (gedragslijn)	sheep (schaap)
2	constitutional (staatsrechtelijk)	horse (paard)
3	desideratum (desideratum)	koe (koe)
4	unanimity (eenstemmigheid)	closet (kast)
5	match (stroken)	wit (wit)
6	character (karakter)	boot (boot)
7	point of departure (uitgangspunt)	wagon (wagon)
8	rationale (grondgedachte)	indigestion (maag)
9	effect (uitwerking)	chest (kist)
10	constitutional (grondwettelijk)	neighbourhood (buurt)

Table 1
Top 10 translated most abstract and concrete terms in the period between 1917 and 1923.

sentence tokenization turned out problematic and currently the quality of the OCR does not permit reliable improvement. In future work we hope to improve the sentence tokenization in order to get a more fine-grained picture of the abstractness dynamics within a speech. In this paper we use paragraph-level abstractness. We consider this an effective intermediate level in between sentences and speeches. Paragraph abstractness scores are calculated by taking the

median of all individual word abstractness scores in the paragraph. Because the scores in the paragraphs are not distributed normally, we take the median instead of the mean.

3.3. Improving Abstractness Signals

The method described so far produces abstractness scores on the level of speech paragraphs. Given the availability of rich metadata on speakers, parties, roles and dates, this offers the opportunity for various types of historical analysis. Temporal trends, differences between groups, or abstractness variations in specific debates could be studied with the scores we have extracted. However, upon looking closely at specific paragraphs and debates, we found that two additional steps were required to improve abstractness signals.

First, close reading abstract paragraphs revealed that many of them dealt with parliamentary procedure. This comes as no surprise, given the "explicitly procedural character of parliamentary politics" described by Kari Palonen and others [31, 32]. The addressing of members, the filing of motions and amendments and standardized types of debates (such as questions to the minister or specialized budget talks) is structured by rules and customs that produce highly stratified forms of language. It appears that this procedural language consists mostly of abstract terms. To verify this suspicion, we trained a topic model with 75 topics using Mallet. The paragraph-level topic distributions could then be compared with the paragraph abstractness scores. Figure 3 shows the fifteen topics that have the highest negative and positive correlation with abstractness. The distributions of the topics at the top of the figure have a correlation coefficient of around -0.2, meaning that they (weakly) correlate with concreteness. The bottom topics have a correlation of around 0.2, demonstrating some correlation with abstractness. Their top terms clarify these patterns. It makes sense that topics about agriculture ("potato, farmer, stock") housing ("house, building, build") correlate negatively with abstractness. The figure shows that many of the topics that correlate positively with abstractness are procedural topics. They pertain to questions, motion, legislative jargon and the more general language of opposition. Topics that show a (moderate) correlation with concreteness are often about specific policy areas, such as housing, defense or agriculture.

This topic abstractness was a first surprising finding in our research. In the subsequent analysis, we used this information to filter out procedural language. By classifying topics as "semantic" or "procedural" (based on their top terms), we kept only those paragraphs where the combined topic probability for procedural topics was lower than 0.5. In this way, a better signal was acquired because the "abstractness bias" encapsulated in procedural language was largely removed. This is not to say that procedural language cannot be used for our analysis, or plays no role in the overall abstractness of debates, but for this first exploratory dive in the abstractness data, we prefered to focus on speech paragraphs that dealt with specific policy debates.

The second step in improving the signal was the calculation of the distance between the average paragraph abstractness and the average session abstractness. This follows from the observed link between topics and abstractness. Because sessions are often dedicated to only a small number of topics, local dynamics are likely to be obscured by the general abstractness of the session. For this reason, we calculated the distance from the average session abstractness and the paragraph abstractness.

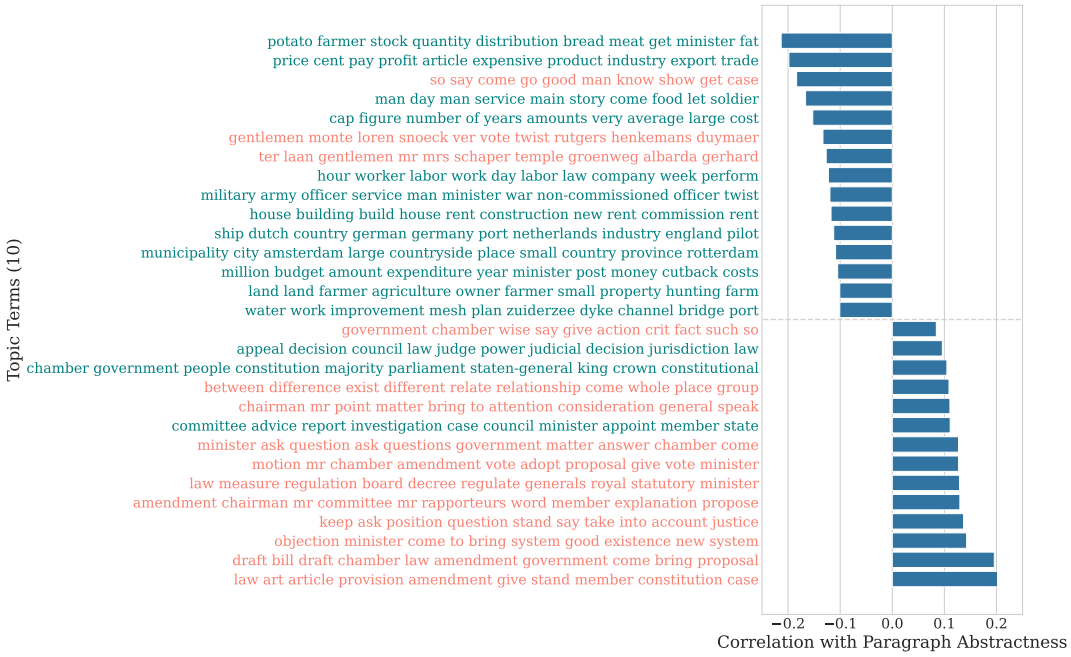

Figure 3: Pearson correlation coefficients between topic distributions and abstractness scores for words in the period between 1917 and 1923. The top 15 most negative and positively correlating topics are shown. The (translated) terms for procedural topics are colored red. The "semantic" topics (those not pertaining to parliamentary procedure but to specific policy areas or issues) are colored blue.

Figure 4 shows the aforementioned distances to the session mean averaged on the speaker level plotted against the general mean abstractness per speaker. The first noteworthy feature is the strong linear correlation. Speakers with a high average abstractness are also relatively abstract compared to the overall abstractness of the sessions in which they speak. This tells us that speaker-level abstractness cannot be explained by the (topic) abstractness of a session alone. General abstractness comes with local abstractness, meaning that even when we take into account the fact that some sessions are more abstract than others, speaker level differences persist.

4. Speaker Abstractness in Context

With abstractness scores now available on the paragraph-level and the signals improved with the topic filtering and session-paragraph distances, parliamentary abstractness can be analyzed on several levels. The data allows an analysis of abstractness at the level of parties or gender, or can be used to track abstractness through time. In this section, we take a brief look at speaker-level abstractness as an example of the type of analysis that can be performed with this method. The section focuses on the time period between 1917 and 1923. We do so because this period is currently under study in the broader project, and we thus know most of this period's political context.

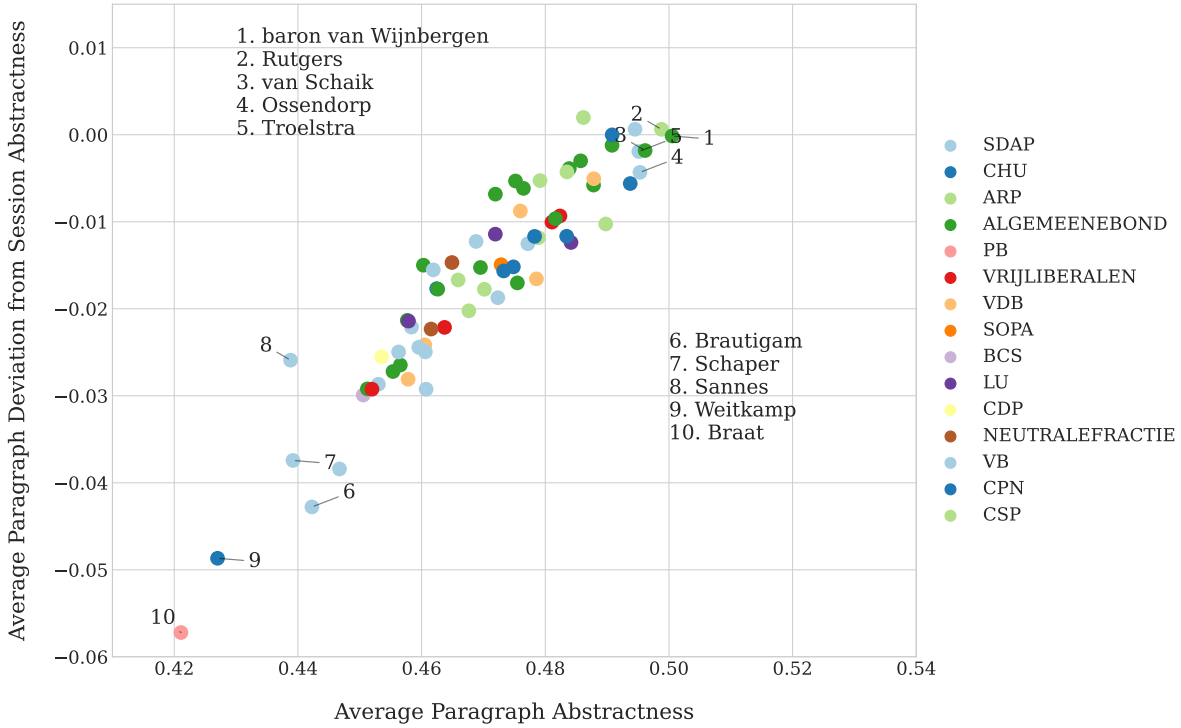

Figure 4: Speaker-level mean paragraph abstractness (x-axis) plotted against the mean deviation from session abstractness (y-axis) on the level of individual speakers in the period between 1917 and 1923. Only speakers with > 250 paragraphs are considered. The five most abstract and concrete speakers are annotated in the plot. Colors indicate party affiliation.

The period between 1917 and 1923 is marked by several fundamental transformations in Dutch politics. The 1917 and 1919 constitutional reforms led to universal suffrage and proportional representation. This left its mark on the party political divisions in parliament. Socialists and confessional parties gained ground, while the power of the formerly dominant liberals declined. Parliamentary debate in the period revolved to a considerable extend around these democratic issues. It also included frequent deliberation over the new areas in which the post-First World War administration was now active: education, welfare and labour [33].

Figure 4 shows the average abstractness and the average distance from the session mean for all speakers with more than 250 paragraphs. Alongside the points, the names of the top most concrete and abstract speakers are included in the plot. By investigating explanations for the appearance of specifically these names, more becomes clear about what exactly our method is capturing. The most abstract speakers are Van Wijnbergen (Catholic), Rutgers (Protestant), Van Schaik (Catholic), Ossendorp and Troelstra (Socialist). The most concrete speakers are Brautigam, Schaper, Sannes (Socialist), Weitkamp (Protestant) and Braat (Agricultural League).

A first important observation when looking at these speakers in their historical context is the persistent influence of topics. In the previous section, it was shown how procedural language impacts abstractness ratings. Non-procedural topics, however, also play a large role

in the constitution of the speaker abstractness average. In the case of Van Wijnbergen and Ossendorp, their frequent involvement in debates about education, a generally abstract topic, seems to explain much of their relatively high abstractness averages. Similarly, the involvement of Brautigam and Sannes in debates about food distribution, housing and shipping, and the involvement of Weitkamp and Braat in agricultural debates shows the close link between what policy area is discussed by a member and how concrete or abstract his (parliament hosted mostly men in this period) average score is. The focus of specific speakers on a limited number of policy areas thus logically explains their high or low average scores.

However, thematic preferences measured through topics, do not explain all individual differences. In fact, a closer look at the topics and individual speeches (subjected to close reading) indicates that the abstractness scores also go beyond thematic preference. Looking at the speeches of the ten members mentioned in Figure 4 and taking into account their individual background and reputation suggests that their abstract and concrete averages might also stem rhetorical style. Troelstra, the famous socialist leader, was known for his ideological style. His speeches are hardly about housing or agriculture, but about capitalism and revolution, and when he discusses concrete matters, he still does so in a relatively abstract way. Similarly, Victor Rutgers, the leader of the protestant A.R.P. (Anti-Revolutionary Party) seems to talk in a relatively abstract way. Even his most concrete paragraphs display a large number of abstract terms. In a similar fashion, Weitkamp and Braat, whose concreteness mainly stems from their involvement in agricultural debates (Braat was the only representative of the populist"Agricultural League"), also seem to speak in concrete ways in other areas. Both speakers were known as atypical members, not conforming to the parliamentary etiquette and even maintaining a way of speaking that was considered rude by their colleagues. Especially the latter, nicknamed "Boer Braat" ("Farmer Braat") was considered a disgrace to parliamentary dignity due to his "un-parliamentary" manners [34]. This potential biographical explanation of concrete style also brings us to the overall concrete socialists, whose most concrete members (Sannes, Schaper and Brautigam) all had a working-class background could have produced a rhetorical style that is clearly different from aristocratic conservatives and liberals, and hence more concrete. Future work should include a more comprehensive and structure examination of these individual factors, but the result of this first exploration already show interesting similarities that cut through party political divisions.

Another way of substantiating this rhetorical dimension of abstractness is to look at abstractness averages on the speaker-topic level. By calculating the mean abstractness for every topic and for every speaker, and subsequently ranking speaker topics based on their abstractness we observe that the general topic abstractness differed significantly from speaker-level topic abstractness. In other words, housing is a concrete topic overall, but in the case of Schaper it is highly abstract. Similarly, a topic on party politics ("party, politics, cabinet, rightwing") is generally abstract, but in the case of Weitkamp and Braat highly concrete. This is a final indication that our method registers not only variation in the things speakers talked about, but also differences in the way they spoke.

Overall, it remains difficult to disentangle and balance the various potential explanations for speaker abstractness. However, the discussion of the most abstract and concrete members shows that the method proposed in this paper is able to point at new axis of difference. Abstractness cuts through the party political lines that often form the basis for historical inquiry. The

fact that we find socialists among conservative protestants as abstract speakers, and other socialists among farmer representatives shows that much remains unknown when it comes to parliamentary rhetoric.

5. Discussion

All in all, our method is able to capture difference in abstractness on the level of parliamentary speech paragraphs. However, our research has brought forward several issues that should be addressed in further research. First, the dependency on individual word scores inevitably leads to the issue of semantic change. Words change meaning, and so does their abstractness, something which hinders applications of the method to historical data. However, the paper has shown that the embedding procedure, based on a robust number of seed terms, is unexpectedly good in identifying corpus-specific abstract and concrete terms. Moreover, comparing terms between time periods can help in identifying the historically changing abstractness of terms.

This issue also relates to word sense and figurative speech. The word "bank" has multiple senses and the level of abstractness for "ocean bank" and "savings bank" is large. In the current method, multi-sense terms are not tackled. Here, more recent embedding methods that are able to produce contextualized embeddings will be of great value.

Another pressing issue is the unit of analysis: the speech (paragraph). A first obvious drawback is the fact that internal variation is not seen by averaging word-abstractness in paragraphs. Speech paragraphs with highly abstract and concrete parts go unnoticed. Measuring lower-level linguistic units (such as sentences) might mitigate this effect.

The impact of thematic preferences, mapped through topic modelling, remains high, even when removing procedural language. The issue of differentiating between the *what* and the *how* of parliamentary speech remains difficult. Still, the paper has shown how several techniques can be used to mitigate this thematic bias. This has resulted in several indications that abstractness also tells us more about rhetorical style: the *how* of parliamentary debate.

Lastly, in relation to our broader concern with technocratic rhetoric it goes without saying that there is no one-to-one correspondence between abstractness and "technocraticness". Still, this method greatly benefits our research. The brief excursion into speaker-level differences shows the potential of the method in tracing unlikely familiarities and exploring stylistic differences that complement close reading analysis of speeches and histories of political thought. Using the differences and similarities between speakers, we were pointed at parliamentary representatives that had gone unnoticed before. In this way, our method might not yet be optimal for macro-level statistical analysis, but does help in contextualizing other forms of historical research.

6. Conclusion

This paper has proposed a method for measuring abstractness in parliamentary speeches. Based on an annotated set of words, the abstractness levels of unseen corpus-specific terms are estimated. Subsequently, speech paragraphs are scored by taking median scores. The paper has subsequently discussed the use of these scores in historical research and several approaches to

improving the abstractness signals. How abstract a member's speeches are depends to a large extent on their preferred themes and topics. However, individual rhetorical style also plays a role. Balancing these complex historical factors and "variables" has proven to be difficult, yet productive. The paper has demonstrated how methods from computational linguistics can be leveraged in historical research and how this amounts to forms of linguistic variation that remain hard to capture with established computational methods.

References

[1] L. Raphael, Die Verwissenschaftlichung des Sozialen als methodische und konzeptionelle Herausforderung für eine Sozialgeschichte des 20. Jahrhunderts 22 (1996) 165–193.

[2] T. Hellstrom, M. Jacob, Scientification of politics or politicization of science? Traditionalist science-policy discourse and its quarrels with Mode 2 epistemology 14 (2000) 69–77. URL: https://doi.org/10.1080/02691720050199315. doi:10.1080/02691720050199315.

[3] K. Brückweh, R. F. Wetzell, Engineering Society: The Role of the Human and Social Sciences in Modern Societies, 1880-1980, Springer, 2012. arXiv:CEP4X4Ib32QC.

[4] J. Habermas, J. J. Shapiro, The Scientization of Politics and Public Opinion, in: Toward a Rational Society: Student Protest, Science, and Politics, Beacon Press, 1970, pp. 62–80.

[5] P. Pettit, Depoliticizing democracy 17 (2004) 52–65.

[6] F. Moretti, D. Pestre, Bankspeak: The language of World Bank reports 92 (2015) 75–99.

[7] B. J. McKenna, P. Graham, Technocratic Discourse: A Primer 30 (2000-07) 223–251. URL: http://journals.sagepub.com/doi/10.2190/56FY-V5TH-2U3U-MHQK. doi:10.2190/56FY-V5TH-2U3U-MHQK.

[8] P. d. Rooy, historicus, H. t. Velde, W. Kok, Met Kok: over veranderend Nederland, Wereldbibliotheek, 2005. URL: http://www.dbnl.nl/tekst/rooy011metk01_01/.

[9] P. Fawcett, M. V. Flinders, C. Hay, M. Wood (Eds.), Anti-Politics, Depoliticization, and Governance, first edition ed., Oxford University Press, 2017.

[10] J. H. ten Cate, De "mannen van de daad" en Duitsland, 1919-1939., Rijksinst. voor Oorlogsdocumentatie, 1995. arXiv:kXZTnQEACAAJ.

[11] M. Fortescue, The Abstraction Engine: Extracting Patterns in Language, Mind and Brain, John Benjamins, 2017-04-15. URL: https://www.jbe-platform.com/content/books/9789027265845. doi:10.1075/aicr.94.

[12] A. M. Borghi, L. Barca, F. Binkofski, L. Tummolini, Varieties of abstract concepts: Development, use and representation in the brain 373 (2018-08-05) 20170121. URL: https://royalsocietypublishing.org/doi/10.1098/rstb.2017.0121. doi:10.1098/rstb.2017.0121.

[13] M. Bolognesi, C. Burgers, T. Caselli, On abstraction: Decoupling conceptual concreteness and categorical specificity 21 (2020-08-01) 365–381. URL: https://doi.org/10.1007/s10339-020-00965-9. doi:10.1007/s10339-020-00965-9.

[14] J. Charbonnier, C. Wartena, Predicting Word Concreteness and Imagery, in: Proceedings of the 13th International Conference on Computational Semantics - Long Papers, Association for Computational Linguistics, 2019-05, pp. 176–187. URL: https://aclanthology.org/W19-0415. doi:10.18653/v1/W19-0415.

[15] J. Borden, X. A. Zhang, Linguistic Crisis Prediction: An Integration of the Linguistic

Category Model in Crisis Communication 38 (2019-10-01) 650–679. URL: https://doi.org/10.1177/0261927X19860870. doi:10.1177/0261927X19860870.

[16] M. Dragojevic, A. Sink, D. Mastro, Evidence of Linguistic Intergroup Bias in U.S. Print News Coverage of Immigration 36 (2017-09-01) 462–472. URL: https://doi.org/10.1177/0261927X16666884. doi:10.1177/0261927X16666884.

[17] R. J. Heuser, Abstraction: A Literary History, 2019.

[18] M. Marx, J. V. Doornik, A. Nusselder, L. Buitinck, Politicalmashup 1814-2012 - members, parties, proceedings, 2012. URL: https://easy.dans.knaw.nl/ui/datasets/id/easy-dataset:50409. doi:10.17026/DANS-294-YW2Z.

[19] PICCL: Philosophical Integrator of Computational and Corpus Libraries, 2022. URL: https://github.com/LanguageMachines/PICCL, original-date: 2017-03-25T12:01:28Z.

[20] M. Wevers, J. van Eijnatten, J. Verheul, Consuming America : A Data-Driven Analysis of the United States as a Reference Culture in Dutch Public Discourse on Consumer Goods, 1890-1990, 2017-09-15. URL: https://dspace.library.uu.nl/handle/1874/355070.

[21] P. D. Turney, M. L. Littman, Measuring praise and criticism: Inference of semantic orientation from association 21 (2003-10-01) 315–346. URL: https://doi.org/10.1145/944012.944013. doi:10.1145/944012.944013.

[22] J. Dunn, Modeling Abstractness and Metaphoricity 30 (2015-10-02) 259–289. URL: http://www.tandfonline.com/doi/full/10.1080/10926488.2015.1074801. doi:10.1080/10926488.2015.1074801.

[23] M. Köper, S. S. Im Walde, Automatically generated affective norms of abstractness, arousal, imageability and valence for 350 000 german lemmas, in: Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16), 2016, pp. 2595–2598.

[24] P. Turney, Y. Neuman, D. Assaf, Y. Cohen, Literal and Metaphorical Sense Identification through Concrete and Abstract Context, in: Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, 2011-07, pp. 680–690. URL: https://aclanthology.org/D11-1063.

[25] E. Rabinovich, B. Sznajder, A. Spector, I. Shnayderman, R. Aharonov, D. Konopnicki, N. Slonim, Learning Concept Abstractness Using Weak Supervision, 2018-09-04. URL: http://arxiv.org/abs/1809.01285. arXiv:1809.01285.

[26] M. Brysbaert, M. Stevens, S. De Deyne, W. Voorspoels, G. Storms, Norms of age of acquisition and concreteness for 30,000 Dutch words 150 (2014-07-01) 80–84. URL: https://www.sciencedirect.com/science/article/pii/S0001691814000985. doi:10.1016/j.actpsy.2014.04.010.

[27] M. Coltheart, The MRC psycholinguistic database 33 (1981) 497–505.

[28] J. M. Clark, A. Paivio, Extensions of the Paivio, Yuille, and Madigan (1968) norms 36 (2004-08-01) 371–383. URL: https://doi.org/10.3758/BF03195584. doi:10.3758/BF03195584.

[29] R. Rehurek, P. Sojka, Gensim–python framework for vector space modelling 3 (2011).

[30] M. Antoniak, D. Mimno, Evaluating the stability of embedding-based word similarities 6 (2018) 107–119.

[31] C. Wiesner, T. Haapala, K. Palonen, Debates, Rhetoric and Political Action, Palgrave Macmillan UK, 2017. URL: http://link.springer.com/10.1057/978-1-137-57057-4. doi:10.1057/978-1-137-57057-4.

[32] K. Palonen, J. M. Rosales, T. Turkka, The Politics of Dissensus: Parliament in Debate, Ed. Universidad de Cantabria, 2014-05-05. arXiv:dPidAwAAQBAJ.

[33] P. De Rooy, The Nation is Divided into Parties, in: A Tiny Spot on the Earth, Amsterdam University Press, 2015-01-20, pp. 185–228. URL: https://www.degruyter.com/document/doi/10.1515/9789048524150-007/html.

[34] I. van den Broek, De taal van het anti-parlementarisme. Poëzie en politiek in Nederland 1870-1940 120 (2005-01-01) 466–496. URL: https://bmgn-lchr.nl/article/view/URN%3ANBN%3ANL%3AUI%3A10-1-107154. doi:10.18352/bmgn-lchr.6256.