

The ParlaSent Multilingual Training Dataset for Sentiment Identification in Parliamentary Proceedings

Michal Mochtak, Peter Rupnik, Nikola Ljubešić

Dept. of Political Science, Dept. of Knowledge Technologies, Dept. of Knowledge Technologies
Radboud University, Jožef Stefan Institute, Jožef Stefan Institute
michal.mochtak@ru.nl, peter.rupnik@ijs.si, nikola.ljubesic@ijs.si

Abstract

The paper presents a new training dataset of sentences in 7 languages, manually annotated for sentiment, which are used in a series of experiments focused on training a robust sentiment identifier for parliamentary proceedings. The paper additionally introduces the first domain-specific multilingual transformer language model for political science applications, which was additionally pre-trained on 1.72 billion words from parliamentary proceedings of 27 European parliaments. We present experiments demonstrating how the additional pre-training on parliamentary data can significantly improve the model downstream performance, in our case, sentiment identification in parliamentary proceedings. We further show that our multilingual model performs very well on languages not seen during fine-tuning, and that additional fine-tuning data from other languages significantly improves the target parliament's results. The paper makes an important contribution to multiple disciplines inside the social sciences, and bridges them with computer science and computational linguistics. Lastly, the resulting fine-tuned language model sets up a more robust approach to sentiment analysis of political texts across languages, which allows scholars to study political sentiment from a comparative perspective using standardized tools and techniques.

Keywords: sentiment, parliament, multilingual model

1. Introduction

Emotions and sentiment in political discourse are deemed as crucial and influential as substantive policies promoted by the elected representatives (Young and Soroka, 2012). Since the golden era of research on propaganda (Lasswell, 1927; Shils and Janowitz, 1948), several scholars have demonstrated the growing role of emotions on affective polarization in politics with negative consequences for the stability of democratic institutions and social cohesion (Garrett et al., 2014; Iyengar and Ansolabehere, 1995; Mason, 2015). With the booming popularity of online media, sentiment analysis has become an indispensable tool for understanding the positions of viewers, customers, and voters (Soler et al., 2012). It has allowed all sorts of entrepreneurs to know their target audience like never before (Ceron et al., 2019). Experts on political communication argue that the way we receive and process information plays an important role in political decision-making, shaping our judgment with strategic consequences both on the level of legislators and the masses (Liu and Lei, 2018). Emotions and sentiment simply do play an essential role in political arenas, and politicians have been (ab)using them for decades.

Although there is a general agreement among political scientists that sentiment analysis represents a critical component for understanding political communication in general (Young and

Soroka, 2012; Flores, 2017; Tumasjan et al., 2010), the empirical applications outside the English-speaking world are still rare (Rauh, 2018; Mohammad, 2021). Moreover, many of the research applications in social sciences lag behind the latest methodological advancements grounded in computational linguistics. This is especially the case for studies analyzing political discourse in low-resourced languages, where the lack of out-of-the-box tools creates a huge barrier for social scientists to do such research in the first place (Proksch et al., 2019; Mochtak et al., 2020; Rauh, 2018). As a result, many of the applications still rely on dictionary-based methods, which tend to produce potentially skewed results (Hardeniya and Borikar, 2016; Proksch et al., 2019) or approximate sentiment scores based on position-taking stances with relatively high-level generalization (e.g. roll-calls or voting behavior (Abercrombie and Batista-Navarro, 2018a)). Field-specific sentiment identifiers trained using machine learning algorithms are comparatively rare. Part of the reason is the fact that training machine learning models can be prohibitively expensive, especially when it comes to collecting, cleaning, and processing training data. However, recent development in the field of computational linguistic and natural language processing fueled by transformer-based deep learning models has lowered the bar for social scientists substantially. This development has additionally allowed for existing language models

to be adapted to a specific domain by additionally pre-training the language model on non-annotated domain data (Sung et al., 2019).

The paper presents annotated sentence-level datasets in seven European languages (Bosnian, Croatian, Czech, English, Serbian, Slovak, and Slovenian) sampled from parliamentary proceedings of seven European countries (Bosnia-Herzegovina, Croatia, Czech Republic, Serbia, Slovakia, Slovenia, and United Kingdom). The selection of proceedings is driven by the existing gap in low-resourced languages of Central and Eastern Europe and their relevance in a broader comparative perspective. The human-annotated datasets are used in a series of experiments focused on training sentiment identifiers intended for detecting sentiment in political discourse. Apart from methodological and experimental goals the paper has, it also can be read as summary guidelines for social scientists interested in training their own sentiment identifiers with similar scope. The paper is written with the intention of facilitating the process of collecting, cleaning, and processing textual data for political science research with realistic estimates for needed resources. When it comes to substantial findings, the paper shows that 1) additional pre-training of a language model on raw parliamentary data can significantly improve the model performance on the task of sentiment identification; 2) large 561-million-parameter multilingual models perform drastically better than those with half of the parameters; 3) multilingual models work very well also on unseen languages; and finally 4) even when the language-specific training data exist for the parliament proceedings one wants to process, a multilingual model trained on four times the size of the dataset from other languages improves the results on the target parliament significantly.

2. Related work

Despite the boom of computational methods in recent years has shown new ways to perform sentiment analysis with relatively high accuracy, political science is catching up relatively slowly. Abercrombie and Batista-Navarro (2020) found that most of the automated applications focused on parliamentary debates and position-taking exist outside of the mainstream of political science, both a surprise and an opportunity for future research. Addressing the existing gap reflects upon the needs of empirical political science research, which recognizes that people tend to interact with politics through emotions (Masch and Gabriel, 2020). Recent research has found that political leaders are keen to use violent and populist rhetoric to connect with citizens on an emotional

level (Gerstlé and Nai, 2019; Piazza, 2020; Masch and Gabriel, 2020). As an effective campaigning strategy, populist parties in Europe use significantly more negative framing than their less populist counterparts simply because negative emotions work (Widmann, 2021). They are often abused as highly conflicting drivers leading to affective polarization (Druckman et al., 2021; Iyengar et al., 2012), negative partisanship (Abramowitz and Webster, 2016), group-based partisan competition (Mason, 2018), and political sectarianism (Finkel et al., 2020). If connected with the long-run emotional effects on the electorate, the impacts are disastrous. Partisan dehumanization, partisan antipathy, and acceptance of partisan violence are just a few examples of morphed competition infused with an emotionally laden identity fueling hostility, bias, and anger (Webster and Albertson, 2022). Understanding these mechanisms is highly important.

When it comes to actual applications focused on political domain, sentiment analysis can be most often found in research focused on classification (Abercrombie and Batista-Navarro, 2018a,b; Akhmedova et al., 2018; Bansal et al., 2008; Bonica, 2016) and dictionary-based sentiment extraction (Honkela et al., 2014; Onyimadu et al., 2013; Mochtak et al., 2022a; Owen, 2017; Proksch et al., 2019). The first stream of research uses different ML algorithms to develop models able to “classify” textual data into predefined categories (classes). These categories are either generated in an automated way based on known polarity traces, such as yes/no votes assigned to MPs’ speech acts (Salah, 2015; Abercrombie and Batista-Navarro, 2018a), or are produced using traditional manual annotation with ground-truth labels (Onyimadu et al., 2013; Rauh, 2018). A majority of these applications fall under the umbrella of supervised learning using a wide range of algorithms, from logistic regression to naïve Bayes, decision trees, nearest neighbor, or boosting. In recent years, many applications in computer science have been significantly tilted towards strategies using neural networks ranging from ‘vanilla’ feed-forward networks to more complex architectures such as transformers pre-trained on raw non-annotated data (Pipalia et al., 2020). In political science, however, dictionary-based strategies are still the dominant approach. They are traditionally focused on counting words with known sentiment affinity in raw text and generalizing their frequencies over the unit of analysis. Although sentiment dictionaries are deemed less accurate and may produce relatively crude estimates, their usage in political and social sciences is quite popular (Mochtak et al., 2022a; Rinker, 2017; Abercrombie and Batista-Navarro,

2020; Proksch et al., 2019). We see that as an opportunity for substantial improvement. The following sections present a new dataset for training a domain-specific sentiment identifier, which builds on a first-of-its-kind domain-specific transformer language model, additionally pre-trained on 1.72B domain-specific words from proceedings of 27 European parliaments. In a series of experiments, we then demonstrate how robust the approach is in various settings, proving its reliability in real-life applications.

3. Dataset construction

3.1. Focus on sentences

The datasets we compile and then use for training different prediction models focus on a sentence-level data and utilize sentence-centric approach for capturing sentiment polarity in text. The focus on sentences as the basic level of the analysis goes against the mainstream research strategies in social sciences which prefer either longer pieces of text (e.g. utterance of ‘speech segment’ or whole documents (Bansal et al., 2008; Thomas et al., 2006)) or generally more coherent messages of shorter nature (Tumasjan et al., 2010; Flores, 2017). However, this approach creates limitations when it comes to political debates in national parliaments, where speeches range from very short comments counting only a handful of sentences to long monologues with thousands of words. Moreover, as longer text may contain a multitude of sentiments, any annotation attempt must generalize them, introducing a complex coder bias that is embedded in any subsequent analysis. The sentence-centric approach attempts to refocus the attention on individual sentences capturing attitudes, emotions, and sentiment positions and use them as lower-level indices of sentiment polarity in a more complex political narrative. Although sentences cannot capture complex meanings as paragraphs or whole documents do, they usually carry coherent ideas with relevant sentiment affinity. This approach stems from a tradition of content analyses used by such projects as Comparative Manifesto (Volkens et al., 2020), the core-sentence approach in mobilization studies (Hutter et al., 2016), or claims analysis in public policy research (Koopmans and Statham, 2006).

Unlike most of the literature which approaches sentiment analysis in political discourse as a proxy for position-taking stances or as a scaling indicator (Abercrombie and Batista-Navarro, 2020; Glavaš et al., 2017; Proksch et al., 2019), a general sentence-level classifier has a more holistic (and narrower) aim. Rather than focusing on a specific policy or issue area, the task is to assign a correct sentiment category to sentence-level data in political discourse with

the highest possible accuracy. Only when a well-performing model exists, knowing whether a sentence is positive, negative, or neutral provides a myriad of possibilities for understanding the context of political concepts as well as their role in political discourse. Furthermore, sentences as lower semantic units can be aggregated to the level of paragraphs or whole documents, which is often impossible the other way around (document → sentences). Although sentences as the basic level of analysis are less frequent in political science research, existing applications include the validation of sentiment dictionaries (Rauh, 2018), ethos mining (Duthie and Budzynska, 2018), opinion mining (Naderi and Hirst, 2016), or detection of sentiment carrying sentences (Onyimadu et al., 2013).

We base our experiments on data sampled from parliamentary proceedings which provide representative and often exhaustive evidence on political discourse in respective countries. In this context, parliamentary debates are considered to be a rich source of information on the position-taking strategies of politicians and one of the best sources of political discourse in general (Lakoff, 2004). As democracy thrives through debate, tracing it becomes essential to understanding politics, its development, and its consequences. In this context, essentially, all democratic parliaments hold a debate before a bill is passed. If public, the debate becomes evidence of how members of parliaments represent their voters and constituencies and their personal beliefs and interests (Chilton, 2004). With all their flaws and shortcomings, parliamentary debates are an important aspect of political representation and an irreplaceable element of democratic systems. They connect voters with their representatives and show how responsive politicians are to people’s wishes (Powell, 2004).

3.2. Background data

In order to compile datasets capturing political discourse, manually annotate them, and then use them for training the classification models for real world applications, we sampled sentences from seven corpora of parliamentary proceedings – Bosnia and Herzegovina (Mochtak et al., 2022c)¹, Croatia (Mochtak et al., 2022a)², Serbia (Mochtak et al., 2022b)³, Czechia (Erjavec et al., 2023)⁴,

¹<https://doi.org/10.5281/zenodo.6517697>

²<https://doi.org/10.5281/zenodo.6521372>

³<https://doi.org/10.5281/zenodo.6521648>

⁴<https://www.clarin.si/repository/xmlui/handle/11356/1486>

Slovakia (Mochtak, 2022)⁵, Slovenia (Erjavec et al., 2023)⁶, and United Kingdom (Erjavec et al., 2023)⁷. The Bosnian corpus contains speeches collected on the federal level. Both chambers are included – House of Representatives (Predstavnički dom / Zastupnički dom) and House of Peoples (Dom naroda). The corpus covers the period from 1998 to 2018 and counts 127,713 speeches. The Czech corpus covers the period of 2013-2021 and counts 154,460 speeches from the Chamber of Deputies, the lower house of the parliament (Poslanecká sněmovna). The Croatian corpus of parliamentary debates covers debates in the Croatian parliament (Sabor) from 2003 to 2020 and counts 481,508 speeches. The Serbian corpus contains 321,103 speeches from the National Assembly of Serbia (Skupština) over the period of 1997 to 2020. The Slovenian corpus covers the period of 2000-2022 and counts 311,354 speeches from the National Assembly (Državni zbor), the lower house of the parliament. The Slovak corpus contains speeches from the period of 1994-2020 from the National Council of the Slovak Republic (Národná rada) and counts 375,024 speeches. Finally, the corpus from British Parliament covers speeches from both the House of Commons and the House of Lords from the period of 2015-2021 counting 552,103 speeches.

3.3. Data sampling

Speeches were segmented into sentences and words using the CLASSLA-Stanza (Ljubešić and Dobrovoljc, 2019; Terčon and Ljubešić, 2023) and Trankit (Nguyen et al., 2021) pipelines with tokenizers available for Czech, Croatian, Serbian, Slovak, Slovene, and English languages. This step was necessary in order to extract individual sentences as the basic unit of our analysis. In the next step, we filtered out only sentences presented by actual speakers, excluding moderators of the parliamentary sessions. Sentences are preserved within their country-defined pools with the exception of Bosnia and Herzegovina, Croatia, and Serbia which were merged together as representatives of one language family (i.e. the corpora were treated as one sampling pool). We are, from now on, referring to this pool as BCS (Bosnian-Croatian-Serbian)⁸. As we want to sample what

⁵<https://doi.org/10.5281/zenodo.7020474>

⁶<https://www.clarin.si/repository/xmlui/handle/11356/1486>

⁷<https://www.clarin.si/repository/xmlui/handle/11356/1486>

⁸The main reason for keeping these three parliaments in one pool is previous work on annotating sentiment in parliamentary proceedings (Mochtak et al., 2022b) which consisted of 2,600 instances jointly sampled from the three underlying parliaments.

can be understood as “average sentences”, we further subset each sentence-based corpus to only sentences having the number of tokens within the first and third frequency quartile (i.e. being within the interquartile range) of the original corpora. Having the set of “average sentences”, we used common sentiment lexicons available for each of the languages (Glavaš et al., 2012; Chen and Skiena, 2014, 2016), and applied them as seed words for sampling sentences for manual annotation. These seed words are used for stratified random sampling which gives us 867 sentences with negative seed word(s), 867 sentences with positive seed word(s), and 866 sentences with neither positive nor negative seed words (supposedly having neutral sentiment) per sampling pool. We sample 2,600 sentences in total for manual annotation per corpus. Under this setting, the sentences inherit all metadata information of their parent documents.

We further sample two random datasets, one from the BCS collection of parliaments, another from the English parliament, not applying the sentiment seed list, but rather aiming at a random selection of sentences, still satisfying the criterion of the sentence length falling into the interquartile range. The primary use case for these two datasets is testing various sentiment identification approaches, therefore we wanted for their sentiment distribution to follow the one occurring in the underlying parliaments. The sampling pipeline is identical to seed-based datasets but without the seed words component. For our experiments, we again sample 2,600 average-length sentences cleaned off of entries from proceedings’ moderators (see above). From now on, we refer to these two datasets as the BCS-test and the EN-test sets.

3.4. Annotation schema

The annotation schema for labelling sentence-level data was adopted from Batanović et al. (2020) who propose a six-item scale for annotation of sentiment polarity in a short text. The schema was originally developed and applied to SentiComments.SR, a corpus of movie comments in Serbian and is particularly suitable for low-resourced languages. The annotation schema contains six sentiment labels (Batanović et al., 2020, p. 6):

- +1 (Positive in our dataset) for sentences that are entirely or predominantly positive
- -1 (Negative in our dataset) for sentences that are entirely or predominantly negative
- +M (M_Positive in our dataset) for sentences that convey an ambiguous sentiment or a mixture of sentiments, but

lean more towards the positive sentiment in a positive-negative classification

- -M (`M_Negative` in our dataset) for sentences that convey an ambiguous sentiment or a mixture of sentiments, but lean more towards the negative sentiment in a positive-negative classification
- +NS (`P_Neutral` in our dataset) for sentences that only contain non-sentiment-related statements, but still lean more towards the positive sentiment in a positive-negative classification
- -NS (`N_Neutral` in our dataset) for sentences that only contain non-sentiment-related statements, but still lean more towards the negative sentiment in a positive-negative classification

The different naming convention we have applied in our dataset serves primarily practical purposes: obtaining the 3-way classification by taking under consideration only the second part of the string (if an underscore is present).

The benefit of the whole annotation logic is that it was designed with versatility in mind allowing reducing the sentiment label set in subsequent processing if needed. That includes various reductions considering polarity categorization, subjective/objective categorization, or change of the number of categories. This is important for various empirical tests we perform in the following sections.

3.5. Data annotation

Data were annotated in multiple iterations. Each seed-based dataset was annotated by two annotators, both being native speakers or having an excellent command of the language to be annotated. Annotation was done via a custom-built, locally-hosted online app using *prodigy* with consistent logging allowing monitoring of the annotation process systematically (Montani and Honnibal). Each annotator went through approximately ten hours of training before the actual annotation. The annotation was done iteratively in several rounds, with automated oversight of the coding process in order to minimize any systematic disagreement. Trained annotators were able to annotate up to 75 sentences per hour on average, resulting in 35 person-hours per annotator, per dataset (feedback and reconciliation not included). When it comes to BCS and English test sets, data were annotated by only one highly trained annotator. Similarly to previous setting, annotation of both test sets were followed by quality control procedures adjusted for just one annotator (e.g., consistency monitoring; pace; consultations).

Despite the relatively smooth annotation process across the datasets, the inter-annotator agreement (IAA) measured using Krippendorff's alpha (KA) has not reached high values. This is consistent across datasets supporting the argument that sentiment perception is a highly subjective matter (Bermingham and Smeaton, 2011; Mozetič et al., 2016) and, despite the effort to eliminate hard disagreements, the results often reflect upon a hard call the annotators had to make. Monitoring and reconciliation of the disagreements further showed that most of the disagreements are not substantially wrong and can be considered relevant under the provided context (i.e. shorter text snippets). The summary of Krippendorff's alpha and agreement rates (accuracy) across datasets and annotation schemas is presented in Table 1.⁹

The final distributions of the three-class labels after reconciliation across datasets are presented in Table 2. The reconciliation was done by the original annotators after the main annotation round was finished. As the reconciliation was done by mutual agreement, the whole process was administrated online to avoid any immediate peer pressure. Annotators were first asked to mark annotations for which they do not have problem to agree with their colleague, indicating their reconciliation position without the need for actual deliberation (i.e., discussion). This approach helped to eliminate disagreements which can be considered easy. Annotators then met in person or online and discussed instances which they could not agree on in the first round. Each dataset contained a handful of hard disagreements that were not possible to reconcile without an external reviewer (one of the authors of this paper). The presented distributions show that the negative class is often the most populous category. This observation, however, does not entirely hold true in the Slovene dataset and English test set (see Table 2). We theorize that this might indicate potential nuances in political nature of different parliaments and an existence of different patterns in their political culture. Neutral sentiment appears to be more dominant there.

3.6. Dataset encoding

The final dataset is encoded as a JSONL document, each line in the dataset representing a JSON object. The dataset is encoded in seven files, five files representing the 2,600 training instances per language (group) and two files representing our two BCS and English test sets.

Each of the training datasets contains the following metadata:

⁹We do not report KA for the BCS and English test sets here as only one annotator performed the annotations.

Dataset	ACC (6 classes)	KA (6 classes)	ACC (3 classes)	KA (3 classes)
BCS	62.0%	0.502	77.7%	0.639
CZ	68.1%	0.531	77.4%	0.612
SK	63.4%	0.506	75.4%	0.607
SL	64.1%	0.502	73.7%	0.531
EN	66.0%	0.543	78.4%	0.656

Table 1: Krippendorff’s alpha (KA) and agreement rates (ACC) across datasets for 6-fold and 3-fold annotation schemas.

Dataset	Negative	Neutral	Positive
all	8232	6691	3277
BCS	1314	773	513
CZ	1398	866	336
SK	1253	895	452
SL	1010	1409	181
EN	1269	680	651
BCS-test	1147	1006	447
EN-test	841	1062	697

Table 2: Distribution of the three-class labels across datasets.

- `sentence` that is annotated.
- `country` of origin of the sentence
- annotation of `annotator1` with one of the labels from the annotation schema presented in Section 3.4
- annotation of `annotator2` following the same annotation schema
- annotation given during `reconciliation` of hard disagreements
- the three-way `label` (positive, negative, neutral) where +NS and -NS labels are mapped to the neutral class
- the `document_id` the sentence comes from
- the `sentence_id` of the sentence
- the `term` inside which the speech was given
- the `date` the speech was given
- the `name`, `party`, `gender`, `birth_year` of the speaker
- the `split` the instance has been assigned to (training set containing 2054 instances, development set 180 instances, and testing set 366 instances)
- the `ruling` status of the party while the speech was given (opposition or coalition)

The EN-test and BCS-test sets differ in their structure minimally, containing only the `annotator1` attribute, and missing the `annotator2` and `reconciliation` attributes, as single annotator performed the annotation of test sets. Furthermore, there is no `split` attribute as the purpose of these datasets is testing various algorithms, while the training datasets can also be used for language-specific experiments, therefore requiring the train-dev-test split.

The dataset is made available through the CLARIN.SI repository at <http://hdl.handle.net/11356/1585> and is available under the CC-BY-SA 4.0 license.

4. Experiments

In this section, we present our experiments through which we aim to answer the following three research questions:

Q1: Does our newly released XLM-R-parla model, which is an XLM-R-large model additionally pre-trained on 1.72 billion words of parliamentary proceedings, model the sentiment phenomenon better than the original XLM-R models?

Q2: How well does our model work on languages not seen during fine-tuning?

Q3: If one wants to process data from some parliament that is covered in the ParlaSent (or any other) dataset, is it advisable to train a language-specific (and parliament-specific) model, or rather train a multilingual model containing the whole ParlaSent dataset?

To make the most out of our rather complex 6-level annotation schema, we set-up all our experiments as regression tasks, the six levels being modified into integer values from 0 to 5. For evaluation, we use primarily the R^2 score, which quantifies the proportion of the prediction variance that can be explained through gold labels, due to its sensitivity to differences in system performance. We also report mean average error (MAE) as a simple-to-understand metric in terms of an average error per instance, knowing that we are predicting values on a scale from 0 to 5, maximum MAE thereby being 5, and acceptable error for most use cases being somewhere below 1.

model \ test	R^2		MAE	
	BCS	en	BCS	EN
Dummy	-0.012	-0.165	1.522	1.645
XLm-R-base	0.500	0.561	0.868	0.808
XLm-R-large	0.605	0.653	0.706	0.694
XLm-R-parla	0.615	0.672	0.705	0.675

Table 3: Results of the first research question comparing the additionally-pretrained XLm-R-parla model with the vanilla XLm-R models and a random baseline.

4.1. The XLm-R-parla model

In this subsection we are answering our first question – whether our newly released model XLm-R-parla¹⁰ performs better than the vanilla XLm-R models of size base and large (Conneau et al., 2019).

The XLm-R-parla model is based on the XLm-R-large model¹¹ due to our preliminary experiments showing that XLm-R-large models outperform XLm-R-base models on this task. The XLm-R-parla model was additionally pre-trained for only 8 thousand steps with a batch size of 1024 sequences of 512 tokens of text. The model was pre-trained on a merge of the ParlaMint 3.0 dataset (Erjavec et al., 2023) and the EuroParl dataset (Koehn, 2011), together covering 30 languages, with 1,717,113,737 words of running text. Important to mention is that our pre-training dataset consists of all the languages contained inside the ParlaSent dataset.

Our hypothesis for this question is that the additional adaptation of the XLm-R-large model to texts as they occur in parliamentary proceedings will significantly improve our predictions of sentiment in parliamentary proceedings.

We fine-tune both the XLm-R-base and XLm-R-large, as well as the XLm-R-parla model, with the same hyperparameter settings, that have proven in preliminary experiments to perform well for our task: learning rate of 8e-6, batch size of 32, and 4 epochs over our whole training dataset (2,600 instances per each of the five parliament(pool)s, 13,000 instances in total). We test each model separately on the BCS and the English test set, each comprising of 2,600 instances. We compare these three models with a dummy baseline returning always the mean value of the training data. We perform five runs for each set-up and report the mean result.

The results in Table 3 show that the mean dummy, as expected, gives a R^2 value of around 0, while the MAE is around 1.5, which means that, if we always predicted a mean value from our training

¹⁰<https://huggingface.co/classla/xlm-r-parla>

¹¹<https://huggingface.co/xlm-roberta-large>

data, on average, we would be “only” 1.5 points away from the correct value. This result represents the baseline result any reasonable system should improve over. The XLm-R-base model does exactly that, lowering the MAE to between 0.81 and 0.87, depending on the test set.

The XLm-R large model, identical to our preliminary experiments, drastically improves over the base-sized model, which simply shows that the task at our hands is a rather complex one and that the extra capacity delivers around 0.1 points better results in R^2 (scale 0-1), which can be called a drastic improvement. The MAE score, much less sensitive to changes in the quality of predictions, still shows an error lower on average of 0.1 to 0.15 points (scale 0-5).

Once we compare the original XLm-R-large model with the additionally pre-trained XLm-R-parla model, we can observe that the additional pre-training has paid off, with minor, but consistent improvements on all metrics.

As expected, all systems perform better on the English test set due to much more English data seen during pre-training than that of Bosnian, Croatian, or Serbian.

We can conclude by offering an answer to the first research question – the additional pre-training of a multilingual model on parliamentary data does improve the performance on our task.

4.2. Performance on unseen languages

Here, we are answering our second question – how is the performance of our best-performing model XLm-R-parla on a language that the model has not seen during fine-tuning?

Our initial hypothesis is that there would be a minor impact on whether the model was fine-tuned on the testing language or not.

We perform two ablation experiments, in one skipping BCS data from the training dataset, and in the other skipping English data, and we evaluate both models on the BCS and the English test sets. Therefore, in the two additional experiments, we do not train on 13,000 but on 10,400 instances. We keep the same hyperparameter values as with the initial XLm-R-parla experiment described in the previous subsection.

In Table 4, reporting the results of these experiments, to our surprise, we cannot observe any obvious pattern regarding the performance on the language that has been removed from the training data. On the BCS example the model performs even better on the BCS test set regarding the R^2 evaluation metric and worse on the EN dataset. If we look at the MAE metric, the results are slightly more expected, the model performing worse on both test sets, the drop being still more significant on the English test set. The inconsistency between the two metrics on the BCS

training set	R^2		MAE	
	BCS	en	BCS	en
ParlaSent	0.615	0.672	0.705	0.675
ParlaSent $\setminus \{BCS\}$	0.630	0.659	0.727	0.704
ParlaSent $\setminus \{EN\}$	0.596	0.655	0.728	0.756

Table 4: Experiments on removing the BCS and English data from the training data, evaluating on the BCS and English test data, to check for performance on languages not seen during fine-tuning.

test set is very likely due to R^2 penalizing outliers more harshly than MAE does.

In the experiment where the English training data is removed, the results are a little more consistent, with performance on both test sets being similarly worse, the English test set getting an extra hit on the MAE metric but not on the R^2 metric.

Overall, we cannot observe a hit on the performance of the models if the testing language gets removed from the training data to a greater extent than what is observed on the other test set. Therefore, we can conclude that the performance drops due to less training data, not due to the target language not being present in the training data.

4.3. Monolingual vs. multilingual training

In this subsection, we answer our third research question – whether target language performance is better if the model is trained on that language only, or rather if it is trained on all five parliament (group)s.

We hypothesize that results might be evened out. On the monolingual side, there is a drastic similarity between the training and the test data, not just due to the same language used, but also due to data coming from the same parliament, each parliament surely having many specific features a system might use to predict sentiment. On the multilingual side, the argument relies on five times more training data than in the monolingual setting.

In this set of experiments, we train and evaluate only on the training datasets, which have a train-dev-test split, as already reported in Section 3.6. In the case of the monolingual setting, we train the model only on the 2054 training instances for the specific parliament, while in the case of the multilingual setting, we train on five times that amount of data, i.e., 10,270 instances. We always evaluate on the test portion of the training dataset of the target parliament. We keep our hyperparameters the same as before, with the difference that monolingual models are fine-tuned for ten epochs due to less data available for fine-tuning.

The results in Table 5 show that only in Czech there seems to be a similar performance of the monolingual and the multilingual models, while in all remaining parliaments, there is a consistent benefit of training on all available languages. Given these results, we can conclude that if one wanted to

language	R^2		MAE	
	mono	multi	mono	multi
bcs	0.699	0.737	0.644	0.572
cz	0.564	0.560	0.706	0.665
en	0.707	0.741	0.652	0.599
sk	0.646	0.681	0.665	0.593
sl	0.512	0.520	0.708	0.667

Table 5: Experiments on comparing performance when training on the target language vs. training on all available languages.

annotate the sentiment in a specific parliament for which there is training data available, better results might still be obtained with additional data, written in different languages and coming from different parliaments.

This result also shows that our annotation guidelines were detailed enough for the annotators in the different languages to have comparable annotation criteria, thereby rendering annotations in different languages useful to each other.

5. Conclusion

In this paper, we have presented a new dataset, consisting of sentences coming from seven different parliaments, manually annotated with a six-level schema for sentiment. This is the first of such datasets available for parliamentary proceedings' data. We show the inter-annotator agreement to be reasonably high for such an endeavor. We share 2,600 instances per parliament (group), the Bosnian, Croatian, and Serbian parliaments forming a single BCS parliament group, the remaining four parliaments being those of the Czech Republic, United Kingdom, Slovakia, and Slovenia. Aside from these five training datasets, we also share two additional test sets, one from the BCS group and another from the United Kingdom. The data are shared via the CLARIN.SI repository¹².

In our experiments, we answer three main research questions. The first question relates to whether additional pre-training of a transformer model on raw parliamentary data would improve the performance on the task, which proves to be correct, and therefore, we also present the new XLM-R-parla model¹³, especially suited for

¹²<http://hdl.handle.net/11356/1868>

¹³<https://huggingface.co/classla/>

parliamentary text processing. Whether the model is more potent in the processing of political texts in general, will have to be tested in follow-up work.

The second question we tackle is how well we can expect the final, fine-tuned model, named XLM-R-ParlaSent¹⁴, to perform on languages and parliaments not seen during fine-tuning. We show that the model is very robust to unseen languages and parliaments with no or minor drop in performance. The limitation of this insight is that the languages and parliaments we check this presumption on are linguistically and traditionally rather related to the remaining languages and parliaments in the training data, so caution is advised if more distant languages or parliaments were to be annotated with the XLM-R-ParlaSent model.

The third question relates to whether a model performs better on a specific language and parliament if it is trained on that parliament's data only, or whether the additional, four times larger dataset, coming from different languages and parliaments, is more beneficial. We show that the multilingual multi-parliamentary approach performs better, which is a direct signal that our annotations are not of high quality inside parliaments only, as measured via the inter-annotator agreement, but also between parliaments and languages.

We consider this work to be a very important step in setting up a more robust approach to sentiment analysis of political texts beyond sentiment lexicon approaches, which will find many applications in the downstream research of political and parliamentary communications. It is part of a more general effort to democratize the latest advancements in natural language processing and their relevance in humanities and social sciences.

6. Acknowledgments

The research presented in this paper was conducted within the research project “Basic Research for the Development of Spoken Language Resources and Speech Technologies for the Slovenian Language” (J7-4642), the research project “Embeddings-based techniques for Media Monitoring Applications” (L2-50070, co-funded by the Kliping d.o.o. agency) and within the research programme “Language resources and technologies for Slovene” (P6-0411), all funded by the Slovenian Research and Innovation Agency (ARIS).

xlm-r-parla

¹⁴<https://huggingface.co/classla/xlm-r-parlasent>

7. Bibliographical References

- Gavin Abercrombie and Riza Batista-Navarro. 2018a. ‘Aye’ or ‘No’? Speech-level Sentiment Analysis of Hansard UK Parliamentary Debate Transcripts. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Gavin Abercrombie and Riza Batista-Navarro. 2020. Sentiment and position-taking analysis of parliamentary debates: a systematic literature review. *Journal of Computational Social Science*, 3(1):245–270.
- Gavin Abercrombie and Riza Theresa Batista-Navarro. 2018b. Identifying Opinion-Topics and Polarity of Parliamentary Debate Motions. In *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 280–285, Brussels, Belgium. Association for Computational Linguistics.
- Alan I. Abramowitz and Steven Webster. 2016. The rise of negative partisanship and the nationalization of U.S. elections in the 21st century. *Electoral Studies*, 41(March):12–22.
- Shakhnaz Akhmedova, Eugene Semekin, and Vladimir Stanovov. 2018. Co-operation of Biology Related Algorithms for Solving Opinion Mining Problems by Using Different Term Weighting Schemes. In Kurosh Madani, Dimitri Peaucelle, and Oleg Gusikhin, editors, *Informatics in Control, Automation and Robotics : 13th International Conference, ICINCO 2016 Lisbon, Portugal, 29-31 July, 2016*, Lecture Notes in Electrical Engineering, pages 73–90. Springer International Publishing, Cham.
- Mohit Bansal, Claire Cardie, and Lillian Lee. 2008. The Power of Negative Thinking: Exploiting Label Disagreement in the Min-cut Classification Framework. In *Coling 2008: Companion volume: Posters*, pages 15–18, Manchester, UK. Coling 2008 Organizing Committee.
- Vuk Batanović, Miloš Cvetanović, and Boško Nikolić. 2020. A versatile framework for resource-limited sentiment articulation, annotation, and analysis of short texts. *PLOS ONE*, 15(11):e0242050. Publisher: Public Library of Science.
- Adam Bermingham and Alan F. Smeaton. 2011. On using Twitter to monitor political sentiment and predict election results. Chiang Mai, Thailand.

- Adam Bonica. 2016. [A Data-Driven Voter Guide for U.S. Elections: Adapting Quantitative Measures of the Preferences and Priorities of Political Elites to Help Voters Learn About Candidates](#). *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 2(7):11–32. Publisher: Russell Sage Foundation.
- Andrea Ceron, Luigi Curini, and Stefano M Iacus. 2019. *Politics and big data: nowcasting and forecasting elections with social media*. Routledge, Abingdon, New York. OCLC: 1084404675.
- Yanqing Chen and Steven Skiena. 2014. [Building Sentiment Lexicons for All Major Languages](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 383–389, Baltimore, Maryland. Association for Computational Linguistics.
- Yanqing Chen and Steven Skiena. 2016. False-friend detection and entity matching via unsupervised transliteration. *arXiv preprint arXiv:1611.06722*.
- Paul A. Chilton. 2004. *Analysing political discourse: theory and practice*. Routledge, London.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- James N. Druckman, Samara Klar, Yanna Krupnikov, Matthew Levendusky, and John Barry Ryan. 2021. [Affective polarization, local contexts and public opinion in America](#). *Nature Human Behaviour*, 5(1):28–38. Number: 1 Publisher: Nature Publishing Group.
- Rory Duthie and Katarzyna Budzyska. 2018. [A deep modular RNN approach for ethos mining: 27th International Joint Conference on Artificial Intelligence, IJCAI 2018](#). *Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI 2018*, pages 4041–4047. Publisher: International Joint Conferences on Artificial Intelligence.
- Tomaž Erjavec, Matyáš Kopp, Maciej Ogrodniczuk, Petya Osenova, Darja Fišer, Hannes Pirker, Tanja Wissik, Daniel Schopper, Martin Kirnbauer, Nikola Ljubešić, Peter Rupnik, Michal Mochtak, Henk van der Pol, Griet Depoorter, Kiril Simov, Vladislava Grigorova, Ilko Grigorov, Bart Jongejan, Dorte Haltrup Hansen, Costanza Navarretta, Martin Mölder, Neeme Kahusk, Kadri Vider, Nuria Bel, Iván Antiba-Cartazo, Marilina Pisani, Rodolfo Zevallos, Adina Ioana Vladu, Carmen Magariños, Daniel Bardanca, Mario Barcala, Marcos Garcia, María Pérez Lago, Pedro García Louzao, Ainhoa Vivel Couso, Marta Vázquez Abuín, Noelia García Díaz, Adrián Vidal Miguéns, Elisa Fernández Rei, Xosé Luís Regueira, Sascha Diwersy, Giancarlo Luxardo, Matthew Coole, Paul Rayson, Amanda Nwadukwe, Dimitris Gkoumas, Vassilis Papavassiliou, Prokopis Prokopidis, Maria Gavriilidou, Stelios Piperidis, Noémi Ligeti-Nagy, Kinga Jelencsik-Mátyus, Zsófia Varga, Réka Dodé, Starkaður Barkarson, Tommaso Agnoloni, Roberto Bartolini, Francesca Frontini, Simonetta Montemagni, Valeria Quochi, Giulia Venturi, Manuela Ruisi, Carlo Marchetti, Roberto Battistoni, Roberts Dargis, Ruben van Heusden, Maarten Marx, Lars Magne Tungland, Michał Rudolf, Bartłomiej Nitoń, José Aires, Amália Mendes, Aida Cardoso, Rui Pereira, Väinö Yrjänäinen, Fredrik Mohammadi Norén, Måns Magnusson, Johan Jarlbrink, Katja Meden, Andrej Pančur, Mihael Ojsteršek, Çağrı Çöltekin, and Anna Kryvenko. 2023. [Multilingual comparable corpora of parliamentary debates ParlaMint 3.0](#). <https://www.clarin.eu/content/parlamint>. Accepted: 2023-07-04T13:40:35Z Publisher: CLARIN ERIC.
- Eli J. Finkel, Christopher A. Bail, Mina Cikara, Peter H. Ditto, Shanto Iyengar, Samara Klar, Lilliana Mason, Mary C. McGrath, Brendan Nyhan, David G. Rand, Linda J. Skitka, Joshua A. Tucker, Jay J. Van Bavel, Cynthia S. Wang, and James N. Druckman. 2020. [Political sectarianism in America](#). *Science*, 370(6516):533–536. Publisher: American Association for the Advancement of Science.
- René D. Flores. 2017. Do Anti-Immigrant Laws Shape Public Sentiment? A Study of Arizona’s SB 1070 Using Twitter Data. *American Journal of Sociology*, 123(2):333–384.
- R. Kelly Garrett, Shira Dvir Gvirsman, Benjamin K. Johnson, Yariv Tsfati, Rachel Neo, and Aysenur Dal. 2014. [Implications of pro- and counterattitudinal information exposure for affective polarization](#). *Human Communication Research*, 40(3):309–332. Place: United Kingdom Publisher: Wiley-Blackwell Publishing Ltd.
- Jacques Gerstlé and Alessandro Nai. 2019. [Negativity, emotionality and populist rhetoric in election campaigns worldwide, and their effects on media attention and electoral success](#).

- European Journal of Communication*, 34(4):410–444. Publisher: SAGE Publications Ltd.
- Goran Glavaš, Federico Nanni, and Simone Paolo Ponzetto. 2017. [Unsupervised Cross-Lingual Scaling of Political Texts](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2*, volume 2, pages 688–693, Valencia.
- Goran Glavaš, Jan Šnajder, and Bojana Dalbelo Bašić. 2012. Semi-supervised Acquisition of Croatian Sentiment Lexicon. In Petr Sojka, Aleš Horák, Ivan Kopeček, and Karel Pala, editors, *Text, speech and dialogue: 15th international conference, TSD 2012, Brno, Czech Republic, September 3-7, 2012: proceedings*, pages 166–173. Springer, Berlin, Heidelberg.
- Tanvi Hardeniya and D. A. Borikar. 2016. [Dictionary Based Approach to Sentiment Analysis - A Review](#). *International Journal of Advanced Engineering, Management and Science*.
- Timo Honkela, Jaakko Korhonen, Krista Lagus, and Esa Saarinen. 2014. [Five-Dimensional Sentiment Analysis of Corpora, Documents and Words](#). In *Advances in Self-Organizing Maps and Learning Vector Quantization*, Advances in Intelligent Systems and Computing, pages 209–218, Cham. Springer International Publishing.
- Swen Hutter, Edgar Grande, and Hanspeter Kriesi. 2016. *Politicising Europe: integration and mass politics*. Cambridge University Press, Cambridge. OCLC: 1280510451.
- Shanto Iyengar and Stephen Ansolabehere. 1995. *Going negative*. Free Press, New York.
- Shanto Iyengar, Gaurav Sood, and Yphtach Lelkes. 2012. [Affect, Not Ideology: A Social Identity Perspective on Polarization](#). *Public Opinion Quarterly*, 76(3):405–431.
- Ruud Koopmans and Paul Statham. 2006. [Political Claims Analysis: Integrating Protest Event and Political Discourse Approaches](#). *Mobilization: An International Quarterly*, 4(2):203–221.
- George Lakoff. 2004. *Don't Think of an Elephant*. Chelsea Green, White River Junction.
- Harold Dwight Lasswell. 1927. *Propaganda technique in the world war*. Peter Smith, New York.
- Dilin Liu and Lei Lei. 2018. [The appeal to political sentiment: An analysis of Donald Trump's and Hillary Clinton's speech themes and discourse strategies in the 2016 US presidential election](#). *Discourse, Context & Media*, 25:143–152.
- Nikola Ljubešić and Kaja Dobrovoljc. 2019. [What does neural bring? analysing improvements in morphosyntactic annotation and lemmatisation of Slovenian, Croatian and Serbian](#). In *Proceedings of the 7th Workshop on Balto-Slavic Natural Language Processing*, pages 29–34, Florence, Italy. Association for Computational Linguistics.
- Lena Masch and Oscar W. Gabriel. 2020. [How Emotional Displays of Political Leaders Shape Citizen Attitudes: The Case of German Chancellor Angela Merkel](#). *German Politics*, 29:158–179. Publisher: Routledge _eprint: <https://doi.org/10.1080/09644008.2019.1657096>.
- Lilliana Mason. 2015. ["I Disrespectfully Agree": The Differential Effects of Partisan Sorting on Social and Issue Polarization](#). *American Journal of Political Science*, 59(1):128–145. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajps.12089>.
- Lilliana Mason. 2018. *Uncivil Agreement: How Politics Became Our Identity*. University of Chicago Press, Chicago.
- Michal Mochtak, Josip Glaurdić, and Christophe Lesschaeve. 2020. Talking War: Representation, Veterans and Ideology in Post-War Parliamentary Debates. *Government and Opposition*, 57(1):148–170.
- Michal Mochtak, Josip Glaurdić, and Christophe Lesschaeve. 2022a. [Talking War: Representation, Veterans and Ideology in Post-War Parliamentary Debates](#). *Government and Opposition*, 57(1):148–170.
- Michal Mochtak, Peter Rupnik, and Nikola Ljubešić. 2022b. The parlament-bcs dataset of sentiment-annotated parliamentary debates from bosnia-herzegovina, croatia, and serbia. *arXiv preprint arXiv:2206.00929*.
- Saif M. Mohammad. 2021. Sentiment analysis: Automatically detecting valence, emotions, and other affectual states from text. <https://arxiv.org/abs/2005.11882>.
- Ines Montani and Matthew Honnibal. [Prodigy: A modern and scriptable annotation tool for creating training data for machine learning models](#). *Prodigy*.
- Igor Mozetič, Miha Grčar, and Jasmina Smilović. 2016. [Multilingual Twitter Sentiment Classification: The Role of Human Annotators](#). *PLOS ONE*, 11(5):e0155036. Publisher: Public Library of Science.
- Nona Naderi and Graeme Hirst. 2016. Argumentation Mining in Parliamentary

- Discourse. In *Principles and Practice of Multi-Agent Systems*, pages 16–25, Cham. Springer International Publishing.
- Minh Van Nguyen, Viet Dac Lai, Amir Poursan Ben Veyseh, and Thien Huu Nguyen. 2021. [Trankit: A Light-Weight Transformer-based Toolkit for Multilingual Natural Language Processing](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 80–90, Online. Association for Computational Linguistics.
- Obinna Onyimadu, Keiichi Nakata, Tony Wilson, David Macken, and Kecheng Liu. 2013. [Towards Sentiment Analysis on Parliamentary Debates in Hansard](#). In *Revised Selected Papers of the Third Joint International Conference on Semantic Technology - Volume 8388*, JIST 2013, pages 48–50, Berlin, Heidelberg. Springer-Verlag.
- Erica Owen. 2017. [Exposure to Offshoring and the Politics of Trade Liberalization: Debate and Votes on Free Trade Agreements in the US House of Representatives, 2001–2006](#). *International Studies Quarterly*, 61(2):297–311.
- James A. Piazza. 2020. [Politician hate speech and domestic terrorism](#). *International Interactions*, 46(3):431–453. Publisher: Routledge.
- Keval Pipalia, Rahul Bhadja, and Madhu Shukla. 2020. [Comparative Analysis of Different Transformer Based Architectures Used in Sentiment Analysis](#). In *2020 9th International Conference System Modeling and Advancement in Research Trends (SMART)*, pages 411–415.
- Bingham G. Powell. 2004. [Political Representation in Comparative Politics](#). 7:273–296. Publisher: Annual Reviews.
- Sven-Oliver Proksch, Will Lowe, Jens Wäckerle, and Stuart Soroka. 2019. Multilingual Sentiment Analysis: A New Approach to Measuring Conflict in Legislative Speeches. *Legislative Studies Quarterly*, 44(1):97–131.
- Christian Rauh. 2018. Validating a sentiment dictionary for German political language—a workbench note. *Journal of Information Technology & Politics*, 15(4):319–343.
- Tyler Rinker. 2017. [GitHub - trinker/entity: Easy named entity extraction](#).
- Zaher Salah. 2015. [Machine learning and sentiment analysis approaches for the analysis of Parliamentary debates](#). Ph.D. thesis, University of Liverpool. Publisher: University of Liverpool.
- Edward A. Shils and Morris Janowitz. 1948. [Cohesion and Disintegration in the Wehrmacht in World War II](#). *Public Opinion Quarterly*, 12(2):315. Publisher: Oxford University Press (OUP).
- Juan M. Soler, Fernando Cuartero, and Manuel Roblizo. 2012. [Twitter as a Tool for Predicting Elections Results](#). In *2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 1194–1200.
- Chul Sung, Tejas Dhamecha, Swarnadeep Saha, Tengfei Ma, Vinay Reddy, and Rishi Arora. 2019. [Pre-training BERT on domain resources for short answer grading](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6071–6075, Hong Kong, China. Association for Computational Linguistics.
- Luka Terčon and Nikola Ljubešić. 2023. [Classstanza: The next step for linguistic processing of south slavic languages](#).
- Matt Thomas, Bo Pang, and Lillian Lee. 2006. [Get out the vote: Determining support or opposition from Congressional floor-debate transcripts](#). *COLING/ACL 2006 - EMNLP 2006: 2006 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, pages 327–335. _eprint: 0607062.
- Andranik Tumasjan, Timm Sprenger, Philipp Sandner, and Isabell Welp. 2010. Predicting elections with twitter: What 140 characters reveal about political sentiment. *Proceedings of the International AAAI Conference on Web and Social Media*, 4(1):178–185.
- Andrea Volkens, Tobias Burst, Werner Krause, Pola Lehmann, Matthieß Theres, Nicolas Merz, Sven Regel, Bernhard Weßels, and Lisa Zehnter. 2020. [Manifesto Project Database](#).
- Steven W. Webster and Bethany Albertson. 2022. [Emotion and Politics: Noncognitive Psychological Biases in Public Opinion](#). *Annual Review of Political Science*, 25(1):401–418. _eprint: <https://doi.org/10.1146/annurev-polisci-051120-105353>.
- Tobias Widmann. 2021. [How Emotional Are Populists Really? Factors Explaining Emotional Appeals in the Communication of Political Parties](#). *Political Psychology*, 42(1):163–181. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/pops.12693>.

Lori Young and Stuart Soroka. 2012. Affective news: The automated coding of sentiment in political texts. *Political Communication*, 29(2):205–231.

8. Language Resource References

Erjavec, Tomaž and Kopp, Matyáš and Ogrodniczuk, Maciej and Osenova, Petya and Fišer, Darja and Pirker, Hannes and Wissik, Tanja and Schopper, Daniel and Kirnbauer, Martin and Ljubešić, Nikola and Rupnik, Peter and Mochtak, Michal and Pol, Henk van der and Depoorter, Griet and Simov, Kiril and Grigorova, Vladislava and Grigorov, Ilko and Jongejan, Bart and Haltrup Hansen, Dorte and Navarretta, Costanza and Mölder, Martin and Kahusk, Neeme and Vider, Kadri and Bel, Nuria and Antiba-Cartazo, Iván and Pisani, Marilina and Zevallos, Rodolfo and Vladu, Adina Ioana and Magariños, Carmen and Bardanca, Daniel and Barcala, Mario and Garcia, Marcos and Pérez Lago, María and García Louzao, Pedro and Vivel Couso, Ainhoa and Vázquez Abuín, Marta and García Díaz, Noelia and Vidal Miguéns, Adrián and Fernández Rei, Elisa and Regueira, Xosé Luís and Diwersy, Sascha and Luxardo, Giancarlo and Coole, Matthew and Rayson, Paul and Nwadukwe, Amanda and Gkoumas, Dimitris and Papavassiliou, Vassilis and Prokopidis, Prokopis and Gavriilidou, Maria and Piperidis, Stelios and Ligeti-Nagy, Noémi and Jelencsik-Mátyus, Kinga and Varga, Zsófia and Dodé, Réka and Barkarson, Starkaður and Agnoloni, Tommaso and Bartolini, Roberto and Frontini, Francesca and Montemagni, Simonetta and Quochi, Valeria and Venturi, Giulia and Ruisi, Manuela and Marchetti, Carlo and Battistoni, Roberto and Dargis, Roberts and van Heusden, Ruben and Marx, Maarten and Tunland, Lars Magne and Rudolf, Michał and Nitoń, Bartłomiej and Aires, José and Mendes, Amália and Cardoso, Aida and Pereira, Rui and Yrjänäinen, Väinö and Norén, Fredrik Mohammadi and Magnusson, Måns and Jarlbrink, Johan and Meden, Katja and Pančur, Andrej and Ojsteršek, Mihael and Çöltekin, Çağrı and Kryvenko, Anna. 2023. *Multilingual comparable corpora of parliamentary debates ParlaMint 3.0*. [\[link\]](#).

Koehn, Philipp. 2011. *European parliament proceedings parallel corpus 1996-2011*. [\[link\]](#).

Mochtak, Michal. 2022. *SVKCorp: Corpus of Debates in the National Council of the Slovak Republic*. [\[link\]](#).

Mochtak, Michal and Glaurdić, Josip and Lesschaeve, Christophe. 2022a. *CROCorp: Corpus of Parliamentary Debates in Croatia (v1.1.1)*. [\[link\]](#).

Mochtak, Michal and Glaurdić, Josip and Lesschaeve, Christophe. 2022b. *SRBCorp: Corpus of Parliamentary Debates in Serbia (v1.1.1)*. [\[link\]](#).

Mochtak, Michal and Glaurdić, Josip and Lesschaeve, Christophe and Muharemović, Ensar. 2022c. *BiHCorp: Corpus of Parliamentary Debates in Bosnia and Herzegovina (v1.1.1)*. [\[link\]](#).