

BasqueParl: A Bilingual Corpus of Basque Parliamentary Transcriptions

Nayla Escribano¹, Jon Ander González¹, Julen Orbegozo-Terradillos²,
Ainara Larrondo-Ureta², Simón Peña-Fernández²,
Olatz Perez-de-Viñaspre¹ and Rodrigo Agerri¹

¹ HiTZ Center - Ixa, University of the Basque Country UPV/EHU

² Gureiker, University of the Basque Country UPV/EHU

{nayla.escribano, olatz.perezdevinaspre, rodrigo.agerri}@ehu.eus

Abstract

Parliamentary transcripts provide a valuable resource to understand the reality and know about the most important facts that occur over time in our societies. Furthermore, the political debates captured in these transcripts facilitate research on political discourse from a computational social science perspective. In this paper we release the first version of a newly compiled corpus from Basque parliamentary transcripts. The corpus is characterized by heavy Basque-Spanish code-switching, and represents an interesting resource to study political discourse in contrasting languages such as Basque and Spanish. We enrich the corpus with metadata related to relevant attributes of the speakers and speeches (language, gender, party...) and process the text to obtain named entities and lemmas. The obtained metadata is then used to perform a detailed corpus analysis which provides interesting insights about the language use of the Basque political representatives across time, parties and gender.

Keywords: Parliamentary Transcripts, Code-switching, Computational Social Science, Information Extraction

1. Introduction

Parliaments and chambers of public representatives are official agoras for the presentation of ideas and their dialectical confrontation, where people chosen to represent certain groups interact in the public-political space. Minutes gathered in each of these agoras, which compile the aforementioned interactions in public institutions, can be seen as the “black box” of a given society. In this sense, we believe that speeches from public representatives may allow us to understand and interpret the reality of a given historical moment.

In this context, technological advances in Natural Language Processing (NLP), machine learning and other computational methods allow for parliamentary rhetoric to be analysed from a multidisciplinary perspective. It should be observed that public institutions have always generated large amounts of textual data (debates, laws, parliamentary discourses) that have received a special interest since the start of digitization and the emergence of the Web in the last decade of the past century. Digitization of debates allows institutions to increase their public and media impact, generating at the same time valuable textual resources that can be analyzed by computational social science and NLP research. Specifically, relevant corpora have been gathered for different NLP tasks such as sentiment analysis (Thomas et al., 2006; Abercrombie and Batista-Navarro, 2020) or machine translation (Roukos et al., 1995; Koehn, 2005; Hajlaoui et al., 2014).

With this multidisciplinary approach in mind, we gathered and processed a corpus of Basque parliamentary transcriptions for public research. Furthermore, we analyzed the representatives’ speeches across different data attributes that might be of interest for the general public such as language use, gender and party. Indeed,

these analyses could reflect whether political groups and concrete speakers do or do not act in parliament according to their manifested ideas. It could also serve to analyze to what extent the Basque parliament is a reflection of the society it represents.

For a bilingual community such as the Basque Country, analyzing language use is a useful aspect to consider in this regard. Results from an official sociolinguistic poll undertaken in 2016 indicated that 13.4% of the people in the Basque Autonomous Community spoke more Basque than Spanish, even if 33.9% were considered to have Basque language skills¹. In this sense, our newly created corpus of parliamentary transcriptions would allow us to compare language use among the general public with respect to the linguistic behaviour of its political representatives.

In this paper we present BasqueParl, a new bilingual corpus for automatic political discourse analysis. It covers transcriptions from the Parliament of the Basque Autonomous Community for eight years and two legislative terms (2012-2020). Its main characteristic is its Basque-Spanish code-switching speeches, which have been processed to identify the language of each speech fragment. Thus, the contributions of this work include:

1. The creation of BasqueParl, a new publicly available 14M word bilingual corpus for political discourse analysis.
2. Enriching the corpus with metadata (language of each speech fragment and speaker’s year of birth, gender and party) and performing neural lemmatization and NER for Basque and Spanish.

¹https://www.irekia.euskadi.eus/uploads/attachments/9954/VI_INK_SOZLG-EH_eus.pdf

3. A detailed data analysis showing that: (i) Basque is often used in speeches but barely to convey speech content, (ii) women are underrepresented in word production, although this trend has reversed in the last years, among other conclusions.
4. The release of the corpus for public research².

We describe the Basque Parliament and relevant aspects of the legislative terms covered by the corpus in Section 2. We discuss related work in Section 3. Methods employed to process the corpus are presented in Section 4, while the corpus processing is explained in Section 5. Finally, Section 6 shows the main results of the data analysis performed on the corpus.

2. The Basque Parliament

The term Basque Country (Euskal Herria in Basque) is usually used to refer to a number of territories across the Pyrenees between France and Spain, of which the Basque Autonomous Community (EAE) is a region.

The Basque Parliament constitutes the legislative body of the Basque Autonomous Community and controls the activity of the Basque Government. Sessions are guided and moderated by the president of the parliament, who usually uses short and frequent utterances for these purposes. Speeches can be produced in either Basque or Spanish, being many debates bilingual. Table 1 presents an example of a bilingual code-switching speech where the speaker changes from a language to another even inside sentences (normal font for Basque, true type for Spanish). It should be noted that the large majority of the bilingual speeches do not include as much Basque usage as this one.

Considered transcriptions include parliamentary sessions from December 3, 2012 to February 7, 2020, covering two legislative terms, namely, 2012-2016 and 2016-2020. Along these years, various political parties have taken part in debates: *Euzko Alderdi Jeltzalea-Partido Nacionalista Vasco (EAJ-PNV)*, *Euskal Herria Bildu (EH Bildu)*, *Partido Socialista de Euskadi-Euskadiko Ezkerra (PSE-EE)*, *Partido Popular (PP)*, *Elkarrekin Podemos (EP)* and *Unión, Progreso y Democracia (UPyD)*. However, other parties and institutions also took part, such as the ombudsman of the Basque Country (Ararteko) or *Ezker Batua (EB)* by means of a regional parliamentary speaker.

EAJ-PNV is the main government party and the party of the parliament president throughout all the considered years, and is therefore the author of the largest number of speeches. UPyD had representation until 2016, and that year EP gained access to the parliament. The rest of the groups maintain their presence in parliament during all the covered time.

3. Related Work

The analysis of parliamentary discourse has received special interest in the last years and more transcrip-

²<https://github.com/ixa-ehu/basqueparl>

tions and corpora are being released. Some parliaments make their transcriptions and voting data freely available for public observation, such as the UK Parliament³, whereas other parliamentary sessions are released by independent sources, like GovTrack⁴ for the US. However, transcriptions of the Basque Parliament sessions are not publicly available.

Several corpora have been built on this kind of transcripts for different purposes. There has been a special interest in tagging Sentiment Analysis on debate speeches, given the vast presence of opinions towards proposed motions. With this objective, Abercrombie and Batista-Navarro (2020) present the ParlVote dataset, comprising English speeches of the UK Parliament, and perform different experiments on Sentiment Analysis. Similar approaches have been made on US congressional floor (Thomas et al., 2006).

Parliamentary transcriptions have also been used for Machine Translation by building multilingual parallel corpora. The Canadian Hansard Corpus contains speeches from the Canadian Parliament in French and English (Roukos et al., 1995), whereas EuroParl (Koehn, 2005) and DCEP (Hajlaoui et al., 2014) cover debates from the European Parliament nowadays in 21 and 23 European languages, respectively.

Among non-parallel multilingual corpora, ParlaMint (Erjavec et al., 2021) gathers sessions from 17 parliaments in their respective European languages (including Spanish parliament transcriptions), with transcriptions from 2015 to mid-2020 and special attention the COVID-19 period. Additionally, DutchParl (Marx and Schuth, 2010) covers Dutch transcriptions of parliaments from the Netherlands, Flanders and Belgium, including Dutch-French bilingual speeches from the Belgium Federal texts. This last bilingual example is perhaps the closest to our Basque-Spanish code-switching speeches which characterized our newly released corpus. In this sense, BasqueParl provides the first resource of this kind for a large language such as Spanish, the second-most native spoken language (after Chinese), and Basque, a pre-indoeuropean isolate language spoken by around 700K people.

4. Methodology

We used different systems and techniques to process and enrich the corpus with the aim of preparing it for data analysis. These resources are described below.

4.1. Language Identification

Since the corpus is bilingual and speeches switch from a language to another as shown in Table 1, we decided to identify the language of the corpus units. To that end, we used the *langdetect* language detection library⁵. *Langdetect* compares the n-grams of a text to

³<https://www.theyworkforyou.com/>

⁴<https://www.govtrack.us/>

⁵<https://github.com/fedelopez77/langdetect>

Bai, zure baimenarekin hemendik.

Ba zure desioak, Guanache andrea, gureak ere badira. Harritu nau eta ez nau harritu hitza berriro hartzeak, zeren hitz egiten nengoen bitartean esan diozu albokoari *le voy a contestar*. *Le voy a contestar*, ondo iruditzen, zure eskubidean zaude, baino beno, ez dut uste inongo astakeriarik esan dudarik.

Gauzak egiten dira eta uste dut nik, nik ere eskubidea dudala Gobernuak eta beste erakundeek egiten dutena esateko. Zeren beti *ver el vaso medio vacío o medio lleno*, pues cambia un poco la perspectiva y vernos siempre en modo Gobierno, creo que no es nada objetivo. Se hacen cosas, se harán cosas y esta vez creo que me deberían reconocer que de la iniciativa primera a lo que hemos acordado, no nos hemos dejado nada o creo que casi nada. Entonces, bueno, sólo quería aclarar eso *eta eskerrak berriro*.

Eta ziur egon emakumea dokumentu horietan ez bada agertzen hitzetan, zeren uste dut hori ez dela garrantzitsuena, bai politiketan egongo dela eta dagoela.

Eskerrik asko.

Table 1: Example of bilingualism and code-switching in Basque Parliament speeches.

n-grams of previously built language profiles and provides the probabilities of the closest languages according to a distance metric.

4.2. Lemmatization and Named Entity Recognition

We performed lemmatization and Named Entity Recognition (NER) using Flair, which is both a deep learning system based on a BiLSTM architecture and a rather effective type of character-based contextual word embeddings (Akbik et al., 2018; Akbik et al., 2019). The system and embeddings have demonstrated high performance in Sequence Labelling tasks such as Part-of-Speech, NER or SRL.

Agerri et al. (2020) demonstrate that text representation models trained on an appropriate monolingual corpus for Basque outperform large multilingual transformer-based models such as mBERT (Devlin et al., 2019) or XLM-RoBERTa (Conneau et al., 2020). Similarly, Agerri (2020) presents the Flair-Oscar monolingual model for Spanish, which obtained the best results at the CAPITEL 2020 shared task for NER in Spanish (Zamorano and Anke, 2020). These two models have been used to perform lemmatization and NER for both languages in the BasqueParl corpus.

5. Corpus Processing

In order to prepare the corpus for data analysis, we performed several pre-processing steps using manual, rule-based methods and machine learning techniques. We also developed a demo to explore the pre-processed corpus.

5.1. Preprocessing

Basic information about each speech includes date, speech identifier, speaker, birth, gender and party. The first three fields were explicitly indicated in the original transcriptions, but we had to normalize variations (e.g. dates in different format or same person mentioned with various name forms) and correct mistakes (e.g. same speech identifier for two different speeches). In the case of gender, we built a set of rules to identify

it from text. For example, *andrea (Ms.)* in the original speaker name would suggest that the speaker is a woman. On the other hand, we extracted the speaker's year of birth and party by gathering a list of politicians and their parties from official sources⁶⁷. In exceptional cases in which some of this data was impossible to retrieve (e.g. the speaker name may be lost in the transcription process) or was not appropriate (e.g. the speech is assigned to an organization and not a person), we marked the corresponding field as none/unknown.

Additionally, we applied machine and deep learning techniques to identify language, lemmatize the text and detect named entities on originally separated speech paragraphs. We used *langdetect* to identify each paragraph written in Spanish, assuming that the non-Spanish paragraphs were Basque. In cases like Table 1, the system decides for each paragraph if the text is in Spanish or not. Then, we separated those paragraphs into sentences using the *segtok* tokenizer⁸ and, according to the language detected in the paragraph, applied the corresponding Flair-based lemmatization and NER models for Basque or Spanish. Due to the agglutinative morphology of Basque, lemmatization is required to obtain named entities without their inflected forms. For both languages, we only kept named entities referring to persons, locations or organizations.

We also developed a list of stopwords for each language removing lemmas and named entities which were not of interest for our data analysis. In addition to stopwords gathered from already existing lists for Basque⁹ and Spanish¹⁰, we filtered out lemmas appearing more than 1000 times. These lemmas mostly referred to

⁶https://www.legebiltzarra.eus/comparla/e_comparla_alf_ACT.html

⁷<https://www.jjggalava.eus/eu/web/jjggalava/batzarkideen-historikoa>

⁸<https://github.com/fnl/segtok>

⁹<https://github.com/stopwords-iso/stopwords-eu/blob/master/stopwords-eu.json>

¹⁰<https://github.com/6/stopwords-json/blob/master/dist/es.json>

Field	Description
Speech date	Date corresponding to the speech, e.g. <i>2020-02-07</i>
Speech id	Number that identifies the speech within its date, e.g. <i>3</i>
Paragraph id	Number that identifies the paragraph within its speech, e.g. <i>3</i>
Speaker	Family names of the speaker, including their position if any, e.g. <i>Tejeria Otermin LEHENDAKARIA</i>
Birth	Year of birth of the speaker, e.g. <i>1971</i>
Gender	Gender of the speaker, e.g. <i>female</i>
Party	Political group of the speaker, e.g. <i>EAJ-PNV</i>
Language	Language assigned to a paragraph, e.g. <i>Basque</i>
Text	Paragraph of the speech text
Lemmas	Lemmatized paragraph, with and without stopwords
Named entities	Named entities extracted from the paragraph, with and without stopwords

Table 2: Field data for each paragraph in the corpus.

Field	Category	Speeches	Paragraphs	Words	Lemmas	Entities
All	All	41,417	342,666	13,872,105	5,090,573	349,890
Language	Basque	34,571	133,599	2,938,061	1,375,686	122,439
	Spanish	18,016	209,067	10,934,044	3,714,887	227,451
Gender	Female	30,857	185,119	6,452,503	2,424,765	179,030
	Male	10,559	157,481	7,416,852	2,664,596	170,782
	None/Unknown	1	66	2,750	1,212	78
Party	EAJ-PNV	28,530	143,685	4,330,467	1,701,415	145,400
	EH Bildu	3,638	58,391	2,423,245	1,009,715	56,440
	PP	3,562	51,683	2,672,062	858,798	56,376
	PSE-EE	2,834	44,738	2,327,350	769,767	43,994
	EP	1,467	24,235	1,171,948	410,966	24,738
	UPyD	1,368	19,110	909,007	322,220	22,270
	ARARTEKO	10	627	30,382	14,535	524
	EB	1	17	746	283	23
	None/Unknown	7	180	6,898	2,874	125
Speaker	President	22,772	60,326	753,231	315,138	62,317

Table 3: Distribution of the corpus data.

ubiquitous terms related to the Parliament or the Government with low semantic meaning for this particular corpus.

5.2. Quantitative description of the corpus

Table 2 illustrates the information included for each paragraph in BasqueParl. Furthermore, Table 3 shows the distribution of the corpus data, where words are whitespace tokenized and lemmas and entities correspond to those without stopwords. Words inherit all the field data from their paragraph, that is, if a paragraph is set to a language, all its words are also set to that language, even if a word or a sentence belongs to the other one. The same applies to lemmas and entities. All the information is distributed by language, gender and party. We also present the data for the president of the Basque Parliament, who is the author of 55% of speeches, 18% of paragraphs and 5% of words. It should be noted that a single speech might consist of several paragraphs in both Basque and Spanish. We count a speech as belonging to either language if that




Figure 1: Number of speeches over time.

language has been detected in at least one paragraph of the speech. Thus, the same speech can be counted as Basque and Spanish, which means that the sum of Basque and Spanish speeches is larger than the total number of speeches. In addition, the length of the speeches can vary significantly, from 1 paragraph to 236. These considerations will allow us to contrast language, gender and party data at speech and word level

in Section 6.

Figure 1 presents the distribution of speeches along the considered years. The lower numbers of speeches in 2012, 2016 and 2020 reflect changes of legislative terms and, in the case of the first and last years, the beginning and the end of the corpus.

5.3. Demo

We provide the BasqueParl demo¹¹ which allows to explore the results of pre-processing and data analysis according to the fields described in Table 2, such as date, speaker, gender, party or language. Firstly, it shows speech examples and lemma and entity frequencies for the selected categories. Secondly, it provides topic modelling based on the LDA model (Blei et al., 2003), considering documents as the non-stopword lemmas of each month. Finally, it displays scattertext plots¹² of non-stopword lemmas to compare two distinct selections of categories.

6. Data Analysis

The distribution of the corpus data described in Table 3 allows us to perform various analyses by crossing the information of each of the field types. This section reports the main results obtained from such analysis.

6.1. Language

An important aspect for a bilingual society is to study the language use by crossing it with other field types such as party or gender. Figure 2 shows the percentages of language use at speech and word level for the full corpus and ignoring the texts from the president of the parliament. As we mentioned before, a speech belongs to either language if at least one of its paragraphs belongs to that language, being possible for a speech to be in Basque and Spanish at the same time. The large amount of Basque speeches in overall (“All”) suggests that most of the speeches have at least a Basque paragraph, compared to those that have at least a Spanish paragraph. On the contrary, the number of Basque words is significantly lower than in Spanish. As observed in a manual inspection of the corpus, this fact suggests that, even if Basque is used in most of the speeches, the most important content is usually conveyed in Spanish.

The distribution of speeches vary considerably when we filter out those texts belonging to the president of the parliament (“No pres.”), which consist mainly of short and frequent utterances in Basque like turn-takings or calls to order. Basque speeches decrease substantially while Spanish speeches double. This suggests that, although both languages are used often at least once in a speech, Spanish is more common. In contrast, this phenomenon is not reflected when we analyze language use at word level, where the percent-




Figure 2: Language use at speech and word level.

ages in language use remain independently of whether we consider the president’s speeches or not. This data seems to be the most realistic and shows a gap of more than 60 percentage points between Basque and Spanish in language use.

Bilingual	Passive	None
33.9%	19.1%	47.0%

Table 4: Basque language skills in EAE in 2016.

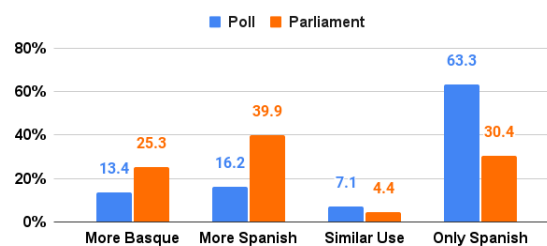


Figure 3: Language preference in EAE according to Sociolinguistic Poll 2016 and in parliament.

Table 4 presents the results of Basque language skills according to the VI Sociolinguistic Poll from 2016¹³, whereas Figure 3 reports language preference in the Basque Autonomous Community (EAE) from the same poll and in parliament at word level. “More Spanish” collapses two poll categories: “less Basque than Spanish” and “some Basque”. In the case of parliament results, “More Basque” corresponds to the percentage of speakers using Basque in more than 55% of their words, “More Spanish” to Spanish use between 55% and 95%, similar use to language use between 45% and 50% and “Only Spanish” to speakers using Spanish in more than 95% of their words. It should be noted that citizens that speak Basque generally also speak Spanish. While language skills are usually higher than overall use according to the poll, the result of our analysis seems to suggest that citizens’ use of Basque is lower than that of their political representatives. In fact,

¹¹<http://legebiltzarra.ixea.eus>

¹²<https://github.com/JasonKessler/scattertext>

¹³https://www.irekia.euskadi.eus/uploads/attachments/9954/VI_INK_SOZLGE-EH_eus.pdf




Figure 4: Language use by gender at speech level.




Figure 5: Language use by gender at word level.




Figure 6: Language use in each party at speech level.

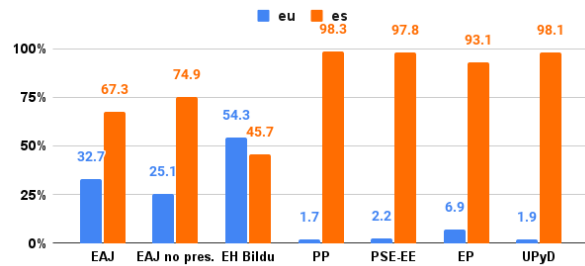


Figure 7: Language use in each party at word level.

speakers that use more Basque than Spanish overcome citizens' result by more than 10 points. However, on a closer look, although the percentage of representatives using only Spanish is half of citizen's ratio, their preference for Spanish doubles the EAE percentage.

6.1.1. Language by gender

This large difference between speech level and word level in language use remains if we look at gender, as illustrated by Figures 4 and 5. Almost all speeches produced by women contain a Basque paragraph and not even a third a Spanish one, while men usually produce Spanish speeches and less often Basque ones. Again, the percentages change if we look at word numbers: the rates of language use get more similar between genders, and Spanish becomes the most frequent by more than 40 percentage points among women and 70 among men. If we ignore president's texts, female and male texts get much closer, although women still tend to use Basque more often than men.

6.1.2. Language by party

The data of the language use by the main political groups¹⁴ in Figures 6 and 7 reflect the same behaviour as for the overall language use. Thus, while Basque use in terms of speech is quite high (especially for EAJ-PNV, EH Bildu and EP), those rates drop wildly at word level, especially for the Spanish unionist parties (PP, PSE-EE and UPyD), for which the use of Spanish is rather non-existent. In the case of EAJ-PNV (conservative Basque Nationalist) without the president, there is also a substantial rise at the use of Spanish at speech level. In terms of words, only one party shows a larger

percentage of Basque use with respect to Spanish (EH Bildu - Left Basque Pro-independence), although the numbers are quite balanced, despite contrary popular opinion. Summarizing, four parties keep Basque usage in words below 10% and thus clearly below citizens' average usage (PP, PSE-EE, EP and UPyD), the Basque party in government EAJ-PNV conveys two thirds of the words in Spanish and only EH Bildu maintains a balanced language use.

6.1.3. Language over time

In order to check language use over time, Figure 8 shows the language use at word level across the years. Results suggest that there is no considerable change in the use of the two languages, since they keep their distance along all the considered years. However, Basque word production between 2012 and 2020 decreases almost 10 points, although these two years present too few texts and this may affect their language use. If we consider the period from 2013 to 2019, we can observe a slight reduction in Basque use.




Figure 8: Language use over time at word level.

¹⁴We exclude EB, Ararteko and None/Unknown parties from Figures 6 and 7 due to their limited amount of data.

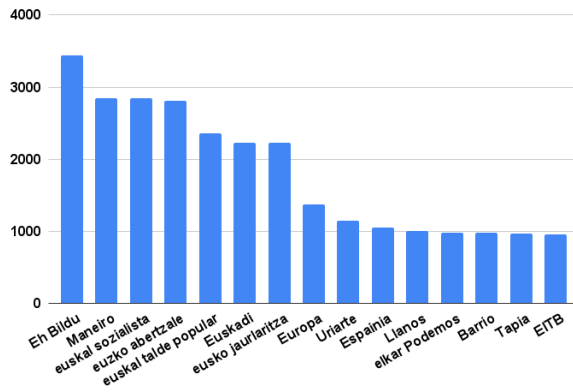


Figure 9: Most frequent Basque entities.




Figure 10: Most frequent Spanish entities.




Figure 11: Gender by party at speaker level.

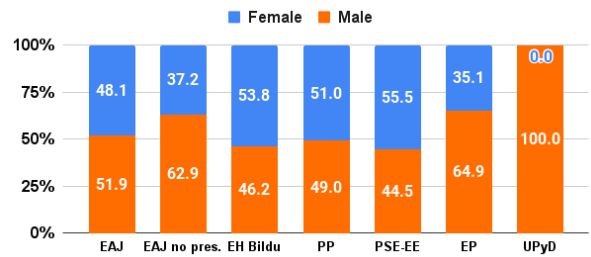


Figure 12: Gender by party at word level.

6.1.4. Entities by language

Lemma and NER allow us to extract lemma and entity frequencies, which can serve further purposes. We present in Figures 9 and 10 the most frequent entities in Basque and Spanish, respectively. As it can be seen, there are common entities of general use in parliament, like locations (e.g. *Euskadi*, *Espainia/España*, *Europa*), official institutions (e.g. *Eusko Jaurlarritza/Gobierno Vasco*, *EITB* for public broadcast service) or political groups (e.g. *EH Bildu*). However, those entities are mentioned in different frequencies: for example, *Euskadi* and *Gobierno Vasco* mentions double the frequency of the next entities in Spanish, whereas their Basque usage is more similar to the rest of the entities. On the other hand, Basque texts present many speaker names (e.g. *Maneiro*, *Uriarte*, *Llanos*) and Spanish speeches add entities referred to other topics (e.g. *ETA*, *Ertzaintza*, *Lanbide*). It must also be noted the difference in absolute frequency, being the mostly used entity in Spanish almost 4 times more frequent than the mostly used Basque entity. These results support the fact that Basque is generally used to start and end speeches and to address other speakers, while Spanish provides most of the speech content.

6.2. Gender

Figure 13 reports percentages regarding the number of speeches or words produced by men and women. Overall, women produce most of the speeches, but less words than men. However, if we filter out the presi-

dent's speeches, females drop drastically in speeches and stay below male percentages at both levels. In fact, the gap between women and men reaches more than 13 percentage points regarding speeches and words, which is three times the gap between the number of female (48%) and male speakers (52%) in parliament. These data would suggest that not only women speak less often, but also that they produce shorter speeches.




Figure 13: Gender at speech and word level.

6.2.1. Gender by party

Figures 11 and 12 show speaker and word gender percentages by party for comparison. While speaker and word level correlate in general, it is observed that two parties have more female presence at word level than the expected by the speaker rate (PSE-EE and PP), reaching almost 10 percentage points more. The rest show more male presence than perhaps expected, rising up to 10 points in the case of EAJ-PNV (without

the president). In the case of UPyD, the only speaker is a man. These results indicate that, in addition to a slightly lower female representation in parliament, the rates of women interventions at speaker and word levels remain substantially lower than those of men. The only exception is EH Bildu, for which female presence is a bit higher than that of their male colleagues.

6.2.2. Gender over time

Finally, Figure 14 illustrate the presence of female representatives in terms of word production over time. Although in 2012 women provided less than one third of the words, by 2020 they produced almost two thirds of them. If we ignore these two years (they gather very few texts compared to the rest), there is a clear trend indicating that more female politicians are speaking more often and more at length over the passing years.




Figure 14: Gender over time at word level.

7. Concluding Remarks

In this paper we present BasqueParl, a publicly available bilingual corpus for political discourse analysis containing Basque and Spanish transcriptions from the Basque Parliament during two legislative terms (2012-2016 and 2016-2020). The code-switching that characterizes most of the speeches offers an interesting opportunity to study language use in political debates. The transcriptions have been processed to enrich it with metadata such as date, speaker, year of birth, gender and party. In addition, lemmas and named entities have been automatically annotated for further analysis.

The corpus data reflects relevant information about the speakers' parliamentary activity. Regarding language use, Basque is often used but barely conveys speech content, one party (EH Bildu) being the exception. If we look at gender, women participate less in parliamentary debate and, overall, their speech content is smaller than we could expect from female representation in parliament. However, this trend is being reversed in the last years.

As far as we know, BasqueParl is the only large resource (around 14M words) of its kind for Spanish and Basque. We hope that its public availability will facilitate multilingual and crosslingual research on NLP tasks related to argumentation, discourse structure, sentiment analysis and fact-checking.

8. Acknowledgements

This work has been partially funded by the UPV/EHU Colab 19/19 project "Tools for the analysis of parliamentary discourses: polarization, subjectivity and affectivity in the post-truth era". Nayla Escribano is funded by the Basque Government grant "Programa Predoctoral de Formación de Personal Investigador No Doctor del Departamento de Educación del Gobierno Vasco". Rodrigo Agerri acknowledges the support received from the RYC-2017-23647 fellowship and from the ANTIDOTE - EU CHIST-ERA project (PCI2020-120717-2) of the Agencia Estatal de Investigación through the INT-Acciones de Programación Conjunta Internacional (MINECO) 2020 call.

9. Bibliographical References

- Abercrombie, G. and Batista-Navarro, R. (2020). ParIVote: A corpus for sentiment analysis of political debates. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5073–5078, Marseille, France, May. European Language Resources Association.
- Agerri, R., Vicente, I. S., Campos, J. A., Barrena, A., Saralegi, X., Etxabe, A. S., and Agirre, E. (2020). Give your text representation models some love: the case for basque. In *LREC*.
- Agerri, R. (2020). Projecting heterogeneous annotations for named entity recognition. In *IberLEF@SEPLN*.
- Akbik, A., Blythe, D., and Vollgraf, R. (2018). Contextual string embeddings for sequence labeling. In *Proceedings of the 27th international conference on computational linguistics*, pages 1638–1649.
- Akbik, A., Bergmann, T., Blythe, D., Rasul, K., Schweter, S., and Vollgraf, R. (2019). Flair: An easy-to-use framework for state-of-the-art nlp. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 54–59.
- Blei, D. M., Ng, A., and Jordan, M. I. (2003). Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022.
- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., and Stoyanov, V. (2020). Unsupervised cross-lingual representation learning at scale. In *ACL*.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Erjavec, T., Ogrodniczuk, M., Osenova, P., Ljubešić, N., Simov, K., Grigorova, V., Rudolf, M., Pančur, A., Kopp, M., Barkarson, S., Steingrímsson, S., van der Pol, H., Depoorter, G., de Does, J., Jongejan, B., Haltrup Hansen, D., Navarretta, C., Calzada Pérez, M., de Macedo, L. D., van Heusden, R., Marx, M., Çöltekin, Ç., Coole, M., Agnoloni, T., Frontini, F.,

- Montemagni, S., Quochi, V., Venturi, G., Ruisi, M., Marchetti, C., Battistoni, R., Sebők, M., Ring, O., Dargis, R., Utko, A., Petkevičius, M., Briedienė, M., Krilavičius, T., Morkevičius, V., Diwersy, S., Luxardo, G., and Rayson, P. (2021). Multilingual comparable corpora of parliamentary debates ParlaMint 2.1. Slovenian language resource repository CLARIN.SI.
- Hajlaoui, N., Kolovratnik, D., Väyrynen, J., Steinberger, R., and Varga, D. (2014). DCEP -digital corpus of the European parliament. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, Reykjavik, Iceland, May. European Language Resources Association (ELRA).
- Koehn, P. (2005). Europarl: A parallel corpus for statistical machine translation. In *Proceedings of Machine Translation Summit X: Papers*, pages 79–86, Phuket, Thailand, September 13-15.
- Marx, M. and Schuth, A. (2010). DutchParl. the parliamentary documents in Dutch. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta, May. European Language Resources Association (ELRA).
- Roukos, S., Graff, D., and Melamed, D. (1995). Hansard French/English LDC95T20.
- Thomas, M., Pang, B., and Lee, L. (2006). Get out the vote: Determining support or opposition from congressional floor-debate transcripts. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pages 327–335, Sydney, Australia, July. Association for Computational Linguistics.
- Zamorano, J. P. and Anke, L. E. (2020). Overview of capitel shared tasks at iberlef 2020: Named entity recognition and universal dependencies parsing. In *IberLEF@SEPLN*.