

# Detecting Unassimilated Borrowings in Spanish: An Annotated Corpus and Approaches to Modeling

Elena Álvarez-Mellado

NLP & IR group  
School of Computer Science, UNED  
elena.alvarez@lsi.uned.es

Constantine Lignos

Michtom School of Computer Science  
Brandeis University  
lignos@brandeis.edu

## Abstract

This work presents a new resource for borrowing identification and analyzes the performance and errors of several models on this task. We introduce a new annotated corpus of Spanish newswire rich in unassimilated lexical borrowings—words from one language that are introduced into another without orthographic adaptation—and use it to evaluate how several sequence labeling models (CRF, BiLSTM-CRF, and Transformer-based models) perform. The corpus contains 370,000 tokens and is larger, more borrowing-dense, OOV-rich, and topic-varied than previous corpora available for this task. Our results show that a BiLSTM-CRF model fed with subword embeddings along with either Transformer-based embeddings pretrained on codeswitched data or a combination of contextualized word embeddings outperforms results obtained by a multilingual BERT-based model.

## 1 Introduction and related work

Lexical borrowing is the process of bringing words from one language into another (Haugen, 1950). Borrowings are a common source of out-of-vocabulary (OOV) words, and the task of detecting borrowings has proven to be useful both for lexicographic purposes and for NLP downstream tasks such as parsing (Alex, 2008a), text-to-speech synthesis (Leidig et al., 2014), and machine translation (Tsvetkov and Dyer, 2016).

Recent work has approached the problem of extracting lexical borrowings in European languages such as German (Alex, 2008b; Garley and Hockenmaier, 2012; Leidig et al., 2014), Italian (Furiassi and Hofland, 2007), French (Alex, 2008a; Chesley, 2010), Finnish (Mansikkaniemi and Kurimo, 2012), Norwegian (Andersen, 2012; Losnegaard and Lyse, 2012), and Spanish (Serigos, 2017), with a particular focus on English lexical borrowings (often called *anglicisms*).

Computational approaches to mixed-language data have traditionally framed the task of identifying the language of a word as a tagging problem, where every word in the sequence receives a language tag (Lignos and Marcus, 2013; Molina et al., 2016; Solorio et al., 2014). As lexical borrowings can be single (e.g. *app*, *online*, *smartphone*) or multi-token (e.g. *machine learning*), they are a natural fit for chunking-style approaches. Álvarez Mellado (2020b) introduced chunking-based models for borrowing detection in Spanish media which were later improved (Álvarez Mellado, 2020a), producing an F1 score of 86.41.

However, both the dataset and modeling approach used by Álvarez Mellado (2020a) had significant limitations. The dataset focused exclusively on a single source of news and consisted only of headlines. The number and variety of borrowings were limited, and there was a significant overlap in borrowings between the training set and the test set, which prevented assessment of whether the modeling approach was actually capable of generalizing to previously unseen borrowings. Additionally, the best results were obtained by a CRF model, and more sophisticated approaches were not explored.

The contributions of this paper are a new corpus of Spanish annotated with unassimilated lexical borrowings and a detailed analysis of the performance of several sequence-labeling models trained on this corpus. The models include a CRF, Transformer-based models, and a BiLSTM-CRF with different word and subword embeddings (including contextualized embeddings, BPE embeddings, character embeddings, and embeddings pretrained on codeswitched data). The corpus contains 370,000 tokens and is larger and more topic-varied than previous resources. The test set was designed to be as difficult as possible; it covers sources and dates not seen in the training set, includes a high number of OOV words (92% of the borrowings in the test set are OOV) and is very borrowing-dense

Media	Topics	Set(s)
ElDiario.es	General newspaper	Train, Dev.
El orden mundial	Politics	Test
Cuarto poder	Politics	Test
El salto	Politics	Test
La Marea	Politics	Test
Píkara Magazine	Feminism	Test
El blog salmón	Economy	Test
Pop rosa	Gossip	Test
Vida extra	Videogames	Test
Espinof	Cinema & TV	Test
Xataka	Technology	Test
Xataka Ciencia	Technology	Test
Xataka Android	Technology	Test
Genbeta	Technology	Test
Microsiervos	Technology	Test
Agencia Sinc	Science	Test
Diario del viajero	Travel	Test
Bebe y más	Parenthood	Test
Vitónica	Lifestyle & sports	Test
Foro atletismo	Sports	Test
Motor pasión	Automobiles	Test

Table 1: Sources included in each dataset split (URLs provided in Appendix A)

(20 borrowings per 1,000 tokens).

The dataset we present is publicly available<sup>1</sup> and has been released under a CC BY-NC-SA-4.0 license. This dataset was used for the ADoBo shared task on automatic detection of borrowings at IberLEF 2021 (Álvarez Mellado et al., 2021).

## 2 Data collection and annotation

### 2.1 Contrasting lexical borrowing with codeswitching

Linguistic borrowing can be defined as the transference of linguistic elements between two languages. Borrowing and codeswitching have been described as a continuum (Clyne et al., 2003).

Lexical borrowing involves the incorporation of single lexical units from one language into another language and is usually accompanied by morphological and phonological modification to conform with the patterns of the recipient language (Onysko, 2007; Poplack et al., 1988).

On the other hand, codeswitches are by definition not integrated into a recipient language, unlike established loanwords (Poplack, 2012). While codeswitches require a substantial level of fluency, comply with grammatical restrictions in both languages, and are produced by bilingual speakers in bilingual discourses, lexical borrowings are words used by monolingual individuals that eventually

<sup>1</sup><https://github.com/lirondos/coalas>

Set	Tokens	ENG	OTHER	Unique
Training	231,126	1,493	28	380
Development	82,578	306	49	316
Test	58,997	1,239	46	987
Total	372,701	3,038	123	1,683

Table 2: Corpus splits with counts

become lexicalized and assimilated as part of the recipient language lexicon until the knowledge of “foreign” disappears (Lipski, 2005).

### 2.2 Data selection

Our dataset consists of Spanish newswire annotated for unassimilated lexical borrowings. All of the sources used are European Spanish online publications (newspapers, blogs, and news sites) published in Spain and written in European Spanish.

Data was collected separately for the training, development, and test sets to ensure minimal overlap in borrowings, topics, and time periods. The training set consists of a collection of articles appearing between August and December 2020 in *elDiario.es*, a progressive online newspaper based in Spain. The development set contains sentences in articles from January 2021 from the same source.

The data in the test set consisted of annotated sentences extracted in February and March 2021 from a diverse collection of online Spanish media that covers specialized topics rich in lexical borrowings and usually not covered by *elDiario.es*, such as sports, gossip or videogames (see Table 1).

To focus annotation efforts for the training set on articles likely to contain unassimilated borrowings, the articles to be annotated were selected by first using a baseline model and were then human-annotated. To detect potential borrowings, the CRF model and data from Álvarez Mellado (2020b) was used along with a dictionary look-up pipeline. Articles that contained more than 5 borrowing candidates were selected for annotation.

The main goal of data selection for the development and test sets was to create borrowing-dense, OOV-rich datasets, allowing for better assessment of generalization. To that end, the annotation was based on sentences instead of full articles. If a sentence contained a word either flagged as a borrowing by the CRF model, contained in a wordlist of English, or simply not present in the training set, it was selected for annotation. This data selection approach ensured a high number of borrowings and

OOV words, both borrowings and non-borrowings. While the training set contains 6 borrowings per 1,000 tokens, the test set contains 20 borrowings per 1,000 tokens. Additionally, 90% of the unique borrowings in the development set were OOV (not present in training). 92% of the borrowings in the test set did not appear in training (see Table 2).

### 2.3 Annotation process

The corpus was annotated with BIO encoding using Doccano (Nakayama et al., 2018) by a native speaker of Spanish with a background in linguistic annotation (see Appendix C). The annotation guidelines (provided in Appendix B) were based on those of Álvarez Mellado (2020a) but were expanded to account for a wider diversity of topics. Following Serigos’s observations and Álvarez Mellado’s work, English lexical borrowings were labeled ENG, other borrowings were labeled OTHER. Here is an example from the training set:<sup>2</sup>

**Benching** [ENG], estar en el banquillo de tu **crush** [ENG] mientras otro juega de titular.

In order to assess the quality of the guidelines and the annotation, a sample of 9,110 tokens from 450 sentences (60% from the test set, 20% from training, 20% from development) was divided among a group of 9 linguists for double annotation. The mean inter-annotation agreement computed by Cohen’s kappa was 0.91, which is above the 0.8 threshold of reliable annotation (Artstein and Poesio, 2008).

### 2.4 Limitations

Like all resources, this resource has significant limitations that its users should be aware of. The corpus consists exclusively of news published in Spain and written in European Spanish. This fact by no means implies the assumption that European Spanish represents the whole of the Spanish language.

The notion of assimilation is usage-based and community-dependant, and thus the dataset we present and the annotation guidelines that were followed were designed to capture a very specific phenomena at a given time and in a given place: unassimilated borrowings in the Spanish press.

In order to establish whether a given word has been assimilated or not, the annotation guidelines rely on lexicographic sources such as the prescriptivist *Diccionario de la Lengua Española* (Real Academia Española, 2020) by the Royal

<sup>2</sup>“Benching: being on your crush’s bench while someone else plays in the starting lineup.”

Set	Precision	Recall	F1
Development			
ALL	74.13	59.72	66.15
ENG	74.20	68.63	71.31
OTHER	66.67	4.08	7.69
Test			
ALL	77.89	43.04	55.44
ENG	78.09	44.31	56.54
OTHER	57.14	8.70	15.09

Table 3: CRF performance on the development and test sets (results from a single run)

Spanish Academy, a dictionary that aims to cover world-wide Spanish but whose Spain-centric criteria has been previously pointed out (Blanch, 1995; Fernández Gordillo, 2014). In addition, prior work has suggested that Spanish from Spain may have a higher tendency of anglicism-usage than other Spanish dialects (McClelland, 2021). Consequently, we limit the scope of the dataset to European Spanish not because we consider that this variety represents the whole of the Spanish-speaking community, but because we consider that the approach we have taken here may not account adequately for the whole diversity in borrowing assimilation within the Spanish-speaking world.

### 2.5 Licensing

Our annotation is licensed with a permissive CC BY-NC-SA-4.0 license. With one exception, the sources included in our dataset release their content under Creative Commons licenses that allow for reusing and redistributing the material for non commercial purposes. This was a major point when deciding which news sites would be included in the dataset. The exception is the source *Microsiervos*, whose content we use with explicit permission from the copyright holder. Our annotation is “stand-off” annotation that does not create a derivative work under Creative Commons licenses, so ND licenses are not a problem for our resource. Table 9 in Appendix A lists the URL and license type for each source.

## 3 Modeling

The corpus was used to evaluate four types of models for borrowing extraction: (1) a CRF model, (2) two Transformer-based models, (3) a BiLSTM-CRF model with different types of unadapted embeddings (word, BPE, and character embeddings) and (4) a BiLSTM-CRF model with previously fine-

	Development			Test		
	Precision	Recall	F1	Precision	Recall	F1
<b>BETO</b>						
ALL	73.36 ± 3.6	73.46 ± 1.6	73.35 ± 1.5	86.76 ± 1.3	75.50 ± 2.8	80.71 ± 1.5
ENG	74.30 ± 1.8	84.05 ± 1.6	78.81 ± 1.6	87.33 ± 2.9	77.99 ± 1.5	82.36 ± 1.5
OTHER	47.24 ± 22.2	7.34 ± 3.7	11.93 ± 4.9	36.12 ± 10.1	8.48 ± 3.5	13.23 ± 3.6
<b>mBERT</b>						
ALL	<b>79.96 ± 1.9</b>	<b>73.86 ± 2.7</b>	<b>76.76 ± 2.0</b>	<b>88.89 ± 1.0</b>	<b>76.16 ± 1.6</b>	<b>82.02 ± 1.0</b>
ENG	80.25 ± 2.6	84.31 ± 1.9	82.21 ± 1.9	89.25 ± 1.6	78.85 ± 1.0	83.64 ± 1.0
OTHER	66.18 ± 16.5	8.6 ± 6.8	14.41 ± 10.0	45.30 ± 11.3	7.61 ± 1.5	12.84 ± 2.3

Table 4: Mean and standard deviation of scores on the development and test sets for BETO and mBERT

tuned Transformer-based embeddings pretrained on codeswitched data. By *unadapted* embeddings, we mean embeddings that have not been fine-tuned for the task of anglicism detection or a related task (e.g. codeswitching).

Evaluation for all models required extracted spans to match the annotation exactly in span and type to be correct. Evaluation was performed with SeqScore (Palen-Michel et al., 2021), using conllev-style repair for invalid label sequences. All models were trained using an AMD 2990WX CPU and a single RTX 2080 Ti GPU.

### 3.1 Conditional random field model

As baseline model, we evaluated a CRF model with handcrafted features from Álvarez Mellado (2020b). The model was built using `pycrfsuite` (Peng and Korobov, 2014), a Python wrapper for `crfsuite` (Okazaki, 2007) that implements CRF for labeling sequential data. The model also uses the `Token` and `Span` utilities from `spaCy` library (Honnibal and Montani, 2017). The following handcrafted binary features from Álvarez Mellado (2020b) were used for the model:

- Bias: active on all tokens to set per-class bias
- Token: the string of the token
- Uppercase: active if the token is all uppercase
- Titlecase: active if only the first character of the token is capitalized
- Character trigram: an active feature for every trigram contained in the token
- Quotation: active if the token is any type of quotation mark (‘ ’ " “ ” « »)
- Suffix: last three characters of the token
- POS tag: part-of-speech tag of the token provided by `spaCy` utilities
- Word shape: shape representation of the token provided by `spaCy` utilities
- Word embedding: provided by Spanish word2vec 300 dimensional embeddings by Cardellino (2019),

one feature per dimension

- URL: active if the token could be validated as a URL according to `spaCy` utilities
- Email: active if the token could be validated as an email address by `spaCy` utilities
- Twitter: active if the token could be validated as a possible Twitter special token: `#hashtag` or `@username`

A window of two tokens in each direction was used for feature extraction. Optimization was performed using L-BFGS, with the following hyperparameter values chosen following the best results from Álvarez Mellado (2020b) were set: `c1 = 0.05`, `c2 = 0.01`. As shown in Table 3, the CRF produced an overall F1 score of 66.15 on the development set (P: 74.13, R: 59.72) and an overall F1 of 55.44 (P: 77.89, R: 43.04) on the test set. The CRF results on our dataset are far below the F1 of 86.41 reported by Álvarez Mellado (2020b), showing the impact that a topically-diverse, OOV-rich dataset can have, especially on test set recall. These results demonstrate that we have created a more difficult task and motivate using more sophisticated models.

### 3.2 Transformer-based models

We evaluated two Transformer-based models:

- BETO base cased model: a monolingual BERT model trained for Spanish (Cañete et al., 2020)
- mBERT: multilingual BERT, trained on Wikipedia in 104 languages (Devlin et al., 2019)

Both models were run using the `Transformers` library by HuggingFace (Wolf et al., 2020). The same default hyperparameters were used for both models: 3 epochs, batch size 16, and maximum sequence length 256. Except where otherwise specified, we report results for 10 runs that use different random seeds for initialization. We perform statistical significance testing using the Wilcoxon rank-sum test.

As shown in Table 4, the mBERT model

(F1: 82.02) performed better than BETO (F1: 80.71), and the difference was statistically significant ( $p = 0.027$ ). Both models performed better on the test set than on the development set, despite the difference in topics between them, suggesting good generalization. This is a remarkable difference with the CRF results, where the CRF performed substantially worse on the test set than on the development set. The limited number of OTHER examples explains the high deviations in the results for this label.

### 3.3 BiLSTM-CRF

We explored several possibilities for a BiLSTM-CRF model fed with different types of word and subword embeddings. The purpose was to assess whether the combination of different embeddings that encode different linguistic information could outperform the Transformer-based models in Section 3.2. All of our BiLSTM-CRF models were built using `Flair` (Akbik et al., 2018) with default hyperparameters (hidden size = 256, learning rate = 0.1, mini batch size = 32, max number of epochs = 150) and embeddings provided by `Flair`.

#### 3.3.1 Preliminary embedding experiments

We first ran exploratory experiments on the development set with different types of embeddings using `Flair` tuning functionalities. We explored the following embeddings: Transformer embeddings (mBERT and BETO), fastText embeddings (Bojanowski et al., 2017), one-hot embeddings, byte pair embeddings (Heinzerling and Strube, 2018), and character embeddings (Lample et al., 2016).

The best results were obtained by a combination of mBERT embeddings and character embeddings (F1: 74.00), followed by a combination of BETO embeddings and character embeddings (F1: 72.09). These results show that using contextualized embeddings unsurprisingly outperforms non-contextualized embeddings for this task, and that subword representation is important for the task of extracting borrowings that have not been adapted orthographically. The finding regarding the importance of subwords is consistent with previous work; feature ablation experiments for borrowing detection have shown that character trigram features contributed the most to the results obtained by a CRF model (Álvarez Mellado, 2020b).

The worst result (F1: 39.21) was produced by a model fed with one-hot vectors, and the second-worst result was produced by a model fed exclu-

sively with character embeddings. While it performed poorly (F1: 41.65), this fully unlexicalized model outperformed one-hot embeddings, reinforcing the importance of subword information for the task of unassimilated borrowing extraction.

#### 3.3.2 Optimally combining embeddings

In light of the preliminary embedding experiments and our earlier experiments with Transformer-based models, we fed our BiLSTM-CRF model with different combinations of contextualized word embeddings (including English BERT embeddings from Devlin et al.), byte-pair embeddings and character embeddings. Table 5 shows development set results from different combinations of embeddings. The best overall F1 on the development set was obtained by the combination of BETO embeddings, BERT embeddings and byte-pair embeddings. The model fed with BETO embeddings, BERT embeddings, byte-pair embeddings and character embeddings ranked second.

Several things stand out from the results in Table 5. The BETO+BERT embedding combination consistently works better than mBERT embeddings, and BPE embeddings contribute to better results. Character embeddings, however, seem to produce little effect at first glance. Given the same model, adding character embeddings produced little changes in F1 or even slightly hurt the results. Although character embeddings seem to make little difference in overall F1, recall was consistently higher in models that included character embeddings, and in fact, the model with BETO+BERT embeddings, BPE embeddings and character embeddings produced the highest recall overall (77.46). This is an interesting finding, as our results from Sections 3.1 and 3.2 as well as prior work (Álvarez Mellado, 2020b) identified recall as weak for borrowing detection models.

The two best-performing models from Table 5 (BETO+BERT embeddings, BPE embeddings and optionally character embeddings) were evaluated on the test set. Table 6 gives results per type on the development and test sets for these two models. For both models, results on the test set were better (F1: 82.92, F1: 83.63) than on the development set (F1: 81.21, F1: 81.05). Although the best F1 score on the development set was obtained with no character embeddings, when run on the test set the model with character embeddings obtained the best score; however, these differences did not show to be statistically significant. Recall, on the other hand,

Word embedding	BPE embedding	Char embedding	Precision	Recall	F1
mBERT	-	-	82.27	69.30	75.23
mBERT	-	✓	79.45	72.96	76.06
mBERT	multi	✓	81.37	73.80	77.40
mBERT	es, en	-	83.07	74.65	78.64
mBERT	es, en	✓	80.83	77.18	78.96
BETO, BERT	-	✓	81.44	76.62	78.96
BETO, BERT	-	-	81.87	76.34	79.01
BETO, BERT	es, en	✓	85.94	<b>77.46</b>	81.48
BETO, BERT	es, en	-	<b>86.98</b>	77.18	<b>81.79</b>

Table 5: Scores (ALL labels) on the development set using a BiLSTM-CRF model with different combinations of word and subword embeddings (results from a single random seed)

Embeddings	Development			Test		
	Precision	Recall	F1	Precision	Recall	F1
BETO+BERT and BPE						
ALL	<b>85.84</b> ± 1.2	77.07 ± 0.8	<b>81.21</b> ± 0.5	<b>90.00</b> ± 0.8	76.89 ± 1.8	82.92 ± 1.1
ENG	86.15 ± 1.1	88.00 ± 0.6	87.05 ± 0.6	90.20 ± 1.9	79.36 ± 1.1	84.42 ± 1.1
OTHER	72.81 ± 13.7	8.8 ± 0.9	15.60 ± 1.6	62.68 ± 8.0	10.43 ± 0.9	17.83 ± 1.3
BETO+BERT, BPE, and char						
ALL	84.29 ± 1.0	<b>78.06</b> ± 0.9	81.05 ± 0.5	89.71 ± 0.9	<b>78.34</b> ± 1.1	<b>83.63</b> ± 0.7
ENG	84.54 ± 1.0	89.05 ± 0.5	86.73 ± 0.5	89.90 ± 1.1	80.88 ± 0.7	85.14 ± 0.7
OTHER	73.50 ± 11.4	9.38 ± 3.0	16.44 ± 4.8	61.14 ± 7.9	9.78 ± 1.8	16.81 ± 2.9

Table 6: Mean and standard deviation of scores on the development and test sets using a BiLSTM-CRF model with BETO and BERT embeddings, BPE embeddings, and optionally character embeddings

was higher when run with character embeddings (R: 78.34) than when run without them (R: 76.89), and the difference was statistically significant ( $p = 0.019$ ). This finding again corroborates the positive impact that character information can have in recall when dealing with previously unseen borrowings.

### 3.4 Transfer learning from codeswitching

Finally, we decided to explore whether detecting unassimilated lexical borrowings could be framed as transfer learning from language identification in codeswitching. As before, we ran a BiLSTM-CRF model using `Flair`, but instead of using the unadapted Transformer embeddings, we used codeswitch embeddings (Sarker, 2020), fine-tuned Transformer-based embeddings pretrained for language identification on the Spanish-English section of the LinCE dataset (Aguilar et al., 2020).

Table 7 gives results for these models. The two best-performing models were the BiLSTM-CRF with codeswitch and BPE embeddings (F1: 84.06) and the BiLSTM-CRF model with codeswitch, BPE and character embeddings (F1: 84.22). The differences between these two models did not show to be statistically significant, but the difference with the best-performing model with unadapted embeddings from Section 3.3 (F1: 83.63) was sta-

tistically significant ( $p = 0.018$ ). These two best-performing models however obtained worse results on the development set than those obtained by the best-performing models from Section 3.3.

Adding BPE embeddings showed to improve F1 score by around 1 point compared to either feeding the model with only codeswitch (F1: 82.83) or only codeswitch and character embeddings (F1: 83.13), and the differences were statistically significant in both cases ( $p = 0.024$ ,  $p = 0.018$ ).

It should be noted that this transfer learning approach is indirectly using more data than just the training data from our initial corpus, as the codeswitch-based BiLSTM-CRF models benefit from the labeled data seen during pretraining for the language-identification task.

## 4 Error analysis

We compared the different results produced by the best performing model of each type on the test set: (1) the mBERT model, (2) the BiLSTM-CRF with BERT+BETO, BPE and character embeddings and (3) the BiLSTM-CRF model with codeswitch, BPE and character embeddings. We divide the error analysis into two sections. We first analyze errors that were made by all three models, with the aim of discovering which instances of the dataset were

Embeddings	Development			Test		
	Precision	Recall	F1	Precision	Recall	F1
Codeswitch						
ALL	80.21 ± 1.7	74.42 ± 1.5	77.18 ± 0.5	90.05 ± 1.0	76.76 ± 3.0	82.83 ± 1.6
ENG	80.19 ± 1.6	85.59 ± 0.5	82.78 ± 0.5	90.05 ± 3.1	79.37 ± 1.6	84.33 ± 1.6
OTHER	85.83 ± 19.2	4.70 ± 2.3	8.78 ± 4.2	90.00 ± 12.9	6.52 ± 0.0	12.14 ± 0.1
Codeswitch + char						
ALL	81.02 ± 2.5	74.56 ± 1.5	77.62 ± 0.9	89.92 ± 0.7	77.34 ± 2.2	83.13 ± 1.2
ENG	81.00 ± 1.6	85.91 ± 0.9	83.34 ± 0.9	89.95 ± 2.3	80.00 ± 1.2	84.67 ± 1.2
OTHER	73.00 ± 41.6	3.67 ± 2.7	6.91 ± 4.9	68.50 ± 40.5	5.43 ± 2.9	9.97 ± 5.3
Codeswitch + BPE						
ALL	<b>83.62</b> ± 1.6	<b>75.91</b> ± 0.7	<b>79.57</b> ± 0.7	90.43 ± 0.7	78.55 ± 1.3	84.06 ± 0.7
ENG	83.54 ± 0.7	86.86 ± 0.8	85.16 ± 0.8	90.57 ± 1.3	81.14 ± 0.7	85.59 ± 0.7
OTHER	94.28 ± 12.0	7.55 ± 2.1	13.84 ± 3.6	67.17 ± 8.3	8.70 ± 1.4	15.30 ± 2.2
Codeswitch + BPE + char						
ALL	82.88 ± 1.8	75.70 ± 1.3	79.10 ± 0.9	<b>90.60</b> ± 0.7	<b>78.72</b> ± 2.5	<b>84.22</b> ± 1.6
ENG	82.90 ± 1.4	86.57 ± 1.0	84.66 ± 1.0	90.76 ± 2.6	81.32 ± 1.6	85.76 ± 1.6
OTHER	87.23 ± 14.5	7.75 ± 2.8	14.03 ± 4.7	66.50 ± 17.5	8.70 ± 2.3	15.13 ± 3.5

Table 7: Mean and standard deviation of scores on the development and test sets for a BiLSTM-CRF model with combinations of codeswitch embeddings, BPE embeddings, and character embeddings

challenging for all models. We then analyze unique answers (both correct and incorrect) per model, with the aim of gaining insight on what are the unique characteristics of each model in comparison with other models.

#### 4.1 Errors made by all models

##### 4.1.1 Borrowings labeled as O

There were 137 tokens in the test set that were incorrectly labeled as O by all three models. 103 of these were of type ENG, 34 were of type OTHER. These errors can be classified as follows

– Borrowings in upper case (12), which tend to be mistaken by models with proper nouns:

Análisis de empresa basados en **Big Data** [ENG].<sup>3</sup>

– Borrowings in sentence-initial position (9), which were titlecased and therefore consistently mislabeled as O:

**Youtuber** [ENG], mujer y afroamericana: Candace Owen podría ser la alternativa a Trump.<sup>4</sup>

Sentence-initial borrowings are particularly tricky, as models tend to confuse these with foreign named entities. In fact, prior work on anglicism detection based on dictionary lookup (Serigos, 2017) stated that borrowings in sentence-initial position were rare in Spanish and consequently chose to ignore all foreign words in sentence-initial position under the assumption that they could be considered named entities. However, these examples (and the

difficulty they pose for models) prove that sentence-initial borrowings are not rare and therefore should not be overlooked.

– Borrowings that also happen to be words in Spanish (8), such as the word *primer*, that is a borrowing found in makeup articles (*un primer hidratante*, “a hydrating primer”) but also happens to be a fully Spanish adjective meaning “first” (*primer premio*, “first prize”). Borrowings like these are still treated as fully unassimilated borrowings by speakers, even when the form is exactly the same as an already-existing Spanish word and were a common source of mislabeling, especially partial mismatches in multitoken borrowings: *red* (which exists in Spanish meaning “net”) in *red carpet*, *tractor* in *tractor pulling* or *total* in *total look*.

– Borrowings that could pass as Spanish words (58): most of the mislabeled borrowings were words that do not exist in Spanish but that could orthographically pass for a Spanish word. That is the case of words like *burpees* (hypothetically, a conjugated form of the non-existing verb *burpear*), *gimbal*, *mules*, *bromance* or *nude*.

– Other borrowings (50): a high number of mislabeled borrowings were borrowings that were orthographically implausible in Spanish, such as *trenchs*, *multipads*, *hypes*, *riff*, *scrunchie* or *mint*. The fact that none of our models were able to correctly classify these orthographically implausible examples leaves the door open to further exploration of character-based models and investigating character-level perplexity as a source of information.

<sup>3</sup>“Business analytics based on Big Data”

<sup>4</sup>“Youtuber, woman and African-American: Candace Owen could be the alternative to Trump”

#### 4.1.2 Non-borrowings labeled as borrowings

29 tokens were incorrectly labeled as borrowings by all three models. These errors can be classified in the following groups:

- Metalinguistic usage and reported speech: a foreign word or sentence that appears in the text to refer to something someone said or wrote.

Escribir “**icon pack**” [ENG] en el buscador.<sup>5</sup>

- Lower-cased proper nouns: such as websites.

Acceder a la página **flywithkarolg** [ENG]<sup>6</sup>

- Computer commands: the test set included blog posts about technology, which mentioned computer commands (such as *sudo apt-get update*) that were consistently mistaken by our models as borrowings. These may seem like an extreme case—after all, computer commands do contain English words—but they are a good example of the real data that a borrowing-detection system may encounter.

- Foreign words within proper nouns: lower-cased foreign words that were part of multitoken proper nouns.

La serie “10.000 **ships** [ENG]” cuenta la odisea de la princesa Nyermeria.<sup>7</sup>

- Acronyms and acronym expansions:

El entrenamiento HITT (**high intensity interval training** [ENG])<sup>8</sup>

- Assimilated borrowings: certain borrowings that are already considered by RAE’s dictionary as fully assimilated were labeled by all models as anglicisms.

Labios rojos, a juego con el **top** [ENG].<sup>9</sup>

#### 4.1.3 Type confusion

Three tokens of type OTHER were marked by all models as ENG. There were no ENG borrowings that were labeled as OTHER by all three models.

Había **buffet** [ENG] libre.<sup>10</sup>

### 4.2 Unique answers per model

We now summarize the unique mistakes and correct answers made per model, with the aim of understanding what data points were handled uniquely well or badly by each model.

<sup>5</sup>“Type ‘icon pack’ on the search box”

<sup>6</sup>“Access the website flywithkarolg”

<sup>7</sup>“The series ‘10,000 ships’ tells the story of princess Nyermeria”

<sup>8</sup>“HITT training (High-intensity interval training)”

<sup>9</sup>“Red lips, matching top”

<sup>10</sup>“There was a free buffet”

#### 4.2.1 mBERT

There were 46 tokens that were incorrectly labeled as borrowings only by the mBERT model. These include foreign words used in reported speech or acronym expansion (21), proper names (11) and already assimilated borrowings (7).

There were 27 tokens that were correctly labeled only by the mBERT model. The mBERT model was particularly good at detecting the full span of multitoken borrowings as in *no knead bread*, *total white*, *wide leg* or *kettlebell swings* (which were only partially detected by other models) and at detecting borrowings that could pass for Spanish words (such as *fashionista*, *samples*, *vocoder*). In addition, the mBERT model also correctly labeled as ○ 12 tokens that the other two models mistook as borrowings, including morphologically adapted anglicisms, such as *craftear* (Spanish infinitive of the verb *to craft*), *crackear* (from *to crack*) or *lookazo* (augmentative of the noun *look*).

#### 4.2.2 BiLSTM-CRF with unadapted embeddings

There were 23 tokens that were incorrectly labeled as borrowings solely by this model, the most common types being assimilated borrowings (such as *fan*, *clon*) and Spanish words (*fiestones*) (9 each).

32 tokens were correctly labeled as borrowings only by this model. These include borrowings that could pass for Spanish words (*camel*, *canvas*). In addition, this model also correctly labeled as ○ 6 tokens that the other two mistook as borrowings, including old borrowings that are considered today as fully assimilated (such as *films* or *sake*) or the usage of *post* as a prefix of Latin origin (as in *post-producción*), which other models mistook with the English word *post*.

#### 4.2.3 BiLSTM-CRF with codeswitch embeddings

The codeswitch-based system incorrectly labeled 18 tokens as borrowings, including proper nouns (7), such as *Baby Spice*, and fully assimilated borrowings (5), such as *jersey*, *relax* or *tutorial*.

This model correctly labeled 27 tokens that were mistakenly ignored by other models, including multitoken borrowings (*dark and gritty*, *red carpet*) and other borrowings that were non-compliant with Spanish orthographic rules but that were however ignored by other models (*messy*, *athleisure*, *multi-touch*, *workaholic*).



Model	Word emb	BPE emb	Char emb	Development			Test		
				Precision	Recall	F1	Precision	Recall	F1
CRF	w2v (spa)	-	-	74.13	59.72	66.15	77.89	43.04	55.44
BETO	-	-	-	73.36	73.46	73.35	86.76	75.50	80.71
mBERT	-	-	-	79.96	73.86	76.76	88.89	76.16	82.02
BiLSTM-CRF	BETO+BERT	en, es	-	<b>85.84</b>	77.07	<b>81.21</b>	90.00	76.89	82.92
BiLSTM-CRF	BETO+BERT	en, es	✓	84.29	<b>78.06</b>	81.05	89.71	78.34	83.63
BiLSTM-CRF	Codeswitch	-	-	80.21	74.42	77.18	90.05	76.76	82.83
BiLSTM-CRF	Codeswitch	-	✓	81.02	74.56	77.62	89.92	77.34	83.13
BiLSTM-CRF	Codeswitch	en, es	-	83.62	75.91	79.57	90.43	78.55	84.06
BiLSTM-CRF	Codeswitch	en, es	✓	82.88	75.70	79.10	<b>90.60</b>	<b>78.72</b>	<b>84.22</b>

Table 8: Scores (ALL labels) for the development and test sets across all models. For the CRF, results from a single run are reported. For all other models, the score reported is the mean calculated over 10 runs with different random seeds (see Tables 4, 6, and 7 for standard deviations).

The codeswitch-based model also correctly labeled as  $\circ$  16 tokens that the other two models labeled as borrowings, including acronym expansions, lower-cased proper names and orthographically unorthodox Spanish words, such as the ideophone *tiki-taka* or *shavales* (a non-standard writing form of the word *chavales*, “guys”).

## 5 Discussion

Table 8 provides a summary of our results. As we have seen, the diversity of topics and the presence of OOV words in the dataset can have a remarkable impact on results. The CRF model—which in previous work had reported an F1 score of 86—saw its performance drop to a 55 when dealing with our dataset, despite the fact that both datasets consisted of journalistic European Spanish texts.

On the other hand, neural models (Transformer-based and BiLSTM-CRF) performed better. All of them performed better on the test set than on the development set, which shows good generalization ability. The BiLSTM-CRF model fed with different combinations of Transformer-based word embeddings and subword embeddings outperformed multilingual BERT and Spanish monolingual BETO. The model fed with codeswitch, BPE, and character embeddings ranked first and was significantly better than the result obtained by the model fed with BETO+BERT, BPE, and character embeddings.

Our error analysis shows that recall was a weak point for all models we examined. Among false negatives, upper-case borrowings (such as *Big Data*) and borrowings in sentence-initial position (in titlecase) were frequent, as they tend to be mistaken with named entities. This finding suggests that borrowings with capitalized initial should not be overlooked. Similarly, words that exist both

in English and Spanish (like *primer* or *red*) are not rare and were also a common source of error. Adding character embeddings produced a statistically significant improvement in recall, which opens the door to future work.

Concurrently with the work presented on this paper, De la Rosa (2021) explored using supplementary training on intermediate labeled-data tasks (such as POS, NER, codeswitching and language identification) along with multilingual Transformer-based models to the task of detecting borrowings. Alternatively, Jiang et al. (2021) used data augmentation to train a CRF model for the same task.

## 6 Conclusion

We have introduced a new corpus of Spanish newswire annotated with unassimilated lexical borrowings. The test set has a high number of OOV borrowings—92% of unique borrowings in the test set were not seen during training—and is more borrowing-dense and varied than resources previously available. We have used the dataset to explore several sequence labeling models (CRF, BiLSTM-CRF, and Transformer-based models) for the task of extracting lexical borrowings in a high-OOV setting. Results show that a BiLSTM-CRF model fed with Transformer-based embeddings pretrained on codeswitched data along subword embeddings produced the best results (F1: 84.22, 84.06), followed by a combination of contextualized word embeddings and subword embeddings (F1: 83.63). These models outperformed prior models for this task (CRF F1: 55.44) and multilingual Transformer-based models (mBERT F1: 82.02).

## 7 Ethical considerations

In this paper we have introduced an annotated dataset and models for detecting unassimilated borrowings in Spanish. The dataset is openly-licensed, and detailed annotation guidelines are provided (Appendix B). Appendix C includes a data statement that provides information regarding the curation rationale, annotator demographics, text characteristics, etc. of the dataset we have presented. We hope these resources will contribute to bringing more attention to borrowing extraction, a task that has been little explored in the field of NLP but that can be of great help to lexicographers and linguists studying language change.

However, the resources we have presented should not be considered a full depiction of either the process of borrowing or the Spanish language in general. We have identified four important considerations that any future systems that build off this research should be aware of.

**The process of borrowing.** Borrowing is a complex phenomenon that can manifest at all linguistic levels (phonological, morphological, lexical, syntactic, semantic, pragmatic). This work is exclusively concerned with lexical borrowings. Furthermore, in this work we have taken a synchronic approach to borrowing: we deal with borrowings that are considered as such in a given dialect and at a given point in time. The process of borrowing assimilation is a diachronic process, and the notion of what is perceived as unassimilated can vary across time and varieties. As a result, our dataset and models may not be suitable to account for partially assimilated borrowings or even for unassimilated borrowings in a different time period.

**Language variety.** The dataset we have presented is exclusively composed of European Spanish journalistic texts. In addition, the guidelines we have described were designed to capture a very specific phenomena: unassimilated borrowings in the Spanish press. In fact, the annotation guidelines rely on sources such as *Diccionario de la Lengua Española*, a lexicographic source whose Spain-centric criteria has been previously pointed out (Blanch, 1995; Fernández Gordillo, 2014). Consequently, the scope of our work is restricted to unassimilated borrowings in journalistic European Spanish. Our dataset and models may not translate adequately to other Spanish-speaking areas or genres.

**The preeminence of written language.** In our work, the notion of what a borrowing is is heavily influenced by how a word is written. According to our guidelines, a word like *meeting* will be considered unassimilated, while the Spanish form *mitin* will be considered assimilated. These preferences in writing may indirectly reveal how well-established a loanword is or how foreign it is perceived by the speaker. But it is questionable that these two forms necessarily represent a difference in pronunciation or linguistic status in the speaker's mental lexicon. How a word is written can be helpful for the purpose of detecting novel anglicisms in written text, but ideally one would not establish a definition of borrowing solely based on lexicographic, corpus-derived or orthotypographic cues. These are all valuable pieces of information, but they only represent an indirect evidence of the status that the word holds in the lexicon. After all, speakers will identify a word as an anglicism (and use it as such), regardless of whether the word is written in a text or used in speech.

On the other hand, the lexicographic fact that a word came from another language may not be enough as a criterion to establish the notion of borrowing. Speakers use words all the time without necessarily knowing where they came from or how long ago they were incorporated into the language. The origin of the word may just be a piece of trivia that is totally irrelevant or unknown to the speaker at the time of speaking, so the etymological origin of the word might not be enough to account for the difference among borrowings. In fact, what lies at the core of the unassimilated versus assimilated distinction is the awareness of speakers when they use a certain word (Poplack et al., 1988). The notion of what a borrowing is lies within the brain of the speaker, and in this work we are only indirectly observing that status through written form. Therefore our definition of borrowing and assimilation cannot be regarded as perfect or universal.

**Ideas about linguistic purity.** The purpose of this project is to analyze the usage of borrowings in the Spanish press. This project does not seek to promote or stigmatise the usage of borrowings, or those who use them. The motivation behind our research is not to defend an alleged linguistic purity, but to study the phenomenon of lexical borrowing from a descriptive and data-driven point of view.

## Acknowledgments

The authors would like to thank Carlota de Benito Moreno, Jorge Diz Pico, Nacho Esteban Fernández, Gloria Gil, Clara Manrique, Rocío Morón González, Aarón Pérez Bernabeu, Monserrat Rius and Miguel Sánchez Ibáñez for their assessment of the annotation quality.

## References

- Gustavo Aguilar, Sudipta Kar, and Thamar Solorio. 2020. [LinCE: A centralized benchmark for linguistic code-switching evaluation](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1803–1813, Marseille, France. European Language Resources Association.
- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. Contextual string embeddings for sequence labeling. In *COLING 2018, 27th International Conference on Computational Linguistics*, pages 1638–1649.
- Beatrice Alex. 2008a. *Automatic detection of English inclusions in mixed-lingual data with an application to parsing*. Ph.D. thesis, University of Edinburgh.
- Beatrice Alex. 2008b. [Comparing corpus-based to web-based lookup techniques for automatic English inclusion detection](#). In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- Elena Álvarez Mellado. 2020a. An Annotated Corpus of Emerging Anglicisms in Spanish Newspaper Headlines. In *Proceedings of the Fourth Workshop on Computational Approaches to Code Switching*, pages 1–8, Marseille, France. European Language Resources Association.
- Elena Álvarez Mellado. 2020b. *Lázaro: An extractor of emergent anglicisms in Spanish newswire*. Master's thesis, Brandeis University.
- Elena Álvarez Mellado, Luis Espinosa-Anke, Julio Gonzalo Arroyo, Constantine Lignos, and Jordi Porta Zamorano. 2021. Overview of ADoBo 2021: Automatic Detection of Unassimilated Borrowings in the Spanish Press. *Procesamiento del Lenguaje Natural*, 67:277–285.
- Gisle Andersen. 2012. Semi-automatic approaches to Anglicism detection in Norwegian corpus data. In Cristiano Furiassi, Virginia Pulcini, and Félix Rodríguez González, editors, *The anglicization of European lexis*, pages 111–130. John Benjamins.
- Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational Linguistics*, 34(4):555–596.
- Emily M. Bender and Batya Friedman. 2018. [Data statements for natural language processing: Toward mitigating system bias and enabling better science](#). *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Juan M Lope Blanch. 1995. Americanismo frente a españolismo lingüísticos. *Nueva revista de filología hispánica*, 43(2):433–440.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Cristian Cardellino. 2019. Spanish Billion Words Corpus and Embeddings. <https://crscardellino.github.io/SBWCE/>.
- José Cañete, Gabriel Chaperon, Rodrigo Fuentes, Jou-Hui Ho, Hojin Kang, and Jorge Pérez. 2020. Spanish pre-trained BERT model and evaluation data. In *PML4DC at ICLR 2020*.
- Paula Chesley. 2010. Lexical borrowings in French: Anglicisms as a separate phenomenon. *Journal of French Language Studies*, 20(3):231–251.
- Michael Clyne, Michael G Clyne, and Clyne Michael. 2003. *Dynamics of language contact: English and immigrant languages*. Cambridge University Press.
- Javier De la Rosa. 2021. [The futility of STILTs for the classification of lexical borrowings in Spanish](#). In *Proceedings of IberLEF 2021*, Málaga, Spain.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Luz Fernández Gordillo. 2014. La lexicografía del español y el español hispanoamericano. *Andamios*, 11(26):53–89.
- Cristiano Furiassi and Knut Hofland. 2007. The retrieval of false anglicisms in newspaper texts. In *Corpus Linguistics 25 Years On*, pages 347–363. Brill Rodopi.
- Matt Garley and Julia Hockenmaier. 2012. [Beefmoves: Dissemination, diversity, and dynamics of English borrowings in a German hip hop forum](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 135–139, Jeju Island, Korea. Association for Computational Linguistics.
- Juan Gómez Capuz. 1997. Towards a typological classification of linguistic borrowing (illustrated with anglicisms in romance languages). *Revista alicantina de estudios ingleses*, 10:81–94.

- Einar Haugen. 1950. The analysis of linguistic borrowing. *Language*, 26(2):210–231.
- Benjamin Heinzerling and Michael Strube. 2018. **BPEmb: Tokenization-free pre-trained subword embeddings in 275 languages**. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing. <https://spacy.io/>.
- Shengyi Jiang, Tong Cui, Yingwen Fu, Nankai Lin, and Jieyi Xiang. 2021. BERT4EVER at ADoBo 2021: Detection of borrowings in the Spanish language using pseudo-label technology. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2021)*. CEUR Workshop Proceedings.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. **Neural architectures for named entity recognition**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California. Association for Computational Linguistics.
- Sebastian Leidig, Tim Schlippe, and Tanja Schultz. 2014. Automatic detection of anglicisms for the pronunciation dictionary generation: a case study on our German IT corpus. In *Spoken Language Technologies for Under-Resourced Languages*.
- Constantine Lignos and Mitch Marcus. 2013. Toward web-scale analysis of codeswitching. Presented at the 87th Annual Meeting of the Linguistic Society of America.
- John M Lipski. 2005. Code-switching or borrowing? No sé so no puedo decir, you know. In *Selected proceedings of the second workshop on Spanish sociolinguistics*, pages 1–15. Cascadilla Proceedings Project Somerville, MA.
- Gyri Smordal Losnegaard and Gunn Inger Lyse. 2012. A data-driven approach to anglicism identification in Norwegian. In Gisle Andersen, editor, *Exploring Newspaper Language: Using the web to create and investigate a large corpus of modern Norwegian*, pages 131–154. John Benjamins Publishing.
- André Mansikkaniemi and Mikko Kurimo. 2012. Un-supervised vocabulary adaptation for morph-based language models. In *Proceedings of the NAACL-HLT 2012 Workshop: Will We Ever Really Replace the N-gram Model? On the Future of Language Modeling for HLT*, pages 37–40. Association for Computational Linguistics.
- Jacob McClelland. 2021. A brief survey of anglicisms among Spanish dialects. *Colorado Research in Linguistics*.
- Giovanni Molina, Fahad AlGhamdi, Mahmoud Ghoneim, Abdelati Hawwari, Nicolas Rey-Villamizar, Mona Diab, and Tamar Solorio. 2016. **Overview for the second shared task on language identification in code-switched data**. In *Proceedings of the Second Workshop on Computational Approaches to Code Switching*, pages 40–49, Austin, Texas. Association for Computational Linguistics.
- Hiroki Nakayama, Takahiro Kubo, Junya Kamura, Yasufumi Taniguchi, and Xu Liang. 2018. doccano: Text annotation tool for human. <https://github.com/doccano/doccano>.
- Eugenia Esperanza Núñez Nogueroles. 2017. An up-to-date review of the literature on anglicisms in Spanish. *Diálogo de la Lengua*, IX, pages 1–54.
- Naoaki Okazaki. 2007. CRFSuite: a fast implementation of Conditional Random Fields (CRFs). <http://www.chokkan.org/software/crfsuite/>.
- Alexander Onysko. 2007. *Anglicisms in German: Borrowing, lexical productivity, and written codeswitching*, volume 23. Walter de Gruyter.
- Chester Palen-Michel, Nolan Holley, and Constantine Lignos. 2021. **SeqScore: Addressing barriers to reproducible named entity recognition evaluation**. In *Proceedings of the 2nd Workshop on Evaluation and Comparison of NLP Systems*, pages 40–50, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Terry Peng and Mikhail Korobov. 2014. python-crfsuite. <https://github.com/scrapinghub/python-crfsuite>.
- Barbara Plank. 2016. **What to do about non-standard (or non-canonical) language in NLP**. In *Proceedings of the 13th Conference on Natural Language Processing (KONVENS 2016)*, pages 13–20, Bochum, Germany.
- Shana Poplack. 2012. What does the nonce borrowing hypothesis hypothesize? *Bilingualism: Language and Cognition*, 15(3):644–648.
- Shana Poplack, David Sankoff, and Christopher Miller. 1988. The social correlates and linguistic processes of lexical borrowing and assimilation. *Linguistics*, 26(1):47–104.
- Chris Pratt. 1980. *El anglicismo en el español peninsular contemporáneo*, volume 308. Gredos.
- Real Academia Española. 2020. Diccionario de la lengua española, ed. 23.4.
- Real Academia Española. 2022. Fundéu. Fundación del Español Urgente. <http://www.fundeu.es>.

Felix Rodríguez González. 1999. Anglicisms in contemporary Spanish. An overview. *Atlantis*, 21(1/2):103–139.

Sagor Sarker. 2020. Codeswitch. <https://github.com/sagorbrur/codeswitch>.

Jacqueline Rae Larsen Serigos. 2017. *Applying corpus and computational methods to loanword research: new approaches to Anglicisms in Spanish*. Ph.D. thesis, The University of Texas at Austin.

Thamar Solorio, Elizabeth Blair, Suraj Maharjan, Steven Bethard, Mona Diab, Mahmoud Ghoneim, Abdelati Hawwari, Fahad AlGhamdi, Julia Hirschberg, Alison Chang, and Pascale Fung. 2014. **Overview for the first shared task on language identification in code-switched data**. In *Proceedings of the First Workshop on Computational Approaches to Code Switching*, pages 62–72, Doha, Qatar. Association for Computational Linguistics.

Yulia Tsvetkov and Chris Dyer. 2016. Cross-lingual bridges with models of lexical borrowing. *Journal of Artificial Intelligence Research*, 55:63–93.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. **Transformers: State-of-the-art natural language processing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

## A Data sources

See Table 9 for the URLs and licenses of the sources used in the dataset.

## B Annotation guidelines

### B.1 Objective

This document proposes a set of guidelines for annotating emergent unassimilated lexical borrowings, with a focus on English lexical borrowings (or anglicisms). The purpose of these annotation guidelines is to assist annotators to annotate unassimilated lexical borrowings from English that appear in Spanish newswire, i.e. words from English origin that are introduced into Spanish without any morphological or orthographic adaptation.

This project approaches the phenomenon of lexical borrowing from a synchronic point of view, which means that we will not be annotating all words that have been borrowed at some point of the history of the Spanish language (like arabisms), but only those that have been recently imported and have not been integrated into the recipient language (in this case, Spanish).

### B.2 Tagset

We will consider two possible tags for our annotation: ENG, for borrowings that come from the English language (or anglicisms), and OTHER for other borrowings that comply with the following guidelines but that come from languages other than English.

### B.3 Defining an unassimilated lexical borrowing

In this section we provide an overview of what words will be considered as unassimilated lexical borrowings for the sake of our annotation project.

#### B.3.1 Definition and scope

The concept of *linguistic borrowing* covers a wide range of linguistic phenomena. We will first provide a general overview of what lexical borrowing is and what will be understood as an anglicism within the scope of this project.

Lexical borrowing is the incorporation of single lexical units from one language (the donor language) into another language (the recipient language) and is usually accompanied by morphological and phonological modification to conform with the patterns of the recipient language (Haugen, 1950; Onysko, 2007; Poplack et al., 1988).

Media	URL	License
elDiario.es	<a href="https://www.eldiario.es/">https://www.eldiario.es/</a>	CC BY-NC 4.0
Agencia Sinc	<a href="https://www.agenciasinc.es/">https://www.agenciasinc.es/</a>	CC BY 4.0
El Salto	<a href="https://www.elsaltodiario.com/">https://www.elsaltodiario.com/</a>	CC BY-SA 3.0
La Marea	<a href="https://www.lamarea.com/">https://www.lamarea.com/</a>	CC BY-SA 3.0
Cuarto poder	<a href="https://www.cuartopoder.es/">https://www.cuartopoder.es/</a>	CC BY-NC 3.0
Genbeta	<a href="https://www.genbeta.com/">https://www.genbeta.com/</a>	CC BY-NC 3.0
Bebe y más	<a href="https://www.bebesymas.com/">https://www.bebesymas.com/</a>	CC BY-NC 3.0
Diario del viajero	<a href="https://www.diariodelviajero.com/">https://www.diariodelviajero.com/</a>	CC BY-NC 3.0
El blog salmón	<a href="https://www.elblogsalmon.com/">https://www.elblogsalmon.com/</a>	CC BY-NC 3.0
Espinof	<a href="https://www.espinof.com/">https://www.espinof.com/</a>	CC BY-NC 3.0
Motor pasión	<a href="https://www.motorpasion.com/">https://www.motorpasion.com/</a>	CC BY-NC 3.0
Pop rosa	<a href="https://www.poprosa.com/">https://www.poprosa.com/</a>	CC BY-NC 3.0
Vida extra	<a href="https://www.vidaextra.com/">https://www.vidaextra.com/</a>	CC BY-NC 3.0
Vitónica	<a href="https://www.vitonica.com/">https://www.vitonica.com/</a>	CC BY-NC 3.0
Xataka	<a href="https://www.xataka.com/">https://www.xataka.com/</a>	CC BY-NC 3.0
Xataka Ciencia	<a href="https://www.xatakaciencia.com/">https://www.xatakaciencia.com/</a>	CC BY-NC 3.0
Xataka Android	<a href="https://www.xatakandroid.com/">https://www.xatakandroid.com/</a>	CC BY-NC 3.0
El orden mundial	<a href="https://elordenmundial.com/">https://elordenmundial.com/</a>	CC BY-NC-ND
Foro atletismo	<a href="https://www.foroatletismo.com/">https://www.foroatletismo.com/</a>	CC BY-NC 2.1
Pikara Magazine	<a href="https://www.pikaramagazine.com/">https://www.pikaramagazine.com/</a>	CC BY-NC-ND
Microsiervos	<a href="https://www.microsiervos.com/">https://www.microsiervos.com/</a>	Used with written permission

Table 9: Media included in the corpus

Anglicisms are lexical borrowings that come from the English language (Gómez Capuz, 1997; Pratt, 1980; Rodríguez González, 1999; Núñez Nogueroles, 2017). For our annotation project, we will focus on direct, unassimilated, emerging anglicisms, i.e. lexical borrowings from the English language into Spanish that have recently been imported and that have still not been assimilated into Spanish, that is, words like *smartphone*, *influencer*, *hype*, *lawfare* or *reality show*.

Although this project focuses on lexical borrowings from English, we will also consider borrowings from other languages that comply with these guidelines. Borrowings from the English language will be annotated with the tag ENG, while borrowings from other languages shall be annotated with the tag OTHER:

... financiados a través de la plataforma de [crowdfunding](ENG) del club [gourmet](OTHER) que tengas más cerca<sup>11</sup>

Other types of borrowings, such as semantic calques, syntactic anglicisms or literal translations will be considered beyond the scope of these annotation project and will not be covered in these

<sup>11</sup>Examples in these guidelines will display the lexical borrowing that should be labeled between square brackets, with the the corresponding tag in parentheses. Examples with no words marked with brackets will illustrate cases where no lexical borrowing should be tagged.

guidelines.

### B.3.2 Types of lexical borrowing

Lexical borrowings can be adapted (the spelling of the word is modified to comply with the phonological and orthographic patterns of the recipient language, as in *fútbol* or *tuit*) or unadapted (the word preserves its original spelling: *millennial*, *newsletter*, *like*). For this annotation project, we will be focusing on unassimilated lexical borrowings: this means that adapted borrowings will be ignored and only unadapted borrowings will be tagged (see Section B.4.2 for a full description on the differences between adapted and unadapted borrowings).

### B.3.3 Multiword borrowings

Lexical borrowings can be both single-token units (*online*, *impeachment*), as well as multiword expressions (*reality show*, *best seller*). Multitoken borrowings will be labeled as one entity.

imagina ser un '[tech bro]' con millones de dólares (ENG)

The annotation should however distinguish between a multitoken borrowing and adjacent borrowings. A phrase like *signature look* is a multiword borrowing (the full phrase has been borrowed as a single unit) and should be annotated as such.

para recrear su [total look] (ENG)

However, a phrase like *look sporty* follows the NAdj order that is typical of Spanish grammar (but

impossible in English): these are in fact two separate borrowings (*look* and *sporty*) that have been borrowed independently and happen to be collocated in a phrase. The annotation should capture these nuances:

```
un [look] (ENG) [sporty] (ENG)
perfecto
```

### B.3.4 Origin of the borrowings

Establishing the origin of a certain borrowings can sometimes be tricky, as the language of origin can sometimes be disputed. Additionally, certain borrowings might have originated in a certain language, but may have reached the recipient language through another language.

In order to establish the origin of borrowings, the origin attributed by reference dictionaries and institutions (Real Academia Española, 2020, 2022) will be followed.

This means that words like *junior* and *senior* (whose frequency and perhaps even their pronunciation may have changed due to the influence of English) will still be considered as latinisms, as DLE registers their adapted versions (*júnior* and *sénior*) as such (and mentions no English influence). Similarly, the word *barista* might have entered the Spanish language via English, but RAE's Observatorio de Palabras considers it of Italian origin (and should therefore be annotated with OTHER label).

## B.4 What is not an unassimilated lexical borrowing

In the previous section we provided an overview of what words will be considered as an unassimilated lexical borrowing for the sake of our annotation project. In this section we will cover what an unassimilated lexical borrowing is *not*.

There are several phenomena that are close enough to unassimilated borrowing and that can sometimes be mistaken with. In this section we will list what phenomena will not be considered as unassimilated lexical borrowings (and are therefore beyond the scope of our annotation project), as well as provide guidelines in order to distinguish these cases and adjudicate them.

We will focus on three main phenomena: assimilated borrowings, proper names and code-mixed inclusions.

Figure 1 summarizes the decision steps that can be followed when deciding if a certain word should

be labeled or not as a lexical borrowing.

### B.4.1 Assimilated vs unassimilated borrowings

This annotation project aims to capture unassimilated lexical borrowings. As a general rule, all unadapted lexical borrowings should be tagged. This means that direct borrowings that have not gone through any morphological or orthographic modification process should be labeled.

Lexical adaptation, however, is a diachronic process and, as a result, what constitutes an unadapted borrowing is not clear-cut. The following guidelines define what borrowings will be considered as unassimilated (and therefore should be tagged) versus those that have already been integrated into the recipient language (and therefore should not be tagged).

### B.4.2 Adapted borrowings

Words that have already gone through orthographical or morphological adaptation (such as *fútbol*, *líder*, *tuit* or *espóiler*) will be considered assimilated and therefore should not be labeled. Partial adaptations (such as *márketing*, where an accent has been added) will also be excluded.

Borrowings that have not been adapted but whose original spelling complies with graphophonological rules of Spanish (and are therefore unlikely to be further adapted, such as *bar*, *fan*,

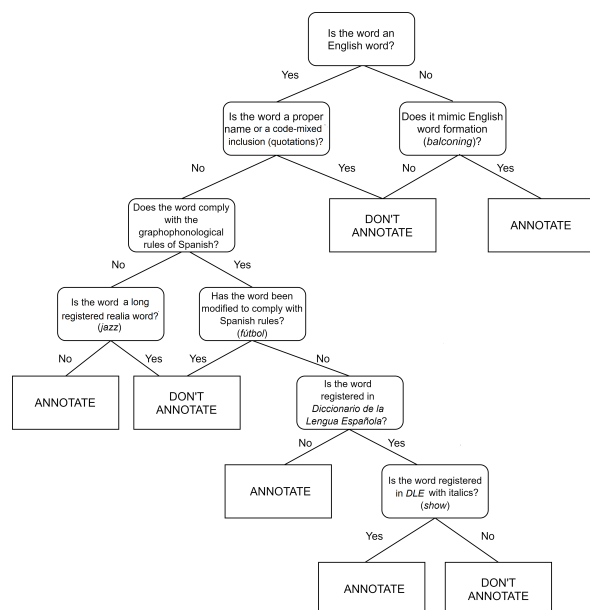


Figure 1: Decision steps to follow during the annotation process to decide whether to annotate a word as an anglicism

*web, internet, club, set* or *videoclip*) will be tagged as a borrowing or not depending on how recent or emergent they are. In order to determine which unadapted, graphophonologically acceptable borrowings are to be annotated, the latest online version of the *Diccionario de la lengua española* (Real Academia Española, 2020) will be consulted (as of February 2021)<sup>12</sup>. If the DLE dictionary already registers the word with that meaning and with no italics or quotation marks, then it will be considered assimilated and therefore should not be tagged.

This means that a word like *set* (when used to refer to a collection of things, a television studio or a part of a tennis match) will be considered assimilated because it is already registered in DLE dictionary with no italics, and therefore should not be labeled as ENG. On the other hand, a word like *nude*, although its spelling also complies with Spanish graphophonological rules, will be considered an unassimilated borrowing because it has not been registered yet in the dictionary, and should therefore be tagged as such.

ganó el primer set

los tonos '[nude]' (ENG)

It should be noted that this guideline only applies to lexical borrowings that comply with graphophonological rules of Spanish. Unadapted lexical borrowings that do not comply with graphophonological rules of Spanish (such as *show, look*, etc) will be tagged as borrowing, regardless of whether the word is included in the dictionary or not (although see section B.4.6 for exceptions to this).

It is important to emphasize that, in order for an unadapted graphophonologically-compliant borrowing to be considered assimilated it should be registered in the dictionary both without italics and with the corresponding meaning. For instance, a word like *top* (that is graphophonologically acceptable in Spanish) is registered in DLE with no italics, but it is only registered with the meaning of a piece of clothing. The word *top* as referring to the upper part of something (as in *top 5*) is not registered. Consequently, the borrowing *top* will be considered assimilated when referring to the piece of clothing, but unassimilated when used to talk about the best elements of a ranking or the upper part of something.

<sup>12</sup><https://dle.rae.es/>

un top estampado

el [top] cinco de artistas  
(ENG)

la [top] desfiló (ENG)

Similarly, the word *post* will not be considered a borrowing when used as a prefix of Latin origin, but will be labeled with ENG when used to refer to something that is published on a social media platform.

el mundo post pandemia

un [post] de Facebook (ENG)

Additionally, assimilated borrowings can still be part of new unassimilated borrowings, in which case they will be labeled as such:

un [boys club] (ENG)

#### B.4.3 Words derived from foreign lexemes

Words derived from foreign lexemes that do not comply with Spanish orthotactics but that have been morphologically derived following the Spanish paradigm (such as *hacktivista, randomizar, shakespeariano*) will be considered assimilated and should therefore not be labeled as a borrowing.

Compound names where one of the lexemes is a borrowing will be labeled as a borrowing or not according to the degree of independence among the lexemes. A verb+noun compound (as *caza-clicks*) will not be labeled as a borrowing, because the elements are not independent from one another. However, noun-noun compounds where each of the lexemes work can work independent from one can be labeled as borrowings:

una casa-[loft] (ENG)

Similarly, prefixed borrowings will be labeled as a borrowing, as long as the borrowing keeps independence from the prefix:

la ex [influencer] (ENG)

For prefixed borrowings, it should be checked whether the prefix can also be considered part of the borrowing:

los [nano influencers] (ENG)



#### B.4.4 Number inflection

Unassimilated borrowings may be incorporated as invariable in number *los master*, with the same plural inflection that they had in the donor language (*los pappardelle*) or may form a new plural that is non-existent in the donor language (*los pappardelles*). For number inflection, we follow the same criteria that DLE (Real Academia Española, 2020) follows: a non-Italian plural like *pizzas* is still regarded as unadapted (and therefore should be written italicized even when the true Italian plural would be *pizze*). Consequently, non-assimilated borrowings that have a non-canonical plural inflection form will still be considered as unassimilated borrowings and labeled as such.

```
una serie de animación de  
[mechas] (OTHER)
```

#### B.4.5 Pseudoanglicisms

Words that do not exist in English (or exist with a different meaning) but were coined following English morphological paradigm to imitate English words (such as *footing* or *balconing*) will be annotated as anglicisms.

```
la imagen del '[balconing]' y  
las excursiones etílicas (ENG)
```

```
practicaba [footing] por la  
calle (ENG)
```

#### B.4.6 Realia words

Borrowings that refer to culture-specific elements (often called *realia words*) that were imported long ago but that have remained unadapted will not be tagged as borrowing. This means that if a borrowing is not adapted (i.e. its form remained exactly as it came from the donor language) but refers to a particular cultural object that came via the original language, that has been registered for a while in Spanish dictionaries and is not perceived as new anymore, then it will not be tagged as a borrowing, even if it does not comply with graphophonologic rules of Spanish.

The purpose of this guideline is to account for cultural terms such as *pizza*, *whisky*, *jazz*, *blues*, *banjo* or *sheriff*. These are all borrowings that are reluctant to be adapted or translated, even when they have been around in the Spanish language for long. The reason is that they refer to cultural inventions (the name was imported along with the object

it refers to), and, given their cultural significance, they never competed with a Spanish equivalent and are seen as assimilated.

Therefore, unadapted borrowings that refer to cultural innovations (such as music, cooking, sport names etc) and that have been registered for long in the Spanish language<sup>13</sup> will not be tagged as emergent borrowings.

It should be noted that this only applies to borrowings that have been around enough time to be registered in dictionaries. A word like *hip hop* is a realia word, but it is still recent enough and has not been registered in the dictionary. In that case, it should be considered as unassimilated and tagged as such.

#### B.4.7 Latinisms

Borrowings that were introduced directly from Latin language (such as *deficit*, *curriculum*, etc) will not be considered emergent and therefore will not be tagged as a borrowing. However, it should be noted that unassimilated borrowings from other languages that happen to have a Latin etymology and that are introduced with a distinct meaning (such as *adlib* or *premium* etc) will still be tagged as borrowings.

### B.5 Borrowings vs names

#### B.5.1 Proper nouns

Non-Spanish proper nouns will not be tagged as borrowings. These include:

- person names: *Bernie Sanders*.
- organization names: *WikiLeaks*.
- product names: *Slack*.
- location names: *Times Square*.
- dates and celebrations: *St. Patrick's Day*, *Black Friday*.
- event names: *Brexit*, *procés*.
- social and political movements: *Black Lives Matter*, *MeToo*.
- treaties and documents: *New Deal*, *Privacy Shield*, *French Tech Visa*.

<sup>13</sup>RAE dictionary <https://dle.rae.es/>, Mapa de diccionarios <https://webfrrl.rae.es/ntllet/SrvltGUILoginNtlletPub> and CREA <http://corpus.rae.es/creanet.html> and CORPES <https://webfrrl.rae.es/CORPES/view/inicioExterno.view> can be consulted

- titles of cultural productions: *Stranger Things*.

### B.5.2 Borrowings in proper nouns

Borrowings that appear as part of proper nouns or named entities (such as book titles or organization names, as in *Los Hermanos Podcast*) will not be labeled as borrowings.

### B.5.3 Proper nouns in borrowings

Multiword borrowings and expressions can sometimes include proper nouns. Even when a proper noun in isolation cannot be considered a borrowing, proper nouns within a borrowed expression will be considered part of the borrowing, as long as the proper noun is part of the borrowing and is used following the grammar rules of the donor language (for example, in an English noun noun compound):

Tecnología [made in Spain]  
(ENG)

[Google cooking] (ENG)

### B.5.4 Names of institutions and political roles

Non-Spanish names that refer to political institutions (such as *Parlament* or *Bundestag*) or to political roles and figures (*lehendakari*, *president*, *conseller*) will be excluded and will not be tagged as borrowings.

### B.5.5 Words derived from proper nouns

Words derived from proper nouns (via metonymy or eponymy) will not be tagged as a borrowing, as long as the relation with the proper noun they come from is transparent to the speaker such as:

- products: *un iPhone*, *un whatsapp*, *un bizum*, *un Scalextric*, *el Satisfyer*.
- works of arts: *un monet*
- characters: *un frankenstein*.

However, borrowings that originated from a proper noun in the donor language but entered the Spanish language as common nouns and are currently recognized as such, will be labeled as borrowings. In order to adjudicate which of these words are still used in Spanish as proper names and which are common nouns, dictionaries and other reference works can be consulted.

### B.5.6 Names of peoples or languages

Names of peoples or languages (such as *inuit*) will not be labeled as borrowings, even if the word is borrowed from another language and is not registered in Spanish dictionaries.

### B.5.7 Fictitious creatures

Unadapted names of fictitious creatures (such as *hobbit* or *troll*) will be labeled as a borrowing.

En un agujero en el suelo  
vivía un [hobbit] (ENG)

### B.5.8 Scientific units

Unadapted borrowings that refer to widespread scientific units (such as *hertz*, *newton*, *byte*, etc) will be considered assimilated and should not be tagged as a borrowing

### B.5.9 Names of species

Scientific names of a species (such as Latin names) will not be tagged as a lexical borrowing (*anisakis*). Names of fruit, vegetable and plant varieties (such as *manzana golden*, *patatas Kennebec* or *aguacate Hass*) will also be excluded.

## B.6 Borrowings vs other code-mixed inclusions

Borrowing (using units from one language in another language) and codeswitching (intertwining segments of different languages in the same discourse) have frequently been described as a continuum (Clyne et al., 2003), with a fuzzy frontier between the two. As a result, it can be difficult to tell the difference between borrowing and other code-mixed inclusions. The following guidelines can assist annotators adjudicate edge cases.

When in doubt while dealing with code-mixed inclusion, the annotator may find it helpful to ask the following question as a rule of thumb: would it make sense to have this non-Spanish word registered in a dictionary of Spanish? If the answer is no (for instance, because the word reflects the literal quotation of what someone said or because the inclusions is metalinguistic usage rather than borrowing), then we are probably not in front of a borrowing but of another type of code-mixed inclusion (and should not be tagged as a borrowing).

### B.6.1 Acronyms and acronym expansions

We consider acronyms to be a different phenomenon from borrowings. Consequently,

acronyms will not be tagged as a borrowing, even if the acronym is of non-Spanish origin

un lector de CD

An acronym however may be tagged as a borrowing if it appears as part of a borrowed multiword expression, as in *CD player*, *peak TV*, *PC gaming*:

un [CD player] (ENG)

Acronym expansions, that is, the expansion of an acronym into the words that form the acronym (that is usually added in between brackets after an acronym has been introduced) will also not be considered a borrowing:

La técnica de PCR (protein chain reaction)

It is important to note that for a sequence to be considered as an acronym expansion it must appear after the acronym has been introduced and serve as a gloss to it (so that it expands what the letters in the acronym stand for). Usages where the full sequence is introduced in the text and later on acronymized for the sake of brevity can still be considered as borrowings.

Utilizaron técnicas de [Machine Learning] (también conocido como ML) (ENG)

### B.6.2 Digits

Similarly to proper nouns, digits in isolation cannot be considered borrowings. As a result, we cannot take for granted that digits within the surroundings of a borrowing will automatically be part of the borrowing.

[top ten] (ENG)

[top] 10 (ENG)

However, if the word order of the tokens makes it clear that the digit is part of a multitoken borrowing (because the order complies with the grammatical structure of an English noun-noun compound), we can label it as part of the borrowing:

los [10% banks] (ENG)

### B.6.3 Metalinguistic usage

Non-Spanish words that appear to refer to the word itself in linguistic discourse and do not cover a lexical gap will not be tagged as a borrowing:

El término viene de la palabra "ghost", que en inglés es 'fantasma'

It should be noted that the newer, less adapted, less transparent a new word is, the more likely that the speaker will be aware of the decoding difficulty it may pose to the reader and will decide to add some sort of metalinguistic strategy or awareness around it, in the form of metacomments, word-pointers, meaning explanations, etc (*known as, so called*). Borrowings with these types of signals will still be considered borrowings, as long as they are covering a lexical gap.

True metalinguistic usage where the foreign word covers no lexical gap but exclusively provides linguistic information (such as etymological information) will not be considered a borrowing.

### B.6.4 Literal quotations

Words or sequences in languages other than Spanish that are reflecting literally what someone said or wrote (as in a quotation, a statement or a slogan) will not be considered a borrowing.

El eslogan 'Make America Great Again'

Es uno de los primeros resultados de Google cuando alguien busca "remote work in Spain" (trabajo en remoto en España).

### B.6.5 Expressions

In general terms, multiword borrowings will be tagged as borrowings. However, phrases and expressions that are not integrated into the sentence will be excluded. This means that autonomous expressions that are rather code switched sentences (rather than real borrowings) that work as a unit totally independently of the rest of the linguistic context (and that we would not expect to be registered in a dictionary) will not be considered or tagged as a borrowing.

La innovación y la competencia tan escasas en la radiotelevisión o peor aún en Internet ("the winners takes all" o "most").

## B.7 Limitations of these guidelines

These guidelines are intended to assist annotators when labeling lexical borrowings. These guidelines, however, were created with a specific goal in mind (to capture unassimilated English lexical borrowings from a corpus of Spanish newswire) and may not be suitable if applied to a project with a different scope. These are some of the shortcomings and limitations that these guidelines may have.

### B.7.1 Text genre

These guidelines were designed to specifically capture borrowings in a corpus of Spanish newswire. Newswire is a very specific genre of text that by no means represent the whole of a language (Plank, 2016).

### B.7.2 Donor language

These guidelines were created with English lexical borrowings in mind, which are the most frequent source of borrowing today in the Spanish press. Although the criteria can be applied to other languages as well (and in fact the annotation tagset we propose includes the tag `OTHER` to account for borrowings from other languages other than English), a more fine-grained approach would require further guidelines.

### B.7.3 Synchronic approach to borrowing

This project approaches emergent, unassimilated lexical borrowing in a synchronic fashion. The process of borrowing and the notion of assimilation is, however, time-dependent. A diachronic approach to lexical borrowing would require a wider scope, a different theoretical framework and an expanded set of criteria.

### B.7.4 Geographic variety

The guidelines in this document were designed to capture borrowings used in Spanish newspapers, that is, written in the variety of Spanish that is spoken in Spain and may not be suitable to account for other dialects. For instance, according to the guidelines we have just introduced, a word like *living* (that is used heavily in some Latin American varieties to refer to the living room) would be considered unassimilated. It is arguable whether these criteria would be suitable for a project that tried to capture emergent lexical borrowings in Argentinian text, for example.

## C Data statement

We document the information concerning our dataset following the data statement format proposed by Bender and Friedman (2018).

**Data set name:** Corpus of Anglicisms in the Spanish Press (COALAS)

**Data set developer:** Elena Álvarez-Mellado

**Dataset license:** Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0)

**Link to dataset:** <https://github.com/lirondos/coalas>

### C.1 Curation rationale

The corpus consist of a collection of Spanish newspaper sentences. These sentences are annotated with unassimilated lexical borrowings.

Data was collected separately for the training, development, and test sets to ensure minimal overlap in borrowings, topics, and time periods (see Section C.5).

To focus annotation efforts for the training set on articles likely to contain unassimilated borrowings, the articles to be annotated were selected by first using a baseline model and were then human-annotated. To detect potential borrowings, the CRF model and data from Álvarez Mellado (2020b) was used along with a dictionary look-up pipeline. Articles that contained more than 5 borrowing candidates were selected for annotation.

The main goal of data selection for the development and test sets was to create borrowing-dense, OOV-rich datasets, allowing for better assessment of generalization. To that end, the annotation was based on sentences instead of full articles. If a sentence contained a word either flagged as a borrowing by the CRF model, contained in a wordlist of English, or simply not present in the training set, it was selected for annotation. This data selection approach ensured a high number of borrowings and OOV words, both borrowings and non-borrowings. While the training set contains 6 borrowings per 1,000 tokens, the test set contains 20 borrowings per 1,000 tokens. Additionally, 90% of the unique borrowings in the development set were OOV (not present in training). 92% of the borrowings in the test set did not appear in training (see Table 2).

### C.2 Language variety

The language of this corpus is Standard Spanish (ISO 639-1 es). Since all the sources in the dataset

are Spanish newspapers, the dialect is standard European Spanish.

### **C.3 Speaker demographic**

No detailed information was collected regarding the demographics of the authors of the collected sentences. However, we can infer that the authors of the text were Spanish journalists aged between 20-65.

### **C.4 Annotator demographic**

The annotator was a 30-35 year-old female graduate student from Spain, who was trained in linguistics and computational linguistics, whose native language is Spanish and who had extensive professional proficiency in English.

### **C.5 Speech situation**

The training set consists of a collection of articles appearing between August and December 2020 in *elDiario.es*, a progressive online newspaper based in Spain. The development set contains sentences in articles from January 2021 from the same source.

The test set consisted of annotated sentences extracted in February and March 2021 from a diverse collection of online Spanish media that covers specialized topics rich in lexical borrowings and usually not covered by *elDiario.es*, such as sports, gossip or videogames (see Table 2).

All the data came from written sources and was presumably edited according to the style guides of each source.

### **C.6 Text characteristics**

All the sentences in this dataset come from journalistic texts. The data in the training set and development set come from a general online newspaper. The data in the test set come from a diverse collection of online media covering specialized topics, which include politics, feminism, economy, gossip, videogames, cinema & TV, technology, science, travel, parenthood, lifestyle, sports and automobiles.

### **C.7 Recording quality**

N/A

### **C.8 Other**

N/A

### **C.9 Provenance appendix**

See Table 9.