

# ADoBo 2021: The futility of STILTs for the classification of lexical borrowings in Spanish

Javier de la Rosa<sup>1,2</sup>[0000–0002–9143–5573]

<sup>1</sup> National Library of Norway AI-Lab, Mo i Rana 8624, Norway  
`versae@nb.no`

<sup>2</sup> LINHD, UNED, Madrid, 28040, Spain  
`versae@linhd.uned.es`

**Abstract.** The first edition of the IberLEF 2021 shared task on automatic detection of borrowings (ADoBo) focused on detecting lexical borrowings that appeared in the Spanish press and that have recently been imported into the Spanish language. In this work, we tested supplementary training on intermediate labeled-data tasks (STILTs) from part of speech (POS), named entity recognition (NER), code-switching, and language identification approaches to the classification of borrowings at the token level using existing pre-trained transformer-based language models. Our extensive experimental results suggest that STILTs do not provide any improvement over direct fine-tuning of multilingual models. However, multilingual models trained on small subsets of languages perform reasonably better than multilingual BERT but not as good as multilingual RoBERTa for the given dataset.

**Keywords:** Borrowings · STILTs · Multilingual Models.

## 1 Introduction

The sociopolitical and sociocultural changes speakers undergo tend to be somehow reflected in their lexicon [21]. The borrowing of words from a donor language to a recipient one is a common mechanism driving language change and word formation [23]. The lexical units being incorporated into the recipient language usually undergo morphological and phonological transformations as to better conform with the features of the recipient language [18]. Interestingly, regardless of the reasons as to why languages borrow words from others, there are a few common patterns that emerge in the process. For example, it seems like content words are much more frequently borrowed than function words, since it is more

---

*IberLEF 2021, September 2021, Málaga, Spain.*

Copyright © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

likely for a language to borrow nouns or verbs rather than prepositions or conjunctions. Nouns in particular might be benefiting from referential transparency and morphosyntactic freedom [14]. However, before a word becomes fully assimilated into a recipient language with the proper morphological or orthographic modifications, it is not uncommon to see the adapted and unadapted versions of borrowings coexisting (e.g., *whisky* and *güisqui* are both correct in the Spanish orthography). Identifying what words enter a language and how they do so is critical for the understanding of the development of a language.

## 2 Related work

Given its global dominance in many domains of our daily lives, automatic approaches for detecting borrowings are mainly focused on words of English origin, that is, anglicism detection [10,7,5,11].

For Spanish, a loanword identification algorithm for Argentine Spanish was proposed in [20], which provided the lemmatized form of the tokens, identified named entities, and preserved loan phrases. Although somewhat successful in the identification of anglicisms in the chosen news corpus, the algorithm was designed as a binary classifier in which every word was labeled as Spanish or English. This impacted negatively the number of false positives and made the system incapable of reliably identifying code-switching or adapted loans.

More recently, [3] showed that framing the problem as a token classification task for the extraction of emergent anglicisms in Spanish newswire provided very good results. The approach compared a conditional random fields (CRF) classifier built upon handcrafted features with word and character-level embeddings, to a bidirectional long short-term memory neural network with an extra conditional random fields layer on top (BiLSTM-CRF). The CRF-only method outperformed the neural network and achieved an F1 score of 87.82 on the test. Interestingly, the corpora made explicit distinctions between unadapted borrowings of English origin or other origin.

## 3 Language models

Previous methods have not yet relied on modern language models for the detection of borrowings. In this work, we used pre-trained transformer-based language models to approach the problem of borrowing detection as a token classification task. Given the prevalence of morphological and orthographic differences between emergent unadapted borrowings and assimilated words that have been part of a language for a long time, our hypothesis is that if a language model is capable of acquiring enough morphological and syntactical information should also be able to use this information to perform well on the detection of borrowings.

In order to test this hypothesis, we compared the results of fine-tuning BERT-based language models [9] on a corpus of Spanish news articles with and without supplementary training on intermediate labeled-data tasks (STILTs) [17] related to the kind of language features that precisely makes identifying unadapted

borrowings possible. Specifically, named entity recognition, part of speech, code-switching, and language identification. The rationale is that if a model is able to perform well in any of these tasks, it should also do well at detecting borrowings.

Supplementary training is a technique that supplements “language model-style pre-training with further training on data-rich supervised tasks” [17]. When using intermediate tasks such as natural language inference, performance on benchmarks like GLUE [25] improves over BERT [9] or ELMo [16]. Moreover, it seems to excel in situations with very limited training data [24], as it is the case with the labels in the ADoBo corpus [4].

In this sense, recent work has focused on systematically analyzing the performance of STILTs on both token and sequence classification. In [19], the authors experimented with a diverse set of 42 intermediate and 11 target English tasks. For the token classification tasks, supplementary tasks based on POS, NER, and emergent entity recognition datasets [8] showed to benefit from each other. However, the authors mention the risks of not choosing a proper supplementary task as a bad chosen intermediate task can degrade performance on the target task considerably.

## 4 Methods

The corpus, released as part of the IberLEF 2021 shared task on Automatic Detection of Borrowings (ADoBO) [4], contains articles from Spanish newswire which are annotated with direct, unadapted, lexical borrowings following a set of publicly made annotation guidelines with specific assimilated borrowings, proper names and code-mixed situations.

The articles are sentence segmented and split into words. The annotations are made at the word level following the BIO annotation schema with two possible categories: **ENG** for English borrowings, and **OTHER** for lexical borrowings originating in languages other than English. The rest, including punctuation marks, are labeled using **O** and omitted from evaluation. Table 1 shows the distributions of sentences, words, and **ENG** and **OTHER** labels per corpus split. It is specially striking the number of words labeled as **ENG** in the validation set when compared to the validation and test sets. Also worth mentioning the scarcity of words labeled as **OTHER** in general, which made learning the label solely based on this corpus a really challenging task.

**Table 1.** Distributions of sentences, words, and **ENG** and **OTHER** labels.

<b>Split</b>	<b>Sentences</b>	<b>Words</b>	<b>ENG</b>	<b>OTHER</b>
Train	8,216	231,126	1,701	32
Validation	2,025	82,578	424	57
Test	1,811	57,998	1,671	52

The general procedure for applying STILTs follows three steps:

1. First, a model is trained on a semi-supervised task with no labeled data such as a language modeling task to gain some language reasoning capabilities.
2. The model is then further trained on an intermediate task for which plenty of labeled data is available.
3. Finally, the resulting model is fine-tuned further on the target task and evaluated.

Since by nature the task involves words in more than one language, as baselines models we first fine-tuned multilingual BERT (mBERT), XLM RoBERTa (XLM-R) [13], a 5 languages version of mBERT including English, French, Spanish, German and Chinese (mBERT-5lang) [1], and two different monolingual Spanish BERT versions, one extracted from mBERT (mBERT-1lang) [1] and the other pre-trained from scratch (BETO) [6].

As intermediate tasks for supplementary training, we chose mBERT models already fine-tuned on LinCE [2], a benchmark for linguistic code-switching evaluation that includes language identification (LinCE-LID), parts of speech tagging (LinCE-POS), and named entity recognition (LinCE-NER) tasks over Spanish-English code-switched data. A version of the Spanish BERT BETO fine-tuned on CoNLL-2002 for Spanish POS tags [22] (BETO-POS) and a version of XLM-RoBERTa fine-tuned on the same dataset for NER (XLM-R-NER) tags were also included. Distributions of sentences and words for the datasets used in these supplementary tasks are shown in Table 2.

**Table 2.** Distributions of sentences and words (sentences / words) for the datasets used in STILTs.

Dataset	Train	Validation	Test
CoNLL-2002	8,324 / 264,715	1,916 / 52,923	1,518 / 51,533
LinCE-LID	21,030 / 253,221	3,332 / 40,391	8,289 / 97,341
LinCE-NER	27,893 / 217,068	4,298 / 33,345	10,720 / 82,656
LinCE-POS	33,611 / 404,428	10,085 / 122,656	23,527 / 281,579

## 5 Results

We run all experiments on 2x24GB NVIDIA GPU RTX 6000, doing grid searches of hyperparameters for each model with learning rates of 1e-5, 2e-5, 3e-5, and 4e-5, and for 3, 5, and 10 epochs. We used the AdamW optimizer, with no weight decay, and a 10% of steps for warmup. We did 3 runs with different seeds and chose the best model on the validation set while reporting results on the test set. We are reporting precision, recall, and F1 micro scores as obtained by `sequeval` [15].

**Table 3.** Best values of precision (P), recall (R), and F1 score for the validation set on baseline and STILTs models. Best scores in **bold**, second best F1 scores underlined.


Model	ENG			OTHER			Total		
	P	R	F1	P	R	F1	P	R	F1
mBERT	84.47	88.89	86.62	65.22	30.61	41.67	83.19	80.85	82.00
mBERT-1lang	80.65	88.56	84.42	<b>73.91</b>	34.69	<u>47.22</u>	80.22	81.13	80.67
mBERT-5lang	85.98	<b>90.20</b>	<u>88.04</u>	62.50	<b>40.82</b>	<b>49.38</b>	83.85	<b>83.38</b>	<b>83.62</b>
BETO	79.28	86.27	82.63	47.62	20.41	28.57	77.40	77.18	77.29
XLM-R	<b>86.48</b>	89.87	<b>88.14</b>	53.85	28.57	37.33	<b>84.01</b>	81.41	<u>82.69</u>
LinCE-LID <sub>STILT</sub>	29.29	18.95	23.02	12.50	2.04	3.51	<b>28.64</b>	16.62	21.03
XLM-R-NER <sub>STILT</sub>	28.46	<b>23.53</b>	<b>25.76</b>	50.00	2.04	<u>3.92</u>	28.63	<b>20.56</b>	<b>23.93</b>
LinCE-NER <sub>STILT</sub>	23.68	20.59	22.03	33.33	2.04	3.85	23.79	18.03	20.51
BETO-POS <sub>STILT</sub>	24.43	17.65	20.49	<b>100.00</b>	2.04	<b>4.00</b>	24.77	15.49	19.06
LinCE-POS <sub>STILT</sub>	<b>29.65</b>	19.28	<u>23.37</u>	16.67	2.04	3.64	29.27	16.9	<u>21.43</u>

Tables 3 and 4 overwhelmingly show that STILTs have absolutely no positive effect on the classification of borrowings. No matter what kind of supplementary training is used, the STILT models perform almost 4 times worse than the baseline finetuning. Among the best performing models on direct finetuning on the validation set, XLM-R and mBERT-5lang achieve somewhat similar F1 scores in total, with 83.62 F1 points for mBERT-5lang and 82.69 for XLM-R; both above mBERT. The biggest difference lies in the scores for the **OTHER** label, where all mBERT varieties perform better than XLM-R.

**Table 4.** Best values of precision (P), recall (R), and F1 score for the test set on baseline and STILTs models using the best validation parameters. Best scores in **bold**, second best F1 scores underlined.

Model	ENG			OTHER			Total		
	P	R	F1	P	R	F1	P	R	F1
mBERT	89.10	83.13	86.01	46.43	28.26	<u>35.14</u>	88.09	<b>81.17</b>	84.49
mBERT-1lang	88.54	81.68	84.97	50.00	15.22	23.33	88.07	79.30	83.46
mBERT-5lang	90.38	82.65	<u>86.34</u>	45.00	<b>39.13</b>	<b>41.86</b>	88.83	81.09	<u>84.78</u>
BETO	88.18	<b>83.70</b>	85.88	<b>66.67</b>	13.04	21.82	88.02	<b>81.17</b>	84.45
XLM-R	<b>91.47</b>	82.24	<b>86.61</b>	43.75	15.22	22.58	<b>90.8</b>	79.84	<b>84.97</b>
LinCE-LID <sub>STILT</sub>	<b>59.04</b>	15.82	<u>24.95</u>	33.33	6.52	10.91	<b>58.36</b>	15.49	<u>24.48</u>
XLM-R-NER <sub>STILT</sub>	53.02	<b>19.85</b>	<b>28.89</b>	40.00	4.35	7.84	52.88	<b>19.30</b>	<b>28.28</b>
LinCE-NER <sub>STILT</sub>	49.47	15.17	23.22	37.50	6.52	<u>11.11</u>	49.23	14.86	22.83
BETO-POS <sub>STILT</sub>	43.18	13.80	20.92	<b>60.00</b>	6.52	<b>11.76</b>	43.39	13.54	20.64
LinCE-POS <sub>STILT</sub>	58.54	15.50	24.51	30.00	6.52	10.71	57.69	15.18	24.03

Taking the best models when evaluated on the validation test, we evaluated the performance on the test set. As shown in Table 4 and Figure 1, results are similar, with mBERT-5lang performing slightly below XLM-R (84.78 vs 84.97), and presenting an F1 score for the **OTHER** label almost twice as that of the XLM-R.



**Fig. 1.** F1 scores for the different models on the validation and test sets per label and in total.

Given the differences between the number of labeled words in the validation and test sets for each label, we also obtained the best scores based solely on the test set. As seen in Table 5, the highest F1 scores are obtained by mBERT-5lang (87.08) and mBERT (86.33), more than doubling the F1 score for the **OTHER** label with respect to the score obtained by XLM-R. In general, the models achieving higher F1 scores for the **ENG** label also get higher scores in total.

## 6 Conclusions

In this work, we framed the detection of borrowings in Spanish as a token classification task. We hypothesized that supplementary learning could improve the performance of simply fine-tuned models. However, our results strongly suggest that supplementary learning might not be as effective for token classification of borrowings as it is for sequence classification on natural language understanding tasks. This results is inline with recent research. For example, [12] noted that supervised parsing might not be as availing as expected for high-level semantic natural language understanding. Given the very nature of the ADoBo task, in which only a small set of all words tagged as verbs or nouns actually occur to be borrowings, supplementary tasks that flag all parts of speech or all named entities might have a hard time trying to figure out exactly which of those are borrowings. In this sense, although not reported as part of the results, we found that when using language identification, with no further training, LinCE-LID<sub>STILT</sub> correctly assigns the **ENG** label with an F1 score of 44.31. It would be interesting as future work to check how much models trained on supplementary tasks unlearn once they are trained on borrowing detection. This could also be an indication that a hybrid approach merging both language models and handcrafted features useful

**Table 5.** Best values of precision (P), recall (R), and F1 score for the test set on baseline and STILTs models using the best test parameters. Best scores in **bold**, second best F1 scores underlined.

Model	ENG			OTHER			Total		
	P	R	F1	P	R	F1	P	R	F1
mBERT	91.10	84.26	<u>87.55</u>	70.00	30.43	<u>42.42</u>	90.74	82.33	<u>86.33</u>
mBERT-1lang	89.84	81.36	85.39	54.17	28.26	37.14	89.09	79.46	84.00
mBERT-5lang	90.59	<b>86.28</b>	<b>88.38</b>	59.26	<b>34.78</b>	<b>43.84</b>	89.89	<b>84.44</b>	<b>87.08</b>
BETO	90.34	84.58	87.37	47.06	17.39	25.40	89.72	82.18	85.78
XLM-R	<b>92.11</b>	81.92	86.72	<b>71.43</b>	10.87	18.87	<b>91.97</b>	79.38	85.21
LinCE-LID <sub>STILT</sub>	<b>55.65</b>	16.71	<u>25.70</u>	<b>50.00</b>	6.52	<u>11.54</u>	<b>55.56</b>	16.34	<u>25.26</u>
XLM-R-NER <sub>STILT</sub>	53.02	<b>19.85</b>	<b>28.89</b>	40.00	4.35	7.84	52.88	<b>19.30</b>	<b>28.28</b>
LinCE-NER <sub>STILT</sub>	53.16	16.95	25.70	42.86	6.52	11.32	52.99	16.58	25.25
BETO-POS <sub>STILT</sub>	47.47	13.64	21.19	37.50	6.52	11.11	47.25	13.39	20.86
LinCE-POS <sub>STILT</sub>	<b>55.65</b>	16.30	25.22	40.00	<b>8.70</b>	<b>14.29</b>	55.23	16.03	24.85

in language identification such as a character 3-grams could potentially boost the performance of the detection of borrowings. Additionally, adding a CRF or even a LSTM layer on top of the classifiers could see some performance gains for the simple finetuning. Auto-regressive models such as mT5 or GTP-J could also be leveraged for the task. A proper error analysis should be conducted to know exactly where and how the different approaches are failing.

We also believe that re-shuffling the train and validation sets could potentially solve the unbalance issue present in the labels of the corpus, improving the quality of the training data and making model selection more effective.

## References

1. Abdaoui, A., Pradel, C., Sigel, G.: Load What You Need: Smaller Versions of Multilingual BERT. In: SustaiNLP / EMNLP (2020)
2. Aguilar, G., Kar, S., Solorio, T.: LinCE: A Centralized Benchmark for Linguistic Code-switching Evaluation. In: Proceedings of The 12th Language Resources and Evaluation Conference. pp. 1803–1813. European Language Resources Association, Marseille, France (May 2020), <https://www.aclweb.org/anthology/2020.lrec-1.223>
3. Álvarez Mellado, E.: Lázaro: An Extractor of Emergent Anglicisms in Spanish Newswire. Master’s thesis (2020)
4. Álvarez Mellado, E., Espinosa Anke, L., Gonzalo Arroyo, J., Lignos, C., Porta Zamorano, J.: Overview of ADoBo 2021 shared task: Automatic Detection of Unassimilated Borrowings in the Spanish Press. *Procesamiento del Lenguaje Natural* **67** (2021)
5. Andersen, G.: Semi-automatic Approaches to Anglicism Detection in Norwegian Corpus Data. *The Anglicization of European lexis* **10**, 111 (2012)
6. Canete, J., Chaperon, G., Fuentes, R., Pérez, J.: Spanish pre-trained bert model and evaluation data. PML4DC at ICLR **2020** (2020)
7. Chesley, P., Baayen, R.H.: Predicting new words from newer words: Lexical borrowings in French (2010)

8. Derczynski, L., Nichols, E., van Erp, M., Limsopatham, N.: Results of the WNUT2017 shared task on Novel and Emerging Entity Recognition. In: Proceedings of the 3rd Workshop on Noisy User-generated Text. pp. 140–147 (2017)
9. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: 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). pp. 4171–4186. Association for Computational Linguistics, Minneapolis, Minnesota (Jun 2019). <https://doi.org/10.18653/v1/N19-1423>, <https://www.aclweb.org/anthology/N19-1423>
10. Furiassi, C., Hofland, K.: The Retrieval of False Anglicisms in Newspaper Texts. In: Corpus Linguistics 25 Years On, pp. 347–363. Brill Rodopi (2007)
11. Furiassi, C., Pulcini, V., González, F.R.: The Anglicization of European Lexis. John Benjamins Publishing (2012)
12. Glavaš, G., Vulić, I.: Is Supervised Syntactic Parsing Beneficial for Language Understanding Tasks? An Empirical Investigation. In: Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume. pp. 3090–3104. Association for Computational Linguistics, Online (Apr 2021), <https://www.aclweb.org/anthology/2021.eacl-main.270>
13. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: RoBERTa: A Robustly Optimized BERT Pretraining Approach. arXiv preprint arXiv:1907.11692 (2019)
14. Matras, Y.: Language Contact. Cambridge University Press (2020)
15. Nakayama, H.: sequeval: A Python Framework for Sequence Labeling e Evaluation (2018), <https://github.com/chakki-works/sequeval>, [Online; accessed 5-February-2021]
16. Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., Zettlemoyer, L.: Deep Contextualized Word Representations. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers). pp. 2227–2237 (2018)
17. Phang, J., Févry, T., Bowman, S.R.: Sentence Encoders on STILTs: Supplementary Training on Intermediate Labeled-Data Tasks. arXiv preprint arXiv:1811.01088 (2018)
18. Poplack, S., Sankoff, D., Miller, C.: The Social Correlates and Linguistic Processes of Lexical Borrowing and Assimilation (1988)
19. Poth, C., Pfeiffer, J., Rücklé, A., Gurevych, I.: What to Pre-Train on? Efficient Intermediate Task Selection (2021)
20. Serigos, J.R.L., et al.: Applying Corpus and Computational Methods to Loanword Research: New Approaches to Anglicisms in Spanish. Ph.D. thesis (2017)
21. Soto-Corominas, A., De la Rosa, J., Suárez, J.L.: What Loanwords Tell Us about Spanish (and Spain). Digital Studies/Le champ numérique 8(1) (2018)
22. Tjong Kim Sang, E.F.: Introduction to the CoNLL-2002 Shared Task: Language-Independent Named Entity Recognition. In: COLING-02: The 6th Conference on Natural Language Learning 2002 (CoNLL-2002) (2002), <https://www.aclweb.org/anthology/W02-2024>
23. Trask, L.: Why do Languages Change? Cambridge University Press (2009)
24. Vu, T., Wang, T., Munkhdalai, T., Sordani, A., Trischler, A., Mattarella-Micke, A., Maji, S., Iyyer, M.: Exploring and Predicting Transferability across NLP Tasks. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). Association for Computational Linguistics, Online (Nov



- 2020). <https://doi.org/10.18653/v1/2020.emnlp-main.635>, <https://www.aclweb.org/anthology/2020.emnlp-main.635>
25. Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., Bowman, S.: GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In: Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP. pp. 353–355 (2018)

## Availability

Source code for replicating the experiments in this paper are available in a code repository: <https://github.com/versae/adobo-eval>. Checkpoints for the best performing models are also released as PyTorch, Tensorflow, and JAX weights:

- mBERT-5lang: <https://huggingface.co/versae/mbert-5lang-adobo2021>
- XLM-R: <https://huggingface.co/versae/xlm-roberta-adobo2021>
- mBERT-5lang (on test): <https://huggingface.co/versae/mbert-5lang-test-adobo2021>
- mBERT (on test): <https://huggingface.co/versae/mbert-test-adobo2021>