

# Towards a Diagnostic and Predictive Evaluation Methodology for Sequence Labeling Tasks

Elena Álvarez-Mellado, Julio Gonzalo

NLP&IR research group, UNED  
{elena.alvarez, julio}@lsi.uned.es

## Abstract

Standard evaluation in NLP typically indicates that system A is better on average than system B, but it provides little info on how to improve performance and, what is worse, it should not come as a surprise if B ends up being better than A on outside data. We propose an evaluation methodology for sequence labeling tasks grounded in error analysis that provides both quantitative and qualitative information on where systems must be improved and predicts how models will perform on a different distribution. The key is to create test sets that, contrary to common practice, do not rely on gathering large amounts of real-world in-distribution scraped data, but consists in handcrafting a small set of linguistically motivated examples that exhaustively cover the range of span attributes (such as shape, length, casing, sentence position, etc.) a system may encounter in the wild. We demonstrate this methodology on a benchmark for anglicism identification in Spanish. Our methodology provides results that are *diagnostic* (because they help identify systematic weaknesses in performance), *actionable* (because they can inform which model is better suited for a given scenario) and *predictive*: our method predicts model performance on external datasets with a median correlation of 0.85.

**Keywords:** sequence labeling, evaluation, span identification, lexical borrowing, anglicisms

## 1. Introduction

Evaluation in NLP is based on comparing the aggregated scores systems obtain on a given test set for a given task. NLP evaluation, however, is plagued by a lingering issue: the fact that the score obtained by a system on a given test set does not tell us much about how the system can be improved, nor does it anticipate how the system will perform when tested on a different sample. Consequently, an excellent score on a given test set does not guarantee that the system will generalize well to a different unseen distribution (Arango et al., 2019). As a result, recent work has pointed out the need to dive deeper into the results produced by NLP models in order to get a better understanding of what models are capable and not capable of, and anticipate how models generalize to new data (Zhou et al., 2023) and perform on real-world scenarios (Reiter, 2025).

In this paper we propose a new methodology grounded in error analysis to create test sets for sequence labeling tasks. We apply it for the task of retrieving anglicisms from Spanish text and produce BLAS, a Benchmark for Loanwords and Anglicisms in Spanish. As a test set, BLAS is rather atypical. Instead of being a large collection of real-world in-distribution scraped data, BLAS is a small set of linguistically-motivated examples made from scratch and designed to exhaustively cover the range of span variability a model may encounter in the wild (in terms of shape, length, sentence position, casing, etc.) and systematically assess

the different types of errors a model can make. Our methodology produces results that are diagnostic, actionable and predictive: scores on BLAS not only serve as an evaluation benchmark to evaluate competing systems, but can also assist researchers identify which phenomena models are struggling with (thus diagnostic), which model is the best choice for a given scenario (thus actionable) and how the model will perform on a different distribution (thus predictive).

## 2. The benchmark paradigm and its discontents

In this section, we provide an overview of the current landscape of benchmark-centric evaluation in NLP and discuss several of its pitfalls.

### 2.1. The benchmark paradigm

Dataset creation is a crucial aspect of NLP research: the availability of annotated data guides the progress in the field and the way datasets are constructed even shapes the way we frame and approach tasks (Paullada et al., 2021).

Various factors may influence the dataset creation process (language, annotation resources, etc). More frequently than not, the availability of the data itself (the existence of previously compiled corpora or scrapeable texts, etc.) is a major factor in dataset creation and determines its content and the linguistic phenomena that gets represented in

it (Plank, 2016) and, in consequence, the type of phenomena that models will be evaluated on.

When a dataset for a given task is released, systems begin to be trained and evaluated on it. Certain datasets become the standard test set against which models are evaluated, thus becoming a benchmark for the task, such as CoNLL03 (Tjong Kim Sang and De Meulder, 2003) or OntoNotes (Pradhan et al., 2013) for NER. Making progress on that given task now implies obtaining better scores on a particular test set. In other words, solving a general task may now be restricted to making systems capable of modeling the phenomena that appear in one specific test set. This means that the phenomena that our modeling efforts will consider will be heavily influenced by the linguistic attributes that our text sample (the dataset) contained to begin with.

After reaching a certain degree of progress, improvement over scores may become elusive: more work may be required to produce increasingly smaller improvements. Scores can eventually plateau, thus producing a glass ceiling situation (Stanislawek et al., 2019). Researchers and practitioners try to grasp what there is left to improve for that given task and look at the errors that their systems are making on the test set in order to identify the type of examples their models are failing at (Rueda et al., 2024). With the development of new techniques, significant improvement on the scores may finally be reached. If the scores are high enough so that there is little room for improvement, the dataset can be considered saturated and sometimes the task will even be claimed to be solved (Kiela et al., 2021).

Eventually, a mismatch between the optimistic scores obtained by models in lab conditions over benchmark data and the results obtained by those same models on real-world scenarios will be identified (Lin et al., 2020; Reiter, 2025). As a result, patches will be explored to account for the type of naturalistic data that models are struggling with (Heigold et al., 2018; Ribeiro et al., 2018; Namysl et al., 2021). Finally, a new and more challenging dataset will be released for the task, and the process starts again.

## 2.2. Its discontents

The Sisyphean approach we have just described, although standard, raises the following issues.

**The limitations of identically distributed evaluation.** Under this paradigm models are trained and tested over different sets that, although assumed to be independent, are similarly sampled and identically distributed (Hupkes et al., 2023). Evaluation over an identically-distributed sample

may be suitable in certain restricted or predictable scenarios, but it has been shown that its results can overestimate performance and lead to poor generalization (Gorman and Bedrick, 2019).

Moreover, calling a task “solved” because outstanding scores were obtained on a split of a dataset where most of the data is alike and where the presence of worst-case scenarios or linguistic variations of the phenomenon in question is not ensured may be too optimistic. Unsurprisingly, minor perturbations that pose no challenge for humans (e.g., casing or out-of-domain topics) can drastically degrade performance on allegedly solved tasks (Gururangan et al., 2018; Mayhew et al., 2019), highlighting the lack of robustness of these models (Ma et al., 2023).

**The limits of adversarial test sets.** As a result, some voices have advocated for building adversarial test sets, that is, populating test sets with the most challenging cases possible (Søgaard et al., 2021). As useful as they may be, adversarial test sets are not a silver bullet. It is indeed important to know that a certain model succeeds on the worst case scenarios. But evaluating exclusively on these types of extreme cases will prevent us from actually learning what a model can do: if tested on challenging cases that fool most models, a model that is good but far from perfect may perform as poorly as a dummy baseline. An evaluation that aims to be useful and actionable should be informative about the models’ capabilities, as well as granular enough that it can help us distinguish between a bad model, a mediocre model, a good model and a great model, something that exclusively evaluating on the most difficult cases cannot guarantee (Rodriguez et al., 2021).

**The curse of holistic scores.** Even if the test set is balanced in terms of the difficulty of the examples it contains, model comparison based on a score over the whole test set may be short-sighted too, because by doing so we are losing crucial information: the type of data that the model is succeeding at. Let’s say that we have a task for which state-of-the-art models achieve a recall of 80%. A model that only scores 30% will be deemed mediocre and probably irrelevant. But what if the data that the apparently mediocre model is succeeding at includes precisely the 20% of data that SOTA models are failing at? That would be a remarkable achievement that could illuminate the areas of research future work should concentrate on, but we have no way of knowing that because in score-based evaluation is the overall score that counts, not the section of data it was achieved on.

**Beyond error analysis: the need for success analysis.** In order to better contextualize the information that is portrayed by numeric scores only, qualitative analysis are usually performed, generally through error analysis. Error analysis usually consists in identifying the type of examples where models are failing and consequently establishing areas that require further work. Although it may seem counterintuitive, focusing exclusively on errors as a way of deciding what areas of modeling require improvement may not be the best solution either: error analysis can only be a productive approach towards the improvement of models if it is accompanied by a *success analysis* that contextualizes how prevalent certain phenomena are (Wu et al., 2019). Otherwise we risk the possibility of overstating the importance of some errors and minimizing others, which may lead to misguided decisions.

In a nutshell, it is a wide-known (yet unsolved) fact in NLP that standard evaluation based on providing a holistic score averaged over a test set (that, more often than not, will usually be identically distributed to the data contained in the training set) does not tell us much about what the model can and cannot do, which areas of the problem require further work, or how the model will generalize to a different distribution.

### 3. Our methodology: scope and overview

Inspired by our observations from Section 2, we propose a methodology for creating test sets in which examples are selected based on their linguistic properties. This way of proceeding radically diverges from the usual data-centric approach, in which the readiness of data (and not its linguistic characteristics) guides the data selection process.

The scope of our methodology is restricted to span identification tasks. Span identification tasks are a type of sequence labeling tasks in which relevant spans are retrieved from text. Named entity recognition (NER) or multiword expression detection (MWE) are prime examples of span identification tasks. In span identification tasks, certain formal attributes of the span (such as its length or its distinctiveness) or of the context it appears in will make the task of retrieving the span more or less challenging (Papay et al., 2020). However, none of these issues are taken into account in the standard evaluation of sequence labeling models.

We propose to populate sequence labeling test sets with examples that are selected based on the linguistic attributes of the spans they contain (in terms of shape, length, casing, position within the sentence, etc.), so that scores can be computed over examples that share the same attributes, thus

producing a result that is ascribable to concrete linguistic attributes.

The final aim of our methodology is that, whenever we see a sequence labeling model failing at retrieving a given span, we can answer the following questions: Would that model have succeeded if the same span had appeared in a more obvious context? Would it have been capable of retrieving a more prototypical span from the same context instead? More generally: Which types of spans is our model finding easier to detect and which ones is it consistently failing to retrieve? Which contexts or formal characteristics are serving as cue to the model? How robust is our model when the cue is not present?

## 4. Implementing the methodology

We will demonstrate our methodology by creating BLAS, a sequence labeling test-only dataset for anglicism identification in Spanish.

### 4.1. Anglicism detection

Anglicism detection is the task of retrieving English lexical borrowings (or *anglicisms*) from non-English texts. Anglicisms can be single-item (*app*) or multiword (*machine learning*, *fake news*). The task of automatically retrieving lexical borrowings from text has proven to be useful both for lexicographic purposes and for NLP downstream tasks in various languages (Furiassi and Hofland, 2007; Alex, 2008; Andersen, 2012; Losnegaard and Lyse, 2012; Tsvetkov et al., 2015; Serigos, 2017) and has previously been framed as a sequence labeling task (Álvarez-Mellado and Lignos, 2022; Chiruzzo et al., 2023), in which relevant in-context spans of text are retrieved from sentences (in a similar fashion to how NER or MWE are approached).

Just like in NER or MWE, in anglicism detection the retrievability of the span is heavily influenced by its shape and the context it appears in. Span attributes such as length, ambiguity, casing, surrounding context or sentence position can make the task of retrieving a given span easier or harder. These attributes, however, are not equally represented in available datasets, as most spans tend to share the same type of formal characteristics (Álvarez-Mellado and Gonzalo, 2024). Our aim with BLAS is to produce a test set where different combinations of span types and contexts are well represented and the characteristics that can lead to a model missing a span are exhaustively explored.

### 4.2. Span attribute selection

The first step of our methodology for the data selection process is to define which span attributes

should models be evaluated on (and thus be represented in our dataset). For BLAS, we decided to focus on the following attributes, which previous work on error analysis had identified as relevant for anglicism identification systems (Álvarez-Mellado and Gonzalo, 2024):

- Span length: being single-item or multiword.
- Span shape: complying with the graphotactic expectation of the recipient language or not<sup>1</sup>.
- Sentence position: appearing mid sentence or in sentence initial position.
- Quotations: being surrounded by quotation marks or not.<sup>2</sup>
- Casing: being lowercase, titlecase, etc.
- Adjacency: being collocated to another span.
- Ambiguity: the same sequence can be seen labeled as a span or not, depending on the context.<sup>3</sup>

These will be the linguistic attributes that our test set will evaluate.

### 4.3. Data creation

We populated BLAS with sentences that exhaustively cover the different combinations of the span attributes we listed in Section 4.2 (for example, non-ambiguous multiword, graphotactically-compliant uppercase spans that appear in sentence initial position between quotation marks, etc.). Finding examples from real sources that covered all combination of attributes would have been an unmanageable endeavor, as anglicisms are a sparse phenomenon that amount to 1%-2% of the language (Rodríguez González, 2002).

Consequently, instead of populating our benchmark with real-world examples, we decided to have such examples written from scratch by a linguist (see Appendix A for guidelines on the writing process). This approach, although unconventional, offered several benefits. First, instead of being confined to evaluating systems on the data that is available on the web, we gained full control of the

<sup>1</sup>Whether the anglicism violates the spelling rules of the recipient language. For instance, the anglicism *streaming* does not comply with the graphotactical expectations that Spanish speakers have (because words in Spanish are not supposed to begin with *str-* or end in *-ing*), but *online* does.

<sup>2</sup>According to orthographical rules of Spanish, lexical borrowings such as anglicisms should be written between quotation marks or in italics.

<sup>3</sup>For example, the word *pie* will be an anglicism when talking about dessert (as in *un pie de limón*) but it will be a native Spanish word when it means “foot”.

linguistic attributes that our dataset would cover. Creating our examples from scratch also avoided copyright issues that usually prevent datasets from being freely reshared. Finally, the fact that the sentences were new and not gathered from any other existing sources ensured that no model had previously been exposed to the data contained in the dataset, thus avoiding data contamination issues (Sainz et al., 2023).

The mechanics behind the creation of BLAS consistently followed the same procedure for all examples: (1) identifying the span attributes we wanted to evaluate (being graphotactically compliant, being ambiguous, being multiword, etc.) (see Section 4.2); (2) curating a minimal but varied list of ten spans that satisfied those attributes (see Appendix A.1); (3) writing sentence examples for each of the selected spans in different sentence positions (see Table 1 and Appendix A.2); and (4) the original sentences underwent a series of systematic casing and punctuation perturbations (all text goes into lowercase, all quotation marks are removed, etc.) (see Table 2 and Appendix A.3). The purpose of these perturbations is to explore all possible punctuation configurations in which a sentence and a span can appear and to control, on top of the span attribute selection, for models’ overreliance on certain orthotypographic cues such as casing and quotation mark presence, which have been shown to affect model performance in sequence labeling tasks (Mayhew et al., 2019; Álvarez-Mellado and Gonzalo, 2024).

### 4.4. Benchmark description

The result of the process described in Section 4.3 is BLAS (Benchmark of Loanwords and Anglicisms in Spanish), a collection of 1,836 annotated sentences in Spanish (37,344 tokens) that contains 2,076 spans labeled as anglicisms. Every sentence in BLAS contains one span labeled as anglicism, except for the sentences with adjacent spans, which contain two.

Each sentence in BLAS is characterized by the linguistic attributes of the span it contains (in terms of shape, the context it appears in, its position within the sentence, its casing, etc.). When models are evaluated on these sentences, scores can be computed over subsets of the data that share the same span attributes, which enables ascribing the resulting score to concrete linguistic attributes.

In order to assess its validity, our dataset was provided to a second annotator for double annotation. The Cohen’s kappa IAA at token level was 0.98 (0.99 of pairwise F1 score at span level), which indicates high reliability (Artstein and Poesio, 2008).

Type	Length	Position	# sentences	Example
Compliant	Single-word	Initial	10	<i>Burpees para perder kilos sin salir de casa.</i>
		Mid	10	<i>Los burpees son efectivos para perder peso.</i>
	Multiword	Initial	10	<i>Medal race entre Nigeria y Polonia por la plata.</i>
		Mid	10	<i>Ambos países han llegado a la medal race ajustados de puntos.</i>
Non-compliant	Single-word	Initial	10	<i>Spoilers del episodio siete a continuación.</i>
		Mid	10	<i>Las redes amanecieron con mensajes llenos de spoilers del último capítulo.</i>
	Multiword	Initial	10	<i>Fact checkers confirman que la cifra aportada por el ministerio no es correcta.</i>
		Mid	10	<i>Los fact checkers contrastarán los datos durante el debate.</i>
Mixed compliant	Multiword	Initial	10	<i>Joint ventures de todo el mundo se reúnen en la mayor feria mundial de la industria.</i>
		Mid	10	<i>La fusión de ambas compañías supone un hito en la historia de las joint ventures.</i>
Ambiguous	Single-word	Mid	3	<i>Receta de pie de limón.</i>
	Multiword	Mid	10	<i>La reina escogió un conjunto total red para el evento.</i>
Mixed ambiguous	Multiword	Initial	10	<i>Casual looks con bufanda para esta temporada.</i>
		Mid	10	<i>Los vestidos dejan hueco a casual looks más alegres y desenfadados.</i>
Adjacent	Single-word	Mid	10	<i>La agencia se especializa en campañas de marketing online.</i>
	Multiword	Mid	10	<i>Ahora trabaja como head hunter full time.</i>

Table 1: Number of sentences in BLAS per span type.

Quotation marks	Casing	Example
No	Standard	<i>Los burpees son efectivos para perder peso.</i>
No	All words are lowercase	<i>los burpees son efectivos para perder peso.</i>
No	All words are uppercase	<i>LOS BURPEES SON EFECTIVOS PARA PERDER PESO.</i>
No	All words are capitalized	<i>Los Burpees Son Efectivos Para Perder Peso.</i>
No	Borrowings are uppercase	<i>Los BURPEES son efectivos para perder peso.</i>
No	Borrowings are capitalized	<i>Los Burpees son efectivos para perder peso.</i>
Yes	Standard	<i>Los "burpees" son efectivos para perder peso.</i>
Yes	All words are lowercase	<i>los "burpees" son efectivo para perder peso.</i>
Yes	All words are uppercase	<i>LOS "BURPEES" SON EFECTIVOS PARA PERDER PESO.</i>
Yes	All words are capitalized	<i>Los "Burpees" Son Efectivos Para Perder Peso.</i>
Yes	Borrowings are uppercase	<i>Los "BURPEES" son efectivos para perder peso.</i>
Yes	Borrowings are capitalized	<i>Los "Burpees" son efectivos para perder peso.</i>

Table 2: Casing and quotation mark perturbations considered in BLAS.

## 5. Experiments & results

### 5.1. Experiments across models

The BLAS test set that resulted from Section 4 was used to evaluate six baseline models that had previously been used for anglicism retrieval (Álvarez Mellado, 2025): five already-available supervised models that had been trained on the COALAS dataset from Álvarez-Mellado and Lignos (2022) (a CRF, Transformer-based BETO and mBERT, and two BiLSTM-CRFs with different types of word and subword embeddings), and one additional large language model on a few-shot approach (8B-Llama3, see Appendix B). As BLAS is a test-only dataset, the models were not trained on any split of BLAS, but only evaluated on it.

As all sentences in BLAS are positive examples (they all contain at least one span) and anglicisms are a sparse phenomenon, the appropriate diagnostic metric to analyze results is recall. As we will see in Section 5, recall will also serve as a reliable proxy to F1 score performance.

#### 5.1.1. Overall performance

Table 3 displays overall recall scores obtained by the six models on BLAS. Scores range between 6% and 36%. These scores were obtained over the whole collection of 1,836 sentences in BLAS, which include all span types with all punctuation transformations.

These results are far below the scores reported by Álvarez-Mellado and Lignos (2022) on COALAS test set (the previously available dataset for the task), where the best performing model obtained an average recall of 78. This gap in recall shows cases that models that learned in the supervised fashion from the biases and characteristics of the training dataset struggle to generalize to outside data that is substantially different in terms of linguistic characteristics (even when the language and genre of the text is the same). This shows that, although small in quantity, a careful selection of examples such as the one we performed in BLAS can serve to evaluate aspects and identify weaknesses in performance that go unnoticed when evaluating

	BLAS		COALAS (2022)	
CRF	6.50	#6	43.04	#5
BETO	23.55	#4	75.50	#4
mBERT	23.80	#2	76.16	#3
BiLSTM-CRF (unadapted)	23.55	#4	78.34	#2
BiLSTM-CRF (codeswitch)	23.75	#3	<b>78.72</b>	#1
Llama3	<b>36.37</b>	#1	32.14	#6

Table 3: Span recall per model on BLAS and COALAS.

on large amounts of in-distribution data.

On the other side, the best-performing model on BLAS is Llama3 (R=36.37), well above all the supervised models. This is remarkable, as Llama3 ranked last when tested on COALAS, even below the CRF. These results suggest that LLMs on few-shot prompting are more robust and better equipped than traditional supervised models when dealing with unseen data that is substantially different to the training data available. If we take BLAS as an exhaustive collection of examples for evaluating the ability of models to retrieve anglicisms from a variety of linguistic shapes and contexts, then we can conclude that anglicism retrieval is far from being a solved task in Spanish, at least for the models we have explored.

### 5.1.2. Performance across punctuation

Holistic results such as the ones reported on Table 3, however, do not tell us much about the type of examples models are failing to retrieve, or how to fix them. Table 4 displays the same results than Table 3, but split into the different casing and quotation mark configurations of BLAS.

Two facts stand out from Table 4. First, all models obtained better results on the versions of the dataset where the spans were written between quotation marks. The differences in recall between examples with quotation marks and without them are substantial, with scores dropping even around 20 points or more for some models. Our results reveal that Llama3 was by far the most vulnerable to quotation mark absence, with a loss of more than 40 points in recall when quotation marks were not present on the standard casing split. Note that these differences in scores are fully attributable to the absence of quotation marks, because both splits contain exactly the same sentences, but in one split the spans are surrounded by quotation marks, while in the other they are not. In other words, the very same models that were capable of correctly identifying *burpees* in *Los "burpees" son efectivos para perder peso* failed to do so when the quotation marks were absent (*Los burpees son efectivos para perder peso*).

Second, supervised models consistently obtained catastrophic results on the capitalized sec-

tions of BLAS, with models scoring zero across uppercase configurations. Llama3, on the other hand, while it still saw a drop in performance over the capitalized sections of BLAS (especially when the whole text was affected), produced *less* bad results than the supervised models, which explains its good position in the overall ranking from Table 3.

The results from Table 4 provide us with actionable conclusions that can inform decision-making when choosing which system to use for anglicism detection on different scenarios: if we are expecting well-edited text with standard casing and quotation marks, then BETO would probably be the best model. If quotation marks are not expected, then mBERT would be a better call. This casing vulnerability points to an issue that future work should address: none of these models is reliable when applied to texts with non-standard capitalization. However, if we had to choose one model for such scenario, Llama3 would be our safest bet. If we are expecting lowercase text (such as social media text), then the BiLSTMs are the best choice.

### 5.1.3. Performance across different span types

We now analyze performance across different span types. We will focus on the standard casing configuration (where only proper names and sentence-initial words are capitalized), as it is the most canonical and results over the capitalized sections of BLAS were catastrophic across most models.

Table 5 shows that certain span attributes produce better results than others, like mid-sentence spans. In fact, non-compliant spans that appear in mid-sentence position with quotation marks (such as *spoilers* in *Las redes amancierón con mensajes llenos de "spoilers" del último capítulo*) can essentially be considered solved, with almost all models producing perfect or near perfect scores over those examples.

On the other hand, performance drops substantially when retrieving spans from sentence initial position (such as *spoilers* in *"Spoilers" del episodio siete a continuación*). It is remarkable that the very same spans that were retrieved with a score of 80, 90 or even 100 when they appeared mid-sentence, produced a recall of 20, 10 or even 0 when they appeared at the beginning of the sentence. Again, it should be noted these differences in recall between sentence-initial and mid-sentence spans can be fully ascribed to the differences in the contexts they appear in, as the spans were the same across both groups (see Table 1).

Ambiguous spans also produced substantially worse results than other types of spans, although Llama3 proved to be more robust to ambiguity, as long as quotation marks were present. On the other hand, adjacency was not a challenge for

	CRF		BETO		mBERT		BiLSTM (unadp)		BiLSTM (codeswitch)		Llama3	
	With quot.	W/o quot.	With quot.	W/o quot.	With quot.	W/o quot.	With quot.	W/o quot.	With quot.	W/o quot.	With quot.	W/o quot.
Standard casing	21.97	7.51	<b>68.79</b>	45.09	61.27	<b>49.13</b>	66.47	45.09	60.12	46.24	66.47	23.12
Text is lowercase	31.79	16.18	87.86	56.65	83.82	65.32	<b>90.17</b>	61.27	80.92	<b>68.21</b>	67.63	27.75
Span is titlecase	0.00	0.00	5.20	5.20	4.62	5.78	6.36	5.78	5.78	6.36	<b>56.65</b>	<b>24.86</b>
Text is titlecase	0.00	0.00	7.51	5.78	10.40	5.20	4.05	2.89	6.94	6.94	<b>32.95</b>	<b>9.25</b>
Span is uppercase	0.00	0.58	0.00	0.58	0.00	0.00	0.00	0.00	1.16	2.31	<b>54.91</b>	<b>45.09</b>
Text is uppercase	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.58	0.00	0.00	<b>21.97</b>	<b>5.78</b>

Table 4: Overall span recall on BLAS per model on different casing and quotation mark configurations.

Type	Length	Position	CRF		BETO		mBERT		BiLSTM (unad)		BiLSTM (codeswitch)		Llama3	
			With quot.	W/o quot.	With quot.	W/o quot.	With quot.	W/o quot.	With quot.	W/o quot.	With quot.	W/o quot.	With quot.	W/o quot.
compliant	multi	ini	0.00	0.00	40.00	<b>50.00</b>	20.00	40.00	20.00	40.00	30.00	30.00	<b>90.00</b>	20.00
		mid	20.00	0.00	<b>100.00</b>	80.00	90.00	80.00	<b>100.00</b>	80.00	90.00	<b>90.00</b>	80.00	30.00
	single	ini	0.00	0.00	10.00	<b>20.00</b>	0.00	10.00	20.00	<b>20.00</b>	10.00	<b>20.00</b>	<b>50.00</b>	10.00
		mid	10.00	0.00	70.00	50.00	80.00	80.00	90.00	80.00	<b>100.00</b>	<b>90.00</b>	40.00	10.00
non compliant	multi	ini	10.00	10.00	<b>90.00</b>	50.00	60.00	<b>70.00</b>	30.00	40.00	20.00	<b>70.00</b>	<b>90.00</b>	30.00
		mid	60.00	10.00	<b>100.00</b>	80.00	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	90.00	40.00
	single	ini	0.00	0.00	20.00	20.00	20.00	30.00	20.00	10.00	10.00	<b>40.00</b>	<b>80.00</b>	<b>40.00</b>
		mid	50.00	40.00	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	90.00	<b>100.00</b>	90.00	<b>100.00</b>	<b>100.00</b>	60.00	50.00
mixed compliant	multi	ini	0.00	0.00	30.00	<b>40.00</b>	20.00	30.00	20.00	30.00	20.00	20.00	<b>80.00</b>	20.00
		mid	10.00	0.00	<b>100.00</b>	90.00	90.00	90.00	<b>100.00</b>	<b>100.00</b>	90.00	80.00	70.00	30.00
ambiguous	multi	mid	0.00	0.00	50.00	30.00	50.00	<b>40.00</b>	<b>70.00</b>	20.00	50.00	10.00	<b>70.00</b>	10.00
		mid	0.00	0.00	<b>33.33</b>	<b>33.33</b>	<b>33.33</b>	0.00	<b>33.33</b>	0.00	0.00	0.00	<b>33.33</b>	0.00
mixed ambiguous	multi	ini	0.00	0.00	30.00	<b>40.00</b>	30.00	30.00	20.00	20.00	0.00	0.00	<b>90.00</b>	20.00
		mid	10.00	0.00	70.00	50.00	70.00	50.00	<b>100.00</b>	<b>60.00</b>	60.00	50.00	80.00	40.00
adjacent	multi	mid	45.00	10.00	<b>90.00</b>	20.00	65.00	<b>30.00</b>	80.00	5.00	80.00	15.00	55.00	10.00
		mid	60.00	25.00	95.00	15.00	95.00	25.00	95.00	<b>40.00</b>	<b>100.00</b>	35.00	30.00	15.00

Table 5: Span recall obtained by each model over BLAS over different types of spans and contexts (on standard casing, with and without quotation marks).

supervised models: most of them produced scores over 80, at least with quotation marks. Llama3, on the contrary, obtained results on adjacent spans that were in line or even below the CRF.

## 5.2. Predictive ability of BLAS

The natural question that follows from our results from Section 5.1 is: are the results per attribute obtained on BLAS representative of the ability of the models to perform on those linguistic attributes, not just over BLAS data, but in general? Could we use the scores obtained on BLAS to predict the overall recall that those same models would obtain on a different distribution? For instance, based on the results obtained on BLAS, we know that Llama3 produces a span recall of 90 over non-compliant multiword spans when they appear mid sentence and are written with quotation marks on standard casing. Could we extrapolate those results obtained by Llama3 on BLAS to other datasets and claim that that Llama3 will successfully retrieve 90% of the spans that share those same characteristics when tested on a different distribution?

In order to answer these questions, we ran the six models from Section 5.1 on five external datasets: the test and dev sets from COALAS (Álvarez-Mellado and Lignos, 2022) and the dev set and test sets A and B from CALCS (Álvarez-Mellado, 2020). These are all publicly available datasets for anglicism identification in Spanish that, although similar in genre to BLAS (Spanish journalistic text), show a very different distribution in terms of the linguistic attributes represented (Álvarez-

Mellado and Gonzalo, 2024). In addition, COALAS test set and CALCS test set B have been reported to be substantially different (in terms of out-of-distribution topics and vocabulary) to their training splits.

The purpose of our experiment was to test if the scores the models obtained on BLAS across different span attributes could predict the results those same models would obtain on these external datasets. To do so, we counted the number of anglicism spans per type in each of these external datasets and calculated the expected number of spans that would be successfully retrieved by simply extrapolating from the recall results obtained per type by each model on BLAS. Going back to our example above: there are 20 spans in the COALAS test set that are non-compliant multiword and that appear with standard casing, quotation marks and in mid sentence position. Given that the span recall of Llama3 obtained on BLAS on that particular type of spans is 90, we hypothesized that 90% of those 20 spans in COALAS would be successfully retrieved by Llama3, which amount to 18 true positives and 2 false negatives. We did the same for each type of span in all five datasets for every model, and calculated the expected overall span recall per model on each dataset based on BLAS scores (see Appendix C).

Table 6 compares the true recall and ranking position and the predicted recall and ranking position derived from extrapolating the results produced on BLAS. Results show that our method accurately predicts both recall values (with a median Pearson correlation of 0.85) and system ranking (median

Models	BLAS	COALAS (test set)				COALAS (dev set)				CALCS (test set A)				CALCS (test set B)				CALCS (dev set)			
	R	Predict R		True R		Predict R		True R		Predict R		True R		Predict R		True R		Predict R		True R	
CRF	6.50	26.58	#6	44.31	#5	33.13	#6	68.63	#5	29.00	#6	65.57	#5	24.76	#6	53.17	#5	28.99	#6	64.84	#5
BETO	23.55	86.87	#3	77.99	#4	87.62	#3	84.05	#4	83.58	#4	83.49	#3	79.13	#4	81.75	#4	81.60	#4	86.30	#1
mBERT	23.80	86.86	#4	78.85	#3	87.10	#4	84.31	#3	83.86	#3	85.85	#1	83.17	#3	84.13	#3	82.73	#3	86.30	#1
BiLSTM*	23.55	87.86	#2	80.88	#2	88.73	#2	89.05	#1	84.85	#2	82.08	#4	84.04	#2	84.92	#2	83.06	#2	83.11	#4
BiLSTM**	23.75	92.96	#1	85.76	#1	91.08	#1	86.57	#2	89.43	#1	83.96	#2	89.52	#1	85.71	#1	87.53	#1	85.84	#3
Llama3	36.37	42.93	#5	33.33	#6	55.46	#5	53.27	#6	48.53	#5	53.30	#6	43.17	#5	40.48	#6	48.49	#5	50.68	#6
Correl				0.94	0.89			0.79	0.83			0.85	0.66			0.91	0.94			0.84	0.41

Table 6: Correlation between the predicted and true recall scores and ranking over five datasets (Pearson and Spearman correlation coefficients).

correlation of 0.83). The highest correlations were obtained on COALAS test set and CALCS test set B (0.94 and 0.91, respectively), precisely the two splits that had been reported to be most dissimilar to their training splits. Our predicted rankings over recall also display a strong correlation with F1 score rankings (median correlation 0.83). It should be noted that overall scores on BLAS are not predictive themselves: it is the projection of the system performance across span attributes in BLAS that correlates with true scores (see Tables 3 and 6).

## 6. Discussion

Our experiments show that BLAS can uncover systematic weaknesses in anglicism detection models that may go unnoticed when evaluating on naturalistic in-distribution data. The fine-grained results produced on BLAS can also inform decision-making when choosing which is the best model for a given data scenario (punctuation, capitalization, types of spans, etc). In addition to serving as a test set to evaluate models and identify systematic weaknesses, our results show that BLAS has predictive ability: the scores per attribute obtained by six models on BLAS were able to accurately predict the performance of those models on five external test sets with a median correlation of 0.85.

Although we have built our case on anglicism detection, our methodology could in principle be applied to other span identification tasks, such as NER, where ambiguous and adjacent spans are ubiquitous and span retrievability is also sensitive to context, punctuation and sentence position.

The graphotactic compliance distinction may be specific to anglicism identification, but this linguistic attribute is simply a concrete instantiation of a more general issue that is common to other sequence labeling tasks: the fact that in any span-based task, there will be some spans that will be more salient than others, a feature that has also been reported to be key in NER (Papay et al., 2020; Lin et al., 2021; Vajjala and Balasubramaniam, 2022).

The predictive ability displayed by BLAS fine-grained results suggests that linguistically-aided test set creation informed by error analysis may be

a workable solution to the long quest for evaluation methodologies that can anticipate how a system will perform on new data, a feature that is considered key to measure the generalization abilities of NLP systems (Hupkes et al., 2023; Zhou et al., 2023; Reiter, 2025).

## 7. Relation to previous work

The ideas that guided the creation of BLAS take inspiration from previous work on robustness evaluation in NLP, especially on NER models.

The gap in performance between the results obtained by NLP models when they are evaluated on benchmark data and when they are tested on real-world scenarios has previously been discussed (e.g. Arango et al., 2019; Kiela et al., 2021; Reiter, 2025).

Lin et al. (2020) referred to this gap as the difference in performance between *regular NER* (NER performed in lab conditions on benchmark data) and *open NER* (NER performed on real world data) and pointed out three possible reasons behind this gap: strong name regularity in NER benchmark entities (i.e., most spans have a similar shape), overlap between the entities in the test set and the training set and test instances that are rich in context.

Additional research on NER has added two more possible culprits for the lack of robustness of NER models: label inconsistency (Fu et al., 2020; Bernier-Colborne and Langlais, 2020; Tu and Lignos, 2021) and presence of adjacent spans (Rueda et al., 2024). These phenomena that plague NER performance are no different to the type of issues that we encounter in anglicism detection. With BLAS, we have translated these issues (span regularity, ambiguity, context and adjacency) to the task of anglicism detection by selecting test examples in which these attributes are exhaustively represented.

The idea of applying systematic perturbations to inputs in order to assess models' robustness is not new and has previously been applied to different aspects of NER performance, particularly for casing (Mayhew et al., 2019), span shape (Agarwal et al., 2021; Lin et al., 2021; Vajjala and Balasubra-

maniam, 2022), type or length (Wang et al., 2021). It also builds on previous work on stress test evaluation of models' robustness on specific linguistic phenomena (Naik et al., 2018; Wang et al., 2019; Ribeiro et al., 2020; Fu et al., 2020), but specifically targeting sequence labeling tasks. Our motivation is similar to the work of Kovatchev and Lease (2024), but they opt for describing dataset items in terms of extrinsic features (how a set of systems perform on each of the instances), while our dimensions are intrinsic and linguistically motivated: they do not depend on the set of systems tested.

The idea of building a testbed by creating linguistically-motivated examples by hand draws from other expert-authored benchmarks from the NLU literature (Bowman and Dahl, 2021), such as the FraCaS test suite (Cooper et al., 1996) or the the Winograd Schema Challenge (Levesque et al., 2012). It also follows the recommendations stated in Linzen (2020) of producing test-only benchmarks to avoid unintended statistical regularities, and those from Søgaard et al. (2021) that proposes evaluating models on biased splits as evaluation protocol to determine data characteristics that affect performance. Similarly, our methodology implements and systematizes the spirit behind other evaluation campaigns such as Kiela et al. (2021) and Ettinger et al. (2017), which put the linguistic expertise at the center of NLP evaluation.

## 8. Conclusions

We introduced a new methodology grounded in error analysis to evaluate sequence labeling tasks, in which test sets are populated with examples selected based on their linguistic attributes and sampled according to the types of errors a system may make. We applied this methodology to create BLAS, a test-only dataset for anglicism identification in Spanish. Results on BLAS are more informative than scores on standard test sets, as they diagnose systematic weaknesses in performance, identify which systems work best in a given scenario, and predict how a system will generalize to new data.

## 9. Limitations

**The task** We have demonstrated our methodology on one task only: anglicism identification in Spanish. Although our task is quite niche, we believe that the same methodology could be applied to other span identification tasks, such as NER and MWE, where span retrieval is affected by the very same phenomena we have explored (ambiguity, casing, sentence position, span length, span shape, etc.). Future work should examine error analysis on these other tasks and consider

whether additional attributes or modifications of the current ones are needed.

**The type of problem** The methodology we have presented can only be applied to span identification tasks, where spans are retrieved from in-context sentences. It is not clear how a similar methodology could be devised for text classification problems, or what attributes should be taken into account.

**The metrics** Our work has focused on evaluating the capability of models at retrieving true spans of different types and shapes from different contexts. Consequently, our analysis and the creation of the dataset itself focused on recall. Prioritizing recall is suitable for tasks in which minimizing the number of false negatives is crucial (even if it is at the cost of having some false positives), such as when building technology to extract linguistic phenomena that is rare (Kermes, 2004), like lexical borrowing.

Other tasks, however, may require taking precision into account and may wish to perform a similar evaluation over false positives. To evaluate precision in a similar way as we did with recall with BLAS, we would first need to investigate the type of linguistic phenomena that characterize false positive spans and then populate our test set with non-relevant spans that could pass for one.

It should be noted that predicting a precision score over an external dataset as we did with recall with BLAS would not be feasible with our methodology: in a span-based task, the sum of false negatives and true positives is known, as it is the number of total spans in the goldstandard. Therefore, we can calculate the proportion of spans that are expected to be missed by a model, based on the recall scores it obtained on the benchmark. However, if we tried to do the same with precision, even if we had a precision-based benchmark, it would be unclear what the set of potential false positive spans of the external test set would be.

**The authorship** Sentences in BLAS were authored by a single linguist. Although extensive guidelines were compiled to ensure reproducibility of our method and the double annotation yielded very high IAA, this fact may introduce biases in the data. Further work is required to assess for generalization if a larger multi-authored test set was to be created.

**The dataset** BLAS should not be taken as a definite benchmark for anglicism identification in Spanish, nor as the sole test set against a model should be measured. Quite on the contrary, BLAS is intended as a very specific diagnosing testbed that may contribute documenting and assessing a

model's capability. In that regard, BLAS follows the trend of other auxiliary validation sets (Ghaddar et al., 2021). How relevant the scores obtained by a model on BLAS are will be determined by the use case scenario where the model in question is expected to be deployed.

## 10. Data availability

BLAS was the test set used in the ADoBo<sup>4</sup> shared task on automatic detection of anglicisms at IberLEF 2025 (Álvarez-Mellado et al., 2025) and the dataset is publicly available<sup>5</sup>.

## 11. Bibliographical References

- Oshin Agarwal, Yinfei Yang, Byron C. Wallace, and Ani Nenkova. 2021. [Entity-Switched Datasets: An Approach to Auditing the In-Domain Robustness of Named Entity Recognition Models](#). ArXiv:2004.04123 [cs].
- AI@Meta. 2024. [Llama 3 model card](#).
- Beatrice Alex. 2008. [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. 2020. [An annotated corpus of emerging anglicisms in Spanish newspaper headlines](#). In *Proceedings of the 4th Workshop on Computational Approaches to Code Switching*, pages 1–8, Marseille, France. European Language Resources Association.
- Elena Álvarez-Mellado. 2020. [Observatorio Lázaro - Observatorio del anglicismo en la prensa española](#). Online.
- Elena Álvarez Mellado. 2025. *Lexical borrowing detection as a sequence labeling task: Data, modeling and evaluation methods for anglicism retrieval in Spanish*. Ph.D. thesis, Universidad Nacional de Educación a Distancia (UNED).
- Elena Álvarez-Mellado and Julio Gonzalo. 2024. [Characterizing spans for sequence labeling: A case on anglicism detection](#). *Procesamiento del lenguaje natural*, 73:235–246.
- Elena Álvarez-Mellado and Constantine Lignos. 2022. [Detecting unassimilated borrowings in Spanish: An annotated corpus and approaches to modeling](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3868–3888, Dublin, Ireland. Association for Computational Linguistics.
- Elena Álvarez-Mellado, Jordi Porta-Zamorano, Constantine Lignos, and Julio Gonzalo. 2025. Overview of ADoBo at IberLEF 2025: Automatic detection of anglicisms in Spanish. *Procesamiento del Lenguaje Natural*, 75:373–383.
- 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.
- Aymé Arango, Jorge Pérez, and Barbara Poblete. 2019. [Hate speech detection is not as easy as you may think: A closer look at model validation](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR'19*, page 45–54, New York, NY, USA. Association for Computing Machinery.
- Ron Artstein and Massimo Poesio. 2008. [Survey article: Inter-coder agreement for computational linguistics](#). *Computational Linguistics*, 34(4):555–596.
- Dhananjay Ashok and Zachary C. Lipton. 2023. [PromptNER: Prompting For Named Entity Recognition](#). ArXiv:2305.15444 [cs].
- Gabriel Bernier-Colborne and Phillippe Langlais. 2020. [HardEval: Focusing on challenging tokens to assess robustness of NER](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 1704–1711, Marseille, France. European Language Resources Association.
- Samuel R. Bowman and George Dahl. 2021. [What will it take to fix benchmarking in natural language understanding?](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4843–4855, Online. Association for Computational Linguistics.
- Luis Chiruzzo, Marvin Agüero-Torales, Gustavo Giménez-Lugo, Aldo Alvarez, Yliana Rodríguez, Santiago Góngora, and Thamar Solorio. 2023.

<sup>4</sup><https://adobo-task.github.io>

<sup>5</sup><https://github.com/lirondos/blas>

- Overview of GUA-SPA at IberLEF 2023: Guarani-Spanish Code Switching Analysis. *Procesamiento del Lenguaje Natural*, 71:321–328.
- Allyson Ettinger, Sudha Rao, Hal Daumé III, and Emily M. Bender. 2017. [Towards linguistically generalizable NLP systems: A workshop and shared task](#). In *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Systems*, pages 1–10, Copenhagen, Denmark. Association for Computational Linguistics.
- Jinlan Fu, Pengfei Liu, and Graham Neubig. 2020. [Interpretable multi-dataset evaluation for named entity recognition](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6058–6069, Online. Association for Computational Linguistics.
- 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.
- Abbas Ghaddar, Philippe Langlais, Ahmad Rashid, and Mehdi Rezagholizadeh. 2021. [Context-aware adversarial training for name regularity bias in named entity recognition](#). *Transactions of the Association for Computational Linguistics*, 9:586–604.
- Kyle Gorman and Steven Bedrick. 2019. [We need to talk about standard splits](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2786–2791, Florence, Italy. Association for Computational Linguistics.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, et al. 2024. [The Llama 3 herd of models](#).
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Georg Heigold, Stalin Varanasi, Günter Neumann, and Josef van Genabith. 2018. [How robust are character-based word embeddings in tagging and MT against word scrambling or random noise?](#) In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pages 68–80, Boston, MA. Association for Machine Translation in the Americas.
- Dieuwke Hupkes, Mario Giulianelli, Verna Dankers, Mikel Artetxe, Yanai Elazar, Tiago Pimentel, Christos Christodoulopoulos, Karim Lasri, Naomi Saphra, Arabella Sinclair, et al. 2023. [A taxonomy and review of generalization research in nlp](#). *Nature Machine Intelligence*, 5(10):1161–1174.
- Hannah Kermes. 2004. [Text analysis meets computational lexicography](#). In *COLING 2004: Proceedings of the 20th International Conference on Computational Linguistics*, pages 854–860, Geneva, Switzerland. COLING.
- Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenertorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina Williams. 2021. [Dynabench: Rethinking benchmarking in NLP](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4110–4124, Online. Association for Computational Linguistics.
- Venelin Kovatchev and Matthew Lease. 2024. [Benchmark transparency: Measuring the impact of data on evaluation](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1536–1551, Mexico City, Mexico. Association for Computational Linguistics.
- Bill Yuchen Lin, Wenyang Gao, Jun Yan, Ryan Moreno, and Xiang Ren. 2021. [RockNER: A simple method to create adversarial examples for evaluating the robustness of named entity recognition models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3728–3737, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hongyu Lin, Yaojie Lu, Jialong Tang, Xianpei Han, Le Sun, Zhicheng Wei, and Nicholas Jing Yuan.

2020. [A rigorous study on named entity recognition: Can fine-tuning pretrained model lead to the promised land?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7291–7300, Online. Association for Computational Linguistics.
- Tal Linzen. 2020. [How can we accelerate progress towards human-like linguistic generalization?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5210–5217, Online. Association for Computational Linguistics.
- 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.
- Ruotian Ma, Xiaolei Wang, Xin Zhou, Qi Zhang, and Xuanjing Huang. 2023. [Towards building more robust NER datasets: An empirical study on NER dataset bias from a dataset difficulty view.](#) In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4616–4630, Singapore. Association for Computational Linguistics.
- Stephen Mayhew, Tatiana Tsygankova, and Dan Roth. 2019. [ner and pos when nothing is capitalized.](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6256–6261, Hong Kong, China. Association for Computational Linguistics.
- Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. [Stress test evaluation for natural language inference.](#) In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2340–2353, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Marcin Namysl, Sven Behnke, and Joachim Köhler. 2021. [Empirical error modeling improves robustness of noisy neural sequence labeling.](#) In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 314–329, Online. Association for Computational Linguistics.
- Sean Papay, Roman Klinger, and Sebastian Padó. 2020. [Dissecting span identification tasks with performance prediction.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4881–4895, Online. Association for Computational Linguistics.
- Amandalynne Paullada, Inioluwa Deborah Raji, Emily M. Bender, Emily Denton, and Alex Hanna. 2021. [Data and its \(dis\)contents: A survey of dataset development and use in machine learning research.](#) *Patterns*, 2(11). Publisher: Elsevier.
- 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.
- Ehud Reiter. 2025. [We should evaluate real-world impact.](#) *Computational Linguistics*, pages 1–13.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018. [Semantically equivalent adversarial rules for debugging NLP models.](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 856–865, Melbourne, Australia. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of NLP models with CheckList.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Pedro Rodriguez, Joe Barrow, Alexander Miserlis Hoyle, John P. Lalor, Robin Jia, and Jordan Boyd-Graber. 2021. [Evaluation examples are not equally informative: How should that change NLP leaderboards?](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4486–4503, Online. Association for Computational Linguistics.
- Félix Rodríguez González. 2002. Spanish. In Manfred Görlach, editor, *English in Europe*, chapter 7, pages 128–150. Oxford University Press.
- F. Rodríguez González and A. Lillo Buades. 2009. *Nuevo diccionario de anglicismos*. Biblioteca románica hispánica: Diccionarios. Gredos.
- Andrew Rueda, Elena Alvarez-Mellado, and Constantine Lignos. 2024. [CoNLL#: Fine-grained error analysis and a corrected test set for CoNLL-03 English.](#) In *Proceedings of the 2024*

- Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3718–3728, Torino, Italia. ELRA and ICCL.
- Oscar Sainz, Jon Campos, Iker García-Ferrero, Julen Etxaniz, Oier Lopez de Lacalle, and Eneko Agirre. 2023. [NLP evaluation in trouble: On the need to measure LLM data contamination for each benchmark](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10776–10787, Singapore. Association for Computational Linguistics.
- 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.
- Anders Søgaard, Sebastian Ebert, Jasmijn Bastings, and Katja Filippova. 2021. [We need to talk about random splits](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1823–1832, Online. Association for Computational Linguistics.
- Tomasz Stanislawek, Anna Wróblewska, Alicja Wójcicka, Daniel Ziembicki, and Przemysław Biecek. 2019. [Named entity recognition - is there a glass ceiling?](#) In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 624–633, Hong Kong, China. Association for Computational Linguistics.
- Yulia Tsvetkov, Waleed Ammar, and Chris Dyer. 2015. [Constraint-based models of lexical borrowing](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 598–608, Denver, Colorado. Association for Computational Linguistics.
- Jingxuan Tu and Constantine Lignos. 2021. [TMR: Evaluating NER recall on tough mentions](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 155–163, Online. Association for Computational Linguistics.
- Sowmya Vajjala and Ramya Balasubramaniam. 2022. [What do we really know about state of the art NER?](#) In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 5983–5993, Marseille, France. European Language Resources Association.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. [SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems](#). In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, volume 32, pages 3266–3280.
- Xiao Wang, Qin Liu, Tao Gui, Qi Zhang, Yicheng Zou, Xin Zhou, Jiacheng Ye, Yongxin Zhang, Rui Zheng, Zexiong Pang, Qinzhuo Wu, Zhengyan Li, Chong Zhang, Ruotian Ma, Zichu Fei, Ruijian Cai, Jun Zhao, Xingwu Hu, Zhiheng Yan, Yiding Tan, Yuan Hu, Qiyuan Bian, Zhihua Liu, Shan Qin, Bolin Zhu, Xiaoyu Xing, Jinlan Fu, Yue Zhang, Minlong Peng, Xiaoqing Zheng, Yaqian Zhou, Zhongyu Wei, Xipeng Qiu, and Xuanjing Huang. 2021. [TextFlint: Unified multilingual robustness evaluation toolkit for natural language processing](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 347–355, Online. Association for Computational Linguistics.
- Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, and Daniel Weld. 2019. [Errudite: Scalable, reproducible, and testable error analysis](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 747–763, Florence, Italy. Association for Computational Linguistics.
- Lexin Zhou, Pablo A. Moreno-Casares, Fernando Martínez-Plumed, John Burden, Ryan Burnell, Lucy Cheke, Cèsar Ferri, Alexandru Marcoci, Behzad Mehrbakhsh, Yael Moros-Daval, Seán Ó hÉigeartaigh, Danaja Rutar, Wout Schellaert, Konstantinos Voudouris, and José Hernández-Orallo. 2023. [Predictable Artificial Intelligence](#). ArXiv:2310.06167 [cs].

## 12. Language Resource References

- Cooper, Robin and Crouch, Dick and Van Eijck, Jan and Fox, Chris and Jaspars, Jan and Kamp, Hans and Milward, David and Pinkal, Manfred and Poesio, Massimo and others. 1996. *Using the framework*.
- Kiela, Douwe and Bartolo, Max and Nie, Yixin and Kaushik, Divyansh and Geiger, Atticus and Wu, Zhengxuan and Vidgen, Bertie and Prasad, Grusha and Singh, Amanpreet and Ringshia, Pratik and Ma, Zhiyi and Thrush, Tristan and

Riedel, Sebastian and Waseem, Zeerak and Stenetorp, Pontus and Jia, Robin and Bansal, Mohit and Potts, Christopher and Williams, Adina. 2021. *Dynabench: Rethinking Benchmarking in NLP*. Association for Computational Linguistics.

Levesque, Hector J. and Davis, Ernest and Morgenstern, Leora. 2012. *The Winograd schema challenge*. AAAI Press, KR'12.

Pradhan, Sameer and Moschitti, Alessandro and Xue, Nianwen and Ng, Hwee Tou and Björkelund, Anders and Uryupina, Olga and Zhang, Yuchen and Zhong, Zhi. 2013. *Towards Robust Linguistic Analysis using OntoNotes*. Association for Computational Linguistics.

Rueda, Andrew and Alvarez-Mellado, Elena and Lignos, Constantine. 2024. *CoNLL#: Fine-grained Error Analysis and a Corrected Test Set for CoNLL-03 English*. ELRA and ICCL.

Tjong Kim Sang, Erik F. and De Meulder, Fien. 2003. *Introduction to the CoNLL-2003 Shared Task: Language-Independent Named Entity Recognition*.

## A. Guidelines for the creation of examples in BLAS

In this section we document the criteria that guided the anglicism selection, writing process and punctuation transformations behind the examples in BLAS.

### A.1. Anglicism selection criteria

The first step when writing the examples for BLAS was to select the anglicisms that would serve to evaluate a given linguistic attribute. Anglicisms were selected by hand from the anglicism dictionary *Nuevo diccionario de anglicismos* by Rodríguez González and Lillo Buades (2009) and from Observatorio Lázaro, an online resource that automatically compiles anglicisms used in the Spanish media (Álvarez-Mellado, 2020).

The following criteria were taken into account for the selection:

**Frequency criteria** With BLAS we wanted to evaluate how well models perform when they face anglicisms that are not seen in training. In other words, we wanted to ensure that most of the borrowings in BLAS would be unseen. However BLAS comes with no associated training data and, being a test-only dataset, we cannot anticipate or assume the type of data that future models evaluated on BLAS will be trained on. This means that we

cannot guarantee that the anglicisms in BLAS will be unseen for the models.

To mitigate this limitation, we made a conscious effort to include borrowings in BLAS that are more niche and less widespread. Consequently, during the selection process we prioritized less obvious anglicisms over anglicisms that we would expect to find in a monolingual dictionary of Spanish, or that are so frequent that they could be expected to appear in a generalistic training set (while still keeping a balance in terms of domains represented).

We also consciously avoided extremely rare borrowings and nonce borrowings; to do so, we made sure that the selected borrowings appeared in the Spanish general media and newspapers. For borrowings retrieved from Observatorio Lázaro, this means that we looked for borrowings that had been documented more than 10 times and less than 50 in Spanish newspapers.

**Sentence position criteria** Because we wanted to assess performance on different sentence positions, we prioritized whenever possible to select borrowings that could work equally well mid sentence and in sentence initial position.

**Topic and shape variety** The selection was also curated to ensure that a variety of topics and shapes were covered, so that not all anglicisms belonged to the same domain or had similar character combinations.

**Unassimilation criteria** Our aim with BLAS is that it can serve as a stable benchmark that future models can be evaluated on. This wish clashes against a linguistic fact of life: the fact that a borrowing that is frequent enough may very easily end up being assimilated into the recipient language and thus not perceived as a borrowing anymore. This poses the risk of including borrowings in BLAS that will eventually become assimilated and will not serve our purpose anymore. This is particularly true for graphotactically compliant anglicisms (that is, anglicisms whose spelling does not violate the spelling rules of Spanish), as the fact of being compliant may cause that speakers rapidly forget that the word was borrowed, thus facilitating its assimilation process (while non-compliance preserves the speakers' perception of word being foreign).

To account for this, for compliant borrowings we chose words where, even if their spelling did not violate the graphotactics of Spanish, there was no correspondence between the spelling of the word and its pronunciation according to the Spanish pronunciation expectations. In other words, a word like *barman* is a graphotactically compliant borrowing whose spelling matches the expectations of how the word is pronounced in Spanish (/ˈbar.man/).

This makes the borrowing *barman* a bad choice for BLAS, because due to its shape and pronunciation it makes it a good candidate for rapid assimilation. On the other hand, borrowings like *paper* (/ˈpeɪ.pər/) or *eyeliner* (/aɪ.ˈlaɪ.nər/), while still being fully compliant in their spelling, are currently pronounced in a way that violates the expectations of Spanish, which ensures that they are perceived as foreign by native speakers and that are far from becoming assimilated in the near future. Consequently, for graphotactically compliant borrowings we chose borrowings where there was no correspondence between its spelling and pronunciation.

### A.1.1. Compliant and not-compliant spans

Compliant anglicisms are anglicisms whose spelling complies with the Spanish rules of spelling. On the other hand, non-compliant anglicisms are anglicisms whose spelling violates the rules of spelling. Non-compliant borrowings are easier to spot, because their spelling violates the expectations that Spanish speakers have about how a word should be written. As a result, previous work had already pointed out that models tended to privilege non-compliant borrowings over compliant ones (Álvarez-Mellado and Gonzalo, 2024). To assess differences in performance between compliant and non-compliant borrowings, we selected the following anglicisms:

- 20 graphotactically compliant anglicisms (10 single word anglicisms, 10 multiword anglicisms): for example, *pipeline*, *medal race*.
- 20 graphotactically non-compliant anglicisms (10 single word anglicisms, 10 multiword anglicisms): for example, *cliffhanger*, *fact checker*.
- 10 graphotactically mixed-compliant anglicisms, where at least one token is compliant and one token is non-compliant (all multiword): for example, *tie break*.

### A.1.2. Ambiguous spans

Ambiguous anglicisms are anglicisms that also exist as fully Spanish words (such as *pie*, that is a native word when it means “foot” but a borrowing pronounced as /pai/ when it refers to a dish). Ambiguous spans can easily confuse models, as they can be seen during training as native words or found as native words in Spanish lexicons. Our aim was to evaluate specifically how well models perform on borrowings that are prone to label mismatch.

As a result, we collected a set of anglicisms that are ambiguous with Spanish words (either single word items or multiword items), following the same procedure we described in section A.1.1. The following borrowings were considered:

- 3 single-word ambiguous borrowings: for example, *pie*.
- 10 multiword ambiguous borrowings: for example, *global director*.
- 10 multiword mixed-ambiguous borrowings (where at least one of the words is ambiguous and one of the tokens is not): for example, *escape room*.

We identified very few candidates that were single-word ambiguous (only 3), but we still considered them to be interesting enough to include them in BLAS, as these examples can help us assess models’ performance on an extremely difficult case.

For pure ambiguous borrowings (that is, not mixed ambiguous), we lifted the restriction that borrowings had to work in sentence initial position: given the limited number of ambiguous borrowings and in order to avoid populating our benchmark with unnatural examples (not all borrowings work equally well on sentence initial position and how natural a borrowing in sentence initial position sounds is given by the meaning of the borrowing), we focused exclusively on mid sentence position examples, where all borrowings work equally well.

Consequently, for pure ambiguous borrowings we only wrote one example per selected borrowing, in which the borrowing appeared mid-sentence. We still wrote two examples (one sentence initial, one mid-sentence) for mixed-ambiguous anglicisms, as mixed-ambiguous are not rare and we could easily find good candidates for mixed-ambiguous that could work on sentence initial position.

### A.1.3. Adjacent spans

Finally, we wanted our benchmark to evaluate performance over sentences with adjacent borrowings, ie. two independent borrowings that are collocated next to one another in the sentence (as in *look total black*), as this type of sentences are prone to segmentation errors. We selected the following 20 pairs of borrowings that worked well collocated:

- 10 pairs of collocated borrowings, where both borrowings were single-word borrowings; for example, *marketing online* in *Una campaña de marketing online*<sup>6</sup>.
- 10 pairs of collocated borrowings, where at least one of the borrowings in the pair was multiword; for example, *look total black* in *La actriz lució un look total black*<sup>7</sup>.

<sup>6</sup>“An online marketing campaign.”

<sup>7</sup>“The actress wore a total black look.”

Given that not all borrowings will work equally well collocated with other borrowings and in order to ensure natural sounding sentences, the selection of borrowings in this case was more lenient: more frequent borrowings were allowed in the selection and only mid-sentence position sentences were included, as what we were interested in assessing with these examples was not the retrieval capability of the models but their ability to correctly segment borrowings that may have been previously seen.

Similarly to ambiguous anglicisms, for adjacent anglicisms we only wrote one example per selected pair of borrowings, in which the borrowings appeared mid-sentence.

## A.2. Writing process

BLAS was populated with sentences that contained the anglicisms selected in Appendix A.1.

Sentences in BLAS were created ad hoc by hand by a linguist inspired on real examples from journalistic sources. Sentences in BLAS are characterized by a given configuration of linguistic attributes: a group of sentences contains multiword anglicisms that are non-compliant and that appear mid-sentence, another group of sentences contain the same multiword anglicisms but in sentence initial position, etc.

All examples in BLAS are original and were written from scratch, although some of them were inspired by real world examples and all of them try to mimic the journalistic style of the Spanish newspapers and media. The linguist that authored the examples was a native speaker of Spanish with experience in anglicism annotation and journalistic writing. Figure 1 reflects the writing process behind the examples in BLAS. Table 1 reflects number of sentences per linguistic attribute.

## A.3. Punctuation transformations

The process described in Appendix A.2 yielded 153 sentences. These sentences underwent a series of automatic punctuation transformations. We considered two types of punctuation transformations: casing and quotation marks.

For casing, the following transformations were performed:

- Standard casing (no transformations, text as is).
- all words are lowercase.
- ALL WORDS ARE UPPERCASE.
- All Words Are Capitalized.
- Only BORROWINGS are in uppercase.
- Only Borrowings are capitalized.

For quotation marks, the following perturbations were performed:

- All “anglicisms” appear between quotation marks.
- No anglicism is written with quotation marks.

This amounts to a total of twelve different configurations (six casing transformations with quotation marks, six without them). Each sentence in the original collection of 153 sentences underwent these twelve transformations. Table 2 shows an example of all twelve casing and punctuation transformations for a single sentence.

## B. Prompting methodology with Llama3

Our prompting methodology was inspired by the work of Ashok and Lipton (2023), which offer prompted-based heuristics to apply to generative LLMs for the task of NER. Their experiments show that in addition to the few-shot examples, providing the model with a modular definition of the relevant entity types in the prompt produces better results than pure few-shot prompting.

To that end, we prompted 8B-Llama3 (Grattafiori et al., 2024; AI@Meta, 2024) with a definition of what an anglicism is and a description of the task the model was expected to complete. The definition included a list of four examples of decontextualized anglicisms (*online*, *machine learning*, *podcast*, *blockchain*) and a description of some of the characteristics and contexts an anglicism may appear in (single-word, multiword, adjacent, quoted or unquoted) illustrated by a minimal collection of the following three examples:

- *El evento se retransmitirá por streaming.*<sup>8</sup>
- *La marca ha sacado a la venta un nuevo concept car.*<sup>9</sup>
- *Compré mi smartphone online.*<sup>10</sup>

The definition explicitly mentioned that proper names such as entity names should not be considered anglicisms.

Figure 2 displays the prompt we used, which draws from the style of prompting for NER proposed by Ashok and Lipton.

We chose Llama3 for our experiments because at the time of conducting our experiments Llama3 was among the best performing models while being open source, which ensured that our conclusions could be tested, replicated and easily transferred to other languages. The model we used was the small Llama3 (8B parameter), as it is less resource intensive than larger versions. We queried it through `ollama` library<sup>11</sup>, with default context window size of 2048 and temperature set to 0.0.

<sup>8</sup>“The event will be broadcast live via streaming”.

<sup>9</sup>“The company has released a new concept car”.

<sup>10</sup>“I bought my smartphone online”.

<sup>11</sup><https://ollama.com/>

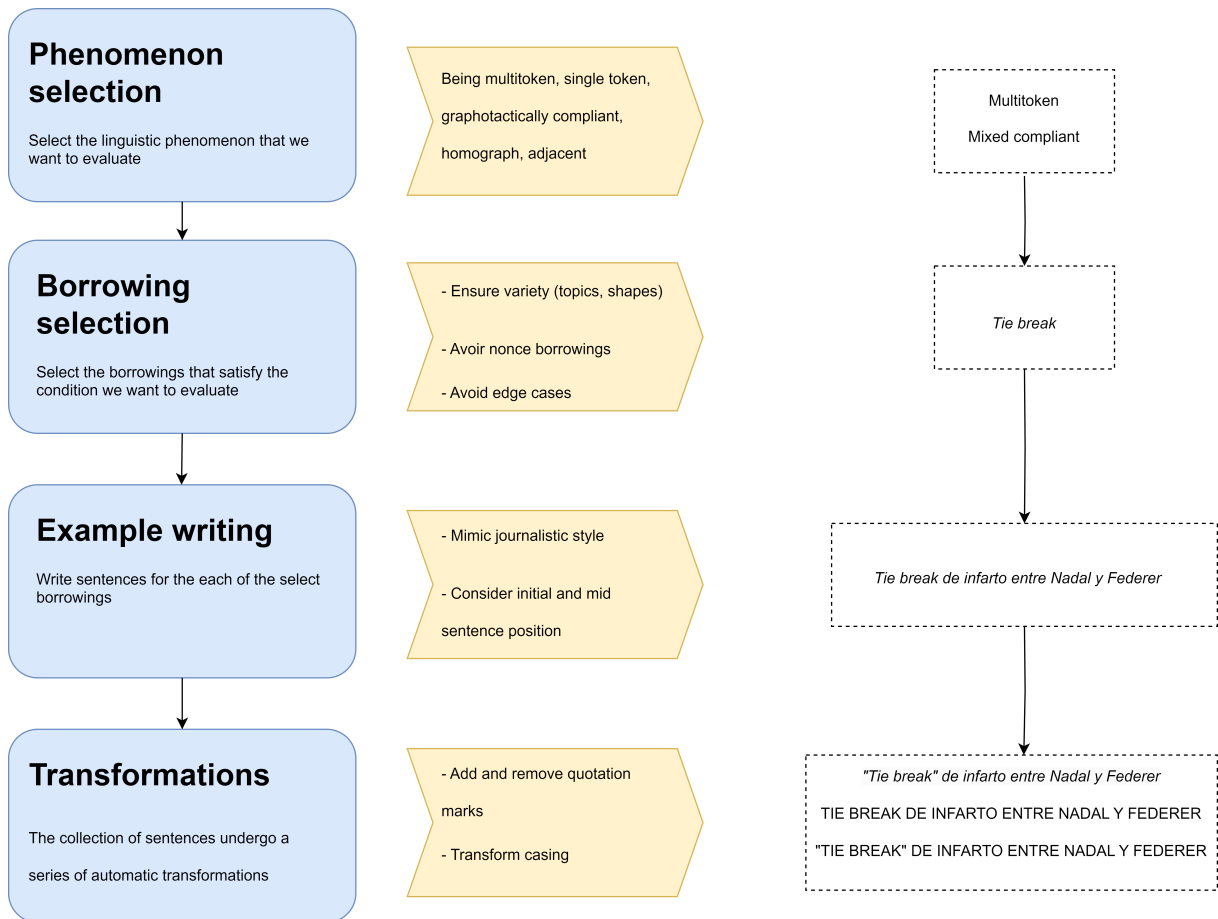


Figure 1: Creation process of the examples in BLAS.

### C. Prediction experiments

Tables 7, 8, 9, 10 and 11 show number of spans per type on each of the five external datasets: COALAS test set and development set (Álvarez-Mellado and Lignos, 2022) and CALCS development and test sets A and B (Alvarez-Mellado, 2020).

Table 12 displays the predicted recall scores and ranking, true recall scores and ranking and true F1 scores and ranking for all six models over the five external datasets, along with Pearson and Spearman correlation coefficients.

Table 13 displays average and median correlation coefficients between the predicted results and true results, along with correlation coefficients between BLAS results and true results.

Casing	Quotation marks	Type	Length	Position	Number
Standard	w/o quot.	non compliant	single	mid	688
Standard	w/o quot.	compliant	single	mid	141
Standard	w/o quot.	non compliant	multi	mid	123
Standard	w/o quot.	mixed compliant	multi	mid	107
Standard	with quot.	non compliant	single	mid	42
Standard	with quot.	non compliant	multi	mid	20
Standard	with quot.	mixed compliant	multi	mid	20
Standard	w/o quot.	adjacent	single	mid	14
Standard	w/o quot.	mixed ambiguous	multi	mid	13
Standard	w/o quot.	non compliant	single	ini	11
Standard	w/o quot.	compliant	multi	mid	9
Standard	with quot.	compliant	multi	mid	9
Spans are titlecase	w/o quot.	non compliant	single	mid	8
Standard	with quot.	compliant	single	mid	7
Standard	w/o quot.	non compliant	multi	ini	4
Text is lowercase	w/o quot.	non compliant	multi	mid	3
Standard	w/o quot.	ambiguous	multi	mid	3
Standard	w/o quot.	mixed compliant	multi	ini	3
Spans are titlecase	w/o quot.	compliant	single	mid	2
Spans are titlecase	w/o quot.	mixed ambiguous	multi	mid	2
Standard	w/o quot.	ambiguous	single	mid	2
Standard	w/o quot.	adjacent	multi	mid	2
Spans are uppercase	w/o quot.	non compliant	single	ini	1
Spans are titlecase	w/o quot.	non compliant	multi	mid	1
Spans are titlecase	with quot.	non compliant	multi	mid	1
Standard	w/o quot.	compliant	single	ini	1
Standard	w/o quot.	mixed ambiguous	multi	ini	1
Spans are titlecase	w/o quot.	mixed compliant	multi	mid	1

Table 7: Number of spans per type in COALAS test set.

Casing	Quotation marks	Type	Length	Position	Number
Standard casing	with quot.	non compliant	single	mid	90
Standard casing	w/o quot.	non compliant	single	mid	79
Standard casing	with quot.	non compliant	multi	mid	30
Standard casing	with quot.	mixed compliant	multi	mid	22
Standard casing	with quot.	compliant	single	mid	19
Standard casing	w/o quot.	compliant	single	mid	17
Standard casing	with quot.	compliant	multi	mid	9
Standard casing	w/o quot.	non compliant	multi	mid	8
Standard casing	w/o quot.	mixed compliant	multi	mid	6
Spans are titlecase	with quot.	mixed compliant	multi	mid	5
Spans are titlecase	with quot.	non compliant	single	mid	4
Spans are titlecase	w/o quot.	mixed compliant	multi	mid	3
Standard casing	with quot.	non compliant	single	ini	2
Spans are titlecase	w/o quot.	non compliant	multi	mid	2
Standard casing	w/o quot.	compliant	multi	mid	2
Standard casing	w/o quot.	non compliant	single	ini	1
Spans are titlecase	with quot.	non compliant	multi	mid	1
Standard casing	w/o quot.	non compliant	multi	ini	1
Standard casing	with quot.	compliant	single	ini	1
Standard casing	w/o quot.	compliant	single	ini	1
Standard casing	with quot.	ambiguous	multi	mid	1
Standard casing	with quot.	ambiguous	single	mid	1
Spans are titlecase	with quot.	mixed ambiguous	multi	mid	1

Table 8: Number of spans per type in COALAS development set.

<b>Casing</b>	<b>Quotation marks</b>	<b>Type</b>	<b>Length</b>	<b>Position</b>	<b>Number</b>
Standard casing	w/o quot.	non compliant	single	mid	67
Standard casing	with quot.	non compliant	single	mid	55
Standard casing	w/o quot.	compliant	single	mid	18
Standard casing	with quot.	compliant	single	mid	15
Standard casing	w/o quot.	non compliant	single	ini	10
Standard casing	w/o quot.	non compliant	multi	mid	9
Standard casing	with quot.	non compliant	multi	mid	8
Standard casing	with quot.	mixed compliant	multi	mid	7
Standard casing	w/o quot.	ambiguous	multi	mid	5
Spans are titlecase	w/o quot.	non compliant	single	mid	3
Standard casing	with quot.	compliant	multi	mid	3
Spans are titlecase	w/o quot.	mixed compliant	multi	mid	3
Standard casing	w/o quot.	mixed compliant	multi	mid	3
Standard casing	w/o quot.	mixed ambiguous	multi	mid	3
Standard casing	with quot.	mixed ambiguous	multi	mid	2
Standard casing	with quot.	non compliant	single	ini	1
Text is lowercase	w/o quot.	non compliant	single	mid	1
Spans are titlecase	w/o quot.	compliant	single	mid	1
Standard casing	with quot.	compliant	single	ini	1
Text is titlecase	w/o quot.	compliant	single	ini	1
Standard casing	with quot.	ambiguous	multi	mid	1
Standard casing	w/o quot.	mixed compliant	multi	ini	1
Standard casing	with quot.	mixed ambiguous	multi	ini	1

Table 9: Number of spans per type in CALCS development set.

Casing	Quotation marks	Type	Length	Position	Number
Standard casing	w/o quot.	non compliant	single	mid	77
Standard casing	with quot.	non compliant	single	mid	47
Standard casing	w/o quot.	compliant	single	mid	19
Standard casing	with quot.	compliant	single	mid	14
Standard casing	w/o quot.	non compliant	multi	mid	7
Standard casing	w/o quot.	mixed compliant	multi	mid	6
Standard casing	with quot.	mixed compliant	multi	mid	6
Standard casing	with quot.	non compliant	multi	mid	5
Standard casing	w/o quot.	non compliant	single	ini	4
Standard casing	with quot.	compliant	multi	mid	4
Standard casing	with quot.	mixed compliant	multi	ini	2
Spans are titlecase	with quot.	mixed compliant	multi	mid	2
Spans are titlecase	w/o quot.	non compliant	single	mid	2
Standard casing	w/o quot.	mixed ambiguous	multi	mid	2
Standard casing	w/o quot.	mixed compliant	multi	ini	1
Spans are titlecase	w/o quot.	mixed compliant	multi	mid	1
Spans are titlecase	with quot.	non compliant	single	mid	1
Text is lowercase	with quot.	non compliant	single	mid	1
Standard casing	with quot.	non compliant	single	ini	1
Standard casing	w/o quot.	non compliant	multi	ini	1
Standard casing	w/o quot.	compliant	multi	mid	1
Standard casing	w/o quot.	compliant	multi	ini	1
Text is lowercase	w/o quot.	compliant	single	mid	1
Standard casing	with quot.	compliant	single	ini	1
Standard casing	w/o quot.	compliant	single	ini	1
Spans are titlecase	w/o quot.	mixed ambiguous	multi	mid	1
Standard casing	with quot.	mixed ambiguous	multi	mid	1
Spans are titlecase	with quot.	mixed ambiguous	multi	mid	1
Standard casing	w/o quot.	ambiguous	multi	mid	1

Table 10: Number of spans per type in CALCS test set A.

Casing	Quotation marks	Type	Length	Position	Number
Standard casing	w/o quot.	non compliant	single	mid	43
Standard casing	w/o quot.	compliant	single	mid	25
Standard casing	with quot.	non compliant	single	mid	20
Standard casing	with quot.	compliant	single	mid	12
Standard casing	with quot.	mixed compliant	multi	mid	4
Standard casing	w/o quot.	non compliant	multi	mid	3
Standard casing	with quot.	non compliant	multi	mid	3
Standard casing	w/o quot.	non compliant	single	ini	2
Standard casing	with quot.	non compliant	single	ini	2
Standard casing	w/o quot.	mixed compliant	multi	mid	2
Standard casing	w/o quot.	ambiguous	multi	mid	2
Standard casing	w/o quot.	compliant	single	ini	1
Standard casing	with quot.	compliant	multi	mid	1
Spans are titlecase	with quot.	compliant	multi	mid	1
Standard casing	w/o quot.	compliant	multi	mid	1
Spans are titlecase	with quot.	non compliant	single	mid	1
titlecased text	with quot.	non compliant	single	ini	1
Spans are titlecase	w/o quot.	mixed compliant	multi	mid	1
Standard casing	with quot.	mixed ambiguous	multi	mid	1

Table 11: Number of spans per type in CALCS test set B.

Un anglicismo es un préstamo crudo incorporado del inglés que se usa dentro de una frase en castellano sin adaptación, palabras como 'online', 'machine learning', 'podcast' o 'blockchain'. Un anglicismo puede estar formado por una única palabra (como 'streaming' en 'El evento se retransmitirá por streaming'), por varias palabras (como en 'La marca ha sacado a la venta un nuevo concept car') o ser dos anglicismos independientes que aparecen uno al lado del otro (como 'smartphone' y 'online' en 'Compré mi smartphone online'), y pueden aparecer entrecomillados o no. Los nombres propios (como los nombres de personas, lugares o títulos de obras de ficción) no son anglicismos."

Voy a darte una frase en castellano, tu misión es identificar si la frase contiene algún anglicismo. La frase puede contener un anglicismo, varios o no contener ninguno. Responde solo con el segmento que cumpla la condición de ser anglicismo (si es que lo hay) y sin añadir ninguna consideración más. Si la frase no contiene ningún anglicismo, responde 'None'. Si la frase contiene más de 1 anglicismo, devuelve todos los anglicismos, con un anglicismo por línea (es decir, añade un salto de línea entre anglicismos). No alucines ni añadas palabras que no estén en la frase original. Esta es la frase:

El problema de la "fast fashion" es triple : consume recursos desmesuradamente , produce toneladas de basura de difícil reciclaje y fomenta la explotación laboral de la industria textil .

"fast fashion"

Figure 2: Prompt provided to Llama3 inspired by Ashok and Lipton's template: Definition of the task with examples, question and task, response by Llama3.

Translation: An anglicism is a lexical borrowing that is incorporated from English into Spanish without adaptation, such as 'online', 'machine learning', 'podcast' or 'blockchain'. An anglicism can be composed of one single word (like 'streaming' in 'El evento se retransmitirá por streaming'), several words (like 'concept car' in 'La marca ha sacado a la venta un nuevo concept car') or two adjacent anglicisms (like 'smartphone' y 'online' en 'Compré mi smartphone online') and can be written between quotation marks or not. Proper foreign names (like person names, location names or titles) are not anglicisms.

You will be given a sentence in Spanish, your mission is identifying if the sentence contains anglicisms. The sentence may contain one anglicism, several anglicisms or none. Reply only with the span that satisfies the requirement of being an anglicism (if there is any), without any other information. If the sentence contains no anglicism simply reply 'None'. If the sentence contains more than one anglicism, reply with one anglicism per line (leave a linebreak between anglicisms). This is the sentence:

El problema de la "fast fashion" es triple : consume recursos desmesuradamente , produce toneladas de basura de difícil reciclaje y fomenta la explotación laboral de la industria textil .

"fast fashion"

Models	BLAS	COALAS (test set)						COALAS (dev set)						CALCS (test set A)						CALCS (test set B)						CALCS (dev set)					
	R	Pred. R	True R	True F1			Pred. R	True R	True F1			Pred. R	True R	True F1			Pred. R	True R	True F1			Pred. R	True R	True F1							
CRF	6.50	26.58 #6	44.31 #5	56.54 #5			33.13 #6	68.63 #5	71.31 #5			29.00 #6	65.57 #5	70.20 #5			24.76 #6	53.17 #5	61.19 #5			28.99 #6	64.84 #5	71.90 #5							
BETO	23.55	86.87 #3	77.99 #4	82.36 #4			87.62 #3	84.05 #4	78.81 #4			83.58 #4	83.49 #3	82.71 #4			79.13 #4	81.75 #4	76.87 #3			81.60 #4	86.30 #1	87.30 #2							
mBERT	23.80	86.86 #4	78.85 #3	83.64 #3			87.10 #4	84.31 #3	82.21 #3			83.86 #3	85.85 #1	84.26 #2			83.17 #3	84.13 #3	76.26 #4			82.73 #3	86.30 #1	86.30 #3							
BiLSTM (unad)	23.55	87.86 #2	80.88 #2	85.14 #1			88.73 #2	89.05 #1	86.73 #1			84.85 #2	82.08 #4	83.86 #3			84.04 #2	84.92 #2	84.25 #1			83.06 #2	83.11 #4	85.25 #4							
BiLSTM (cs)	23.75	92.96 #1	85.76 #1	84.22 #2			91.08 #1	86.57 #2	84.66 #2			89.43 #1	83.96 #2	84.96 #1			89.52 #1	85.71 #1	83.08 #2			87.53 #1	85.84 #3	87.44 #1							
Llama3	36.37	42.93 #5	33.33 #6	39.73 #6			55.46 #5	53.27 #6	40.30 #6			48.53 #5	53.30 #6	29.62 #6			43.17 #5	40.48 #6	17.17 #6			48.49 #5	50.68 #6	30.41 #6							
Correlation			0.94	0.89	0.89	0.83		0.79	0.83	0.61	0.83		0.85	0.66	0.64	0.89		0.91	0.94	0.69	0.83		0.84	0.41	0.63	0.71					

Table 12: Pearson correlation coefficient over true recall scores, F1 scores and ranking obtained by six models over different datasets and the predicted recall scores and rankings obtained by extrapolating BLAS scores.

Dataset	Recall correl. with BLAS prediction	Ranking correl. with BLAS prediction	Recall correl. with BLAS	Ranking correl. with BLAS
COALAS test set	0.94	0.89	-0.04	-0.08
COALAS dev set	0.79	0.83	-0.23	-0.17
CALCS test set A	0.85	0.66	-0.19	-0.03
CALCS test set B	0.91	0.94	-0.10	-0.37
CALCS dev set	0.83	0.43	-0.20	0.06
Average correl.	<b>0.86</b>	<b>0.75</b>	<b>-0.15</b>	<b>-0.12</b>
Median correl.	<b>0.85</b>	<b>0.83</b>	<b>-0.19</b>	<b>-0.08</b>

Table 13: Pearson and Spearman correlation coefficients obtained from comparing recall score distributions and ranking between the true results obtained on a dataset, the predicted results obtained by extrapolating BLAS scores and BLAS results.