

Evaluation Benchmarks for Spanish Sentence Representations

Vladimir Araujo^{1,2,6,7}, Andrés Carvallo^{1,6}, Souvik Kundu³, José Cañete^{4,6},
Marcelo Mendoza^{5,6,7}, Robert E. Mercer³, Felipe Bravo-Marquez^{4,6,7},
Marie-Francine Moens², Alvaro Soto^{1,6}

¹Pontificia Universidad Católica de Chile, Santiago, Chile

²KU Leuven, Leuven, Belgium

³University of Western Ontario, London, Canada

⁴Department of Computer Science, University of Chile, Santiago, Chile

⁵Universidad Técnica Federico Santa María, Santiago, Chile

⁶National Center for Artificial Intelligence Research (CENIA), Santiago, Chile

⁷Millennium Institute for Foundational Research on Data (IMFD), Santiago, Chile

vgaraujo@uc.cl

Abstract

Due to the success of pre-trained language models, versions of languages other than English have been released in recent years. This fact implies the need for resources to evaluate these models. In the case of Spanish, there are few ways to systematically assess the models' quality. In this paper, we narrow the gap by building two evaluation benchmarks. Inspired by previous work (Conneau and Kiela, 2018; Chen et al., 2019), we introduce Spanish SentEval and Spanish DiscoEval, aiming to assess the capabilities of stand-alone and discourse-aware sentence representations, respectively. Our benchmarks include considerable pre-existing and newly constructed datasets that address different tasks from various domains. In addition, we evaluate and analyze the most recent pre-trained Spanish language models to exhibit their capabilities and limitations. As an example, we discover that for the case of discourse evaluation tasks, mBERT, a language model trained on multiple languages, usually provides a richer latent representation than models trained only with documents in Spanish. We hope our contribution will motivate a fairer, more comparable, and less cumbersome way to evaluate future Spanish language models.

Keywords: discourse evaluation, language models, sentence evaluation, representation learning

1. Introduction

Spanish is one of the most widely spoken languages. This fact has drawn the attention of the NLP community to the development of resources for that language. As a result, some pre-trained Spanish language models (Etcheverry and Wonsever, 2016; Che et al., 2018; Cañete et al., 2020; Gutiérrez-Fandiño et al., 2021) have been released in recent years driven by self-supervised approaches. This proliferation of Spanish language models increases the need for annotated datasets to evaluate them. Consequently, Spanish datasets for a wide variety of independent tasks have been proposed (Á. García Cumbreñas et al., 2006; Cruz Mata et al., 2008; Artetxe et al., 2020; Huertas-Tato et al., 2022). However, little effort has been put into creating benchmarks that allow models to be evaluated systematically and fairly.

Recently, Cañete et al. (2020) presented the GLUES benchmark, a compilation of natural language understanding tasks in Spanish. This benchmark aims to evaluate the performance of models by fine-tuning them to a target task (Wang et al., 2018). In contrast, another methodology known as probing tasks aims to assess whether the resulting representations of the models are general-purpose (Conneau et al., 2018a). A probing task is designed in such a way as to isolate some linguistic phenomena, and a classifier is used on top of the representations to verify if the model has en-

coded the linguistic phenomena in question. This type of representation evaluation for Spanish language models is generally carried out using a cross-lingual setting (Ravishankar et al., 2019; Şahin et al., 2020). However, these benchmarks only focus on assessing word representations or basic linguistic knowledge.

On the one hand, the **Spanish SentEval**, inspired by SentEval (Conneau and Kiela, 2018), aims to evaluate representations of independent sentences. Unlike previous work focused on probing tasks for basic linguistic properties (Ravishankar et al., 2019; Şahin et al., 2020), our benchmark comprises four sets of sentence classification tasks with realistic texts from different domains. On the other hand, the **Spanish DiscoEval**, inspired by DiscoEval (Chen et al., 2019), focuses on the evaluation of discourse knowledge in sentence representations. Evaluating discourse involves analyzing a sentence in the context in which it is located. For this reason, we include five sets of tasks based on sentence ordering, discourse relations, and discourse coherence. The overall objective of both benchmarks is to avoid unnecessary re-implementations and the use of multiple evaluation schemes, thus allowing a comparable and fair assessment between models. Furthermore, we compare publicly available Spanish sentence encoders on our Spanish SentEval and Spanish DiscoEval to demonstrate their strengths and weaknesses. The results and subsequent analysis

expose the Spanish language models’ current capabilities, showing that there is still room to improve them in future work. Our code and datasets are available for future experimentation and replicability at <https://github.com/OpenCENIA/Spanish-Sentence-Evaluation>.

2. Sentence Evaluation

SentEval was introduced by Conneau and Kiela (2018) as a tool for evaluating the quality of universal sentence representations. It encompasses a standard pipeline evaluation that uses the representations generated by sentence encoders as features in various downstream tasks. Specifically, SentEval includes stand-alone sentence and sentence pair tasks modeled by classification or regression. For comparison purposes, this framework consists of simple predefined neural architectures to avoid shifting the burden of modeling to their optimization process.

For our Spanish SentEval, we adopt the original pipeline and include datasets equivalent to those in English. Below we describe each task and dataset included in our Spanish version. Additionally, basic statistics for each dataset are shown in Appendix A.

As proposed in Chen et al. (2019), we use $[\cdot, \cdot, \dots]$ to denote concatenation of vectors, \odot for element-wise multiplication, and $|\cdot|$ for element-wise absolute value.

2.1. Sentence Classification (SC)

Sentence classification is one of the most common NLP tasks, with applications ranging from document classification to sentiment analysis. Because of its inherent simplicity, the task offers a straightforward way to evaluate sentence-level representations. For our version, we include a set of binary and multiclass datasets that cover various types of sentence classification tasks.

For sentiment analysis, we include MuchoChine (MC) (Cruz Mata et al., 2008), a movie review dataset, and TASS 2020 (Vega et al., 2020) tasks 1 and 2 consisting of polarity and emotion classification (Plaza del Arco et al., 2020), respectively. Figure 1 shows an example of a MC positive sentiment sentence.

Other types of text classification datasets that we include are the FilmAffinity corpus (FAC) (Sobrevilla Cabezudo et al., 2015) for subjective/objective classification and the Spanish QC dataset (SQC) (Á. García Cumbreiras et al., 2006) for question-type. For all of these tasks, the input to the classifier is the representation of the sentence.

- Una historia policiaca que Scorsese la transforma en una memorable muestra del genero.

Figure 1: SC example. The sentence belongs to the MC dataset and shows a *positive* sentiment.

2.2. Sentence Pair Classification (SPC)

In sentence pair classification, each example in a dataset has two sentences along with the appropriate

target, and the aim is to model the textual interaction between them. We consider entailment and paraphrasing tasks for our Spanish benchmark.

For the entailment task, we include two datasets. The first is the recently released SICK-es (Huertas-Tato et al., 2022) for entailment (SICK-es-E), which was constructed by translating and manually curating the English SICK dataset into Spanish. Due to the lack of NLI tasks in Spanish, the second dataset was constructed using XNLI (Conneau et al., 2018b) and esXNLI (Artetxe et al., 2020). Specifically, we use the XNLI test set as the training set, the XNLI development set as the development set, and the esXNLI set as the test set. We will refer to this as NLI-es (example shown in Figure 2).

For the paraphrasing task, we use PAWS-X (Yang et al., 2019), a cross-lingual paraphrase identification dataset with high lexical overlap. We only use the Spanish text, naming it PAWS-es for ease of reference.

Like English SentEval, we encode the two sentences and use $[|x_1 - x_2|, x_1 \odot x_2]$ as input to the classifier.

Premise: Y yo estaba bien, ¡y eso fue todo!
Hypothesis: Después de que dije que sí, terminé.

Figure 2: Example of SPC from NLI-es. The two sentences show an *entailment*.

2.3. Semantic Similarity (SS)

This task consists of scoring a pair of sentences based on their degree of similarity, even if they are not exact matches. There are two common approaches to evaluating this task and we include them in our Spanish SentEval. The first requires training a model on top of the sentence embeddings. For this approach, we use the SICK-es (Huertas-Tato et al., 2022) for relatedness (SICK-es-R). The second assesses sentence pairs using an unsupervised approach. For this case, we include the Spanish track of STS tasks 2014 (Agirre et al., 2014), 2015 (Agirre et al., 2015) and 2017 (Cer et al., 2017). All of these datasets consist of a pair of sentences labeled with a similarity score between 0 and 5; an example is shown in Figure 3. The objective is to evaluate whether the cosine similarity of two sentence representations correlates with a human-labeled similarity score through Pearson and Spearman correlations.

- Un perro está con un juguete.
- Un perro tiene un juguete.

Figure 3: Example of the STS task. The two sentences are similar with a *score* of 4.8 out of 5.

2.4. Linguistic Probing Tasks (LPT)

Some sentence classification tasks are complex and make it difficult to infer what kind of information is present in the representations. This prompted the creation of X-Probe (Ravishankar et al., 2019), a multi-lingual benchmark of nine probing tasks to evaluate

individual linguistic properties. These tasks were designed to evaluate surface information (SentLen, WC), syntactic information (BiShift, TreeDepth), and semantic information (Tense, SubjNum, ObjNum, SOMO, CoordInv). The former evaluate superficial tasks that could be solved simply by looking at the sentence tokens. The second tests whether the embeddings are sensitive to the syntactic properties of the sentences. The third assesses the semantic understanding of the embedding. We include all the proposed probing tasks in Spanish from X-Probe. We refer to the original paper (Ravishankar et al., 2019) for further information. The input to the classifier is the representation of the sentence, and the output can be binary or multiclass.

- En enero participó en la infructuosa defensa de Forlí frente a César Borgia.

Figure 4: Example of LPT. The task consists of Tense classification. In this case the sentence is in *past tense*.

3. Discourse Evaluation

DiscoEval originally proposed by Chen et al. (2019) includes tasks to evaluate discourse-related knowledge in pretrained sentence representations. DiscoEval adopts the SentEval pipeline with fixed standard hyperparameters to avoid discrepancies. For our Spanish version of DiscoEval, we follow closely the original construction and evaluation methodology. Specifically, DiscoEval includes supervised sentence and sentence group classification tasks modeled by logistic regression or classification. Our datasets were constructed from multiple domains encompassing a wide diversity of text sources. Below we describe the tasks and dataset constructions. Statistics for each dataset are shown in Appendix A.

3.1. Sentence Position (SP)

SP seeks to assess the ability of a model to order ideas in a paragraph. This dataset is constructed by taking five consecutive sentences from a given corpus and randomly moving one of these five sentences to the first position. The task consists of predicting the proper location of the first sentence. We have five classes where class 1 means that the first sentence is in the correct position. But if the class is between 2 and 5, the first sentence corresponds to another position in the paragraph. We create three Spanish versions of different domains for this task using: the first five sentences of Wikipedia articles¹, Chilean university thesis abstracts², and news articles in Spanish from the MLSUM dataset (Scialom et al., 2020).

Figure 5 shows an example of this task for the thesis dataset. The first sentence should be in the second po-

¹We use the latest Spanish Wikipedia articles dump (dumps.wikimedia.org/eswiki/latest/)

²We collected abstracts from public repositories of the Pontificia Universidad Católica de Chile (repositorio.uc.cl) and Universidad de Chile (repositorio.uchile.cl).

sition among these sentences. To make correct predictions, the model needs to be aware of both typical orderings of events and how events are described in language. In the example shown, the model needs to understand that the objective of the thesis has to be described before the main findings of the study.

As proposed by Chen et al. (2019) to train the classifier for this task, we first encode the five sentences into vector representations x_i . As input to the classifier, we concatenate x_1 and $x_1 - x_i$ for $2 \leq i \leq 5$: $[x_1, x_1 - x_2, x_1 - x_3, x_1 - x_4, x_1 - x_5]$. The output is between 1 and 5, which indicates the target position of the first sentence.

- 1) Se encontró que la adición de nanopartículas de sílice aumenta la rigidez del material. (2)
- 2) El objetivo de este trabajo es estudiar el efecto de la incorporación de nanopartículas de sílice en la rigidez de material.
- 3) Las Nanopartículas de sílice fueron sintetizadas utilizando el método sol-gel.
- 4) Las Nanopartículas de menor tamaño tienen un mayor efecto sobre las propiedades del material.
- 5) La rigidez del material aumentó hasta en un 80% con la adición de 30% de nanopartículas de sílice.

Figure 5: SP example of thesis domain. The number inside the circle shows the correct position of the first sentence. This sentence belongs in the 2nd place.

3.2. Binary Sentence Ordering (BSO)

BSO is a binary classification task to determine if the order of two sentences is correct. This task aims to assess the ability of sentence representations to capture local discourse coherence. This data comes from the same three domains of the SP task. However, in this case, we only take the first two sentences of each text. Figure 6 provides an example from the Spanish Wikipedia. The order of the sentences is incorrect as the “Neue Pinakothek” museum should be mentioned before describing the art found inside. In order to find the incorrect ordering in this example, the sentence representations need to be able to provide information if one sentence comes after or before the separation. As English DiscoEval to create the model inputs to train the classifiers, we concatenate the embeddings generated by the sentence encoder of both sentences with their element-wise difference: $[x_1, x_2, x_1 - x_2]$.

- 1) Se centra en el Arte europeo del siglo XIX.
- 2) El Neue Pinakothek es un museo de arte situado en Múnich, Alemania.

Figure 6: Example from the Wikipedia domain of the BSO task. The sequence is in the wrong order.

3.3. Discourse Coherence (DC)

The Discourse Coherence (DC) task is a sentence disentanglement task proposed to determine if a sequence

of six sentences forms a coherent paragraph. We create three versions of this task, two from open-domain dialogue datasets and the other from Wikipedia articles. Given six coherent contiguous sentences, we randomly replace one of them with a sentence from another sequence. Note that we choose the sentence to replace uniformly among positions 2-5. We generate balanced datasets with coherent (positive) and non-coherent (negative) instances, which results in a binary classification task.

For the open-domain dialogue dataset, we use the OpenSubtitles³ corpus (Lison and Tiedemann, 2016) and the Gutenberg Dialogue dataset (Csaky and Recski, 2021). OpenSubtitles is a large corpus, so we randomly retrieve some dialogues and create the splits. In the case of the Gutenberg Dialogue, we use the original splits provided by the author. For the Wikipedia domain, we take only one coherent text from each article. Then we randomly create the splits. In all cases, we discard paragraphs with fewer than six sentences and we select the negative sample from other dialogues or articles in the corresponding domain. Figure 7 shows a dialog to which the second sentence does not belong. Like English DiscoEval, we encode the six sentences as vector representations and concatenate them $([x_1, x_2, x_3, x_4, x_5, x_6])$ as input to the classifier.

- 1) ¡nicolás, ha llegado tu hora!
- 2) **recuerdo que en la galería oscura me ofrecísteis vuestra casa.**
- 3) no; prefiero fumarne una pipa.
- 4) ¿dónde está tu pipa?
- 5) en el chaleco.
- 6) bien; aquí la tienes.

Figure 7: Example of DC from the Gutenberg. The sentence in bold does not belong to the dialogue.

3.4. Sentence Section Prediction (SSP)

SSP is a task to determine the section of a given sentence. This is based on the fact that the writing style can vary throughout a document, showing distinct patterns. The English DiscoEval originally used abstract and other sections of scientific papers to build the dataset. For our Spanish version, we use news articles instead. The news usually has a headline that is a sentence that presents the main idea of the article, a subhead that is a group of sentences that helps to encapsulate the entire piece or informs the reader about the topic, and a body that tells the entire story (Van Dijk, 1983).

We rely on the MLSUM dataset (Scialom et al., 2020), which consists of news articles that have the structure mentioned above. We use subhead and body sentences because the former has sentences summarizing the entire article, while the latter uses broader wording. Figure 8 shows an example of each style. We randomly

sample one sentence from the subhead as a positive instance and one sentence from the body as a negative sample. The task is a binary classification that takes the representation of the sentence as input.

Subhead: Los Reyes presiden este sábado el desfile de las Fuerzas Armadas
Body: Sevilla acoge este sábado el tradicional desfile de las Fuerzas Armadas, que estará presidido por los Reyes de España

Figure 8: Examples of SSP. One sentence is from the subhead, while the other is from the body of a news.

3.5. Discourse Relations (DR)

A direct way to test discourse knowledge is to predict the relations between sentences, which is why the RST Discourse Treebank (Carlson et al., 2001) was used in previous work (Ferracane et al., 2019; Chen et al., 2019). We consider the RST Spanish Treebank (da Cunha et al., 2011) for our Spanish version, which consists of an annotated corpus with rhetorical relations. According to RST (Mann and Thompson, 1988), a text can be segmented into Elementary Discourse Units (EDUs) linked by means of nucleus-satellite (NS) or multi-nuclear (NN) rhetorical relations. In the first, the satellite provides additional information about the nucleus, on which it depends (e.g., Fondo, Condición). In the second, several nuclei elements are connected at the same level, so no element is dependent on any other (e.g., Unión, Lista). For instance, Figure 9 shows an example with a relation NS and NN. A relation can take multiple units, so like Chen et al. (2019), we rely on right-branching trees for non-binary relations to binarize the tree structure and use the 29 coarse-grained relations defined by da Cunha et al. (2011). We adopt the originally proposed training and testing splits.




Figure 9: An RST Spanish Treebank tree with nucleus-satellite (NS) and multi-nuclear (NN) relations.

To evaluate the representations, we first encode all EDUs. We then use averaged EDU representations of subtrees as inputs. Formally, the input to the classifier is $[x_{left}, x_{right}, x_{left} \odot x_{right}, |x_{left} - x_{right}|]$ and the label

³<http://www.opensubtitles.org/>

Models	SentEval				DiscoEval				
	SC	SPC	SS	LPT	SP	BSO	DC	SSP	DR
Sent2Vec	<u>75.11</u>	59.51	76.05	66.89	36.49	54.92	55.77	70.88	36.69
ELMo	71.50	61.62	62.06	<u>69.90</u>	37.13	55.13	58.68	72.60	45.14
ELECTRA	62.80	51.40	42.07	64.20	38.56	56.85	55.18	76.22	37.59
RoBERTa-BNE	72.51	54.57	41.34	68.22	<u>41.82</u>	57.02	56.31	76.83	39.21
BERTIN	73.54	55.47	32.53	67.72	41.66	56.66	55.54	78.42	45.86
BETO	76.34	58.17	55.37	69.38	41.43	<u>57.53</u>	60.89	75.33	<u>47.84</u>
mBERT	70.47	<u>60.05</u>	<u>67.77</u>	71.41	43.21	57.97	63.45	<u>77.80</u>	51.08

Table 1: Results for Spanish SentEval and Spanish DiscoEval by group. The best performing model is in bold, and the runner up method is underlined. The reported numbers are accuracy, except SS, which is Pearson’s correlation.

is the relation of the node. x_{left} and x_{right} are vectors of the left and right subtrees respectively. For instance, the input for the label “NS-Elaboración” from Figure 9 is $x_{left} = x_1$ and $x_{right} = \frac{x_2 + x_3}{2}$.

4. Experiments

4.1. Setup

Parameters We adopt and adapt the original implementation of SentEval and DiscoEval for our Spanish version, so the same hyperparameters can be set. We use the PyTorch version of the classifiers, Adam optimizer with a batch size of 64, and 4 training epochs for all of the experiments.

Datasets We tokenize each dataset with the spaCy tokenizer (Honnibal and Montani, 2017) and save all files using a common file format with UTF-8 encoding.

4.2. Models

We benchmark all of the main Spanish sentence encoders available to date to the best of our knowledge. Sent2Vec⁴ (Pagliardini et al., 2018) which is a bilinear model, ELMo (Che et al., 2018) which is based on bidirectional RNNs. More recent and based on Transformer (Vaswani et al., 2017), we evaluate BETO (Cañete et al., 2020), the Spanish version of BERT (Devlin et al., 2019), BERTIN⁵ and RoBERTa-BNE (Gutiérrez-Fandiño et al., 2021), two versions in Spanish of the RoBERTa model (Liu et al., 2019). Finally, ELECTRA⁶ (Clark et al., 2020) that was trained on a small piece of data as part of a tutorial. We also include the multilingual BERT (mBERT) for further comparison. We use the base version for all models except ELECTRA, which is a small version.

For evaluating Sent2Vec and ELMo, we use their final representation. For the Transformer-based models, as proposed by Chen et al. (2019), we use the average of each layer’s special tokens [CLS] as the sentence representation.

4.3. Results

Table 1 shows the results of the experiments for all of the Spanish SentEval and Spanish DiscoEval tasks av-

eraged for all of the datasets used for each of the tasks (for detailed results, see Appendix B). It can be seen that, in general, the evaluation of all the language models’ latent representations in both the Spanish SentEval and Spanish DiscoEval tasks show a similar behavior compared to their English language representations counterparts (Conneau and Kiela, 2018; Chen et al., 2019). Regarding Spanish SentEval, for the SC tasks, the BETO model latent representation surpasses Sent2Vec, the second-best, by a 1.63 percentage difference (pd). This can be explained since BETO was trained with general domain sentences, capturing a representation capable of generalizing for any domain in the classification task. In general, SPC shows worse results than SC in terms of accuracy for all language models, where the best achieves an accuracy of 61.62%. In the SPC task, it can be seen that ELMo learned representation surpasses the second-best representation (mBERT) by 2.6 pd. The SS task shows that Sent2Vec representation outperforms other representations by more than 12 pd in terms of Pearson’s correlation, indicating that this learned representation can distinguish if a pair of sentences are semantically similar better than the representation learned by mBERT, which was not trained initially for this particular task. Finally, for the LPT task, the mBERT learned representation outperforms the BETO representation by 2.2 pd, showing that, when training a multi-language language model such as mBERT, the model can obtain richer sentence representations for a task that is more challenging than standard text classification.

Concerning the Spanish DiscoEval set of tasks, the SP task arises as one of the most challenging tasks, where mBERT, which is the best performing learned representation, reaches only a 43.21% accuracy, improving by 3.33 pd over RoBERTa-BNE, the runner up model. We observe a similar pattern in the BSO task since the representation learned by mBERT outperforms the second-best model representation by a low margin of 0.7 pd, showing that training a Transformer-based model in multiple languages can obtain a richer representation for the task of ordering two sentences in a paragraph. A similar behavior is observed for the DC task, where the mBERT representation outperforms the runner-up method by more than 4 pd. The SSP task results indicate that the BERTIN learned representation

⁴<https://github.com/BotCenter/spanish-sent2vec>

⁵<https://huggingface.co/bertin-project>

⁶<https://chriskhanhtran.github.io/posts/electra-spanish/>




Figure 10: Performance of DiscoEval and SentEval tasks using the [CLS] representation of layers 1 through 12.

surpasses mBERT by a small margin of 0.7 pd. Finally, for the RST task, mBERT representation shows the best performance in terms of accuracy compared to other language models, outperforming the BETO representa-

tion by 6.7 pd. In general, it is observed that in most of the DiscoEval tasks, mBERT learned the best representation. The exception is for SSP, but the difference in accuracy compared to representation learned by BETO

is small. These results provide evidence that mBERT learns better representations when trained with multiple languages, allowing it to outperform other models on most probing tasks.

4.4. Further Analysis

In this section, we perform a per-layer performance analysis of the representations learned by transformer-based models. These experiments allow verifying which layers are more transferable for downstream tasks. Figure 10 shows the results for all SentEval and DiscoEval groups.

It can be seen that the best performance fluctuates in the last layers, primarily between layers 10 and 12. Moreover, all representations perform well on the early layers for the SSP task, with accuracy levels near 0.7, indicating it is relatively straightforward. Nevertheless, all representations do not yield competitive performance for the SP task reaching a maximum accuracy slightly higher than 0.4, suggesting that they are ineffective at finding positions of a sentence in a discourse. Furthermore, something similar occurs with the DR task, where all representations achieve accuracy close to 0.5 for the last layers, showing that discovering relations between elements of discourse seems non-trivial to solve with the learned latent representations. Regarding the impact of the training data, we see that the representations generalize better when being trained with multiple languages than when using only Spanish text. Evidence of this is given by the performance of the representations learned by mBERT. In most cases, it outperforms other models' representations on several probing tasks. However, BETO representation beats the ones learned by mBERT in the last layers for SC, SS, and SSP, suggesting that for these tasks, representations learned only with Spanish texts seem to be more critical for obtaining an informative latent representation.

Another factor that positions mBERT and BETO as the two best-learned representations is that both were trained with more data, implying a better performance than ELECTRA and BERTIN, which were trained with fewer data. Interestingly, representations learned by RoBERTa-BNE do not get the desired performance compared to other representations, particularly on the early layers and on the last layers on tasks such as DC, DR, SSP, SC, and SPC.

5. Related Work

5.1. Language Model Evaluations

We can find at least four approaches in the work carried out in the evaluation of language models. The first focuses on evaluating the adaptability of a language model to a new domain through fine-tuning. GLUE (Wang et al., 2018) and SuperGLUE (Wang et al., 2019) are examples of this approach that include several downstream tasks. The second involves evaluating the generalization of text representations by in-

corporating a classifier for downstream tasks on top of them. Following this approach, SentEval (Conneau and Kiela, 2018) and DiscoEval (Chen et al., 2019) include tasks at the sentence and discourse level. The third focuses on stress tests (Naik et al., 2018; Aspillaga et al., 2020; Araujo et al., 2021a) that seek to assess the ability of language models to adapt to cases designed to confuse them. The fourth objective is an evaluation from a linguistic perspective (Warstadt et al., 2019; Etinger, 2020; Puccetti et al., 2021) to elucidate the models' actual linguistic capacities or knowledge.

The aforementioned benchmarks are scarce for languages other than English. This, in fact, is the case for Spanish. For instance, regarding the adaptability evaluation for Spanish models, Cañete et al. (2020) recently proposed GLUES, a Spanish version of GLUE. In the case of representation evaluation, most of the work is in a cross-linguistic setting for word (Şahin et al., 2020), sentence (Ravishankar et al., 2019) and discourse (Koto et al., 2021) evaluations. For this reason and following the motivation of works such as RuSentEval (Mikhailov et al., 2021), we provide SentEval and DiscoEval in Spanish, which consists of tasks originally created with texts in Spanish and aimed at evaluating models of that language.

5.2. Sentence Encoders

Pre-trained self-supervised language models have become the de facto sentence encoders. Early work in deep learning introduced ELMo (Peters et al., 2018). With this model, sentence representations are produced by a mean-pooling of all contextualized word representations. After the Transformer model (Vaswani et al., 2017), several models were proposed (Devlin et al., 2019; Liu et al., 2019; Clark et al., 2020). These BERT-type models produce sentence representations using a special token [CLS]. More recently, some models (Lee et al., 2020; Iyer et al., 2020; Araujo et al., 2021b) have been proposed to improve discourse-level representations by incorporating additional components or mechanisms into the vanilla BERT.

Furthermore, due to the success of deep learning sentence encoders, some Spanish models were released. Che et al. (2018) released ELMo for many languages, including Spanish. BETO (Cañete et al., 2020) the Spanish version of BERT (Devlin et al., 2019) was trained on a large Spanish corpus. RoBERTa-BNE (Gutiérrez-Fandiño et al., 2021), the Spanish version of the RoBERTa model (Liu et al., 2019), was trained on a corpus of crawled .es domains.

6. Conclusion

We introduce Spanish SentEval and Spanish DiscoEval, two test suites for evaluating stand-alone and discourse-aware sentence representations. Like the English versions, our work aims to evaluate the representations of current and future Spanish language models. Our benchmarks consist of a single pipeline that attempts a fair and less cumbersome assessment across

multiple tasks with text from different domains. As future work, more tasks could be included in these benchmarks. Likewise, other types of evaluations such as stress or linguistic tests could be carried out to evaluate the actual capacities of the language models taking into account the peculiarities of the Spanish language.

7. Acknowledgements

We want to thank the authors of the datasets for providing us access to their data. This work was supported by the National Center for Artificial Intelligence CENIA FB210017, Basal ANID. Felipe Bravo-Marquez was supported by ANID FONDECYT grant 11200290, U-Inicia VID Project UI-004/20 and ANID -Millennium Science Initiative Program - Code ICN17.002.

8. Bibliographical References

- Á. García Cumberas, M., Ureña López, L. A., and Martínez Santiago, F. (2006). BRUJA: Question classification for Spanish. using machine translation and an English classifier. In *Proceedings of the Workshop on Multilingual Question Answering - MLQA '06*.
- Agirre, E., Banea, C., Cardie, C., Cer, D., Diab, M., Gonzalez-Agirre, A., Guo, W., Mihalcea, R., Rigau, G., and Wiebe, J. (2014). SemEval-2014 task 10: Multilingual semantic textual similarity. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 81–91, Dublin, Ireland, August. Association for Computational Linguistics.
- Agirre, E., Banea, C., Cardie, C., Cer, D., Diab, M., Gonzalez-Agirre, A., Guo, W., Lopez-Gazpio, I., Maritxalar, M., Mihalcea, R., Rigau, G., Uribe, L., and Wiebe, J. (2015). SemEval-2015 task 2: Semantic textual similarity, English, Spanish and pilot on interpretability. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 252–263, Denver, Colorado, June. Association for Computational Linguistics.
- Araujo, V., Carvallo, A., Aspillaga, C., Thorne, C., and Parra, D. (2021a). Stress test evaluation of biomedical word embeddings. In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 119–125, Online, June. Association for Computational Linguistics.
- Araujo, V., Villa, A., Mendoza, M., Moens, M.-F., and Soto, A. (2021b). Augmenting BERT-style models with predictive coding to improve discourse-level representations. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3016–3022, Online and Punta Cana, Dominican Republic, November. Association for Computational Linguistics.
- Artetxe, M., Labaka, G., and Agirre, E. (2020). Translation artifacts in cross-lingual transfer learning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7674–7684, Online, November. Association for Computational Linguistics.
- Aspillaga, C., Carvallo, A., and Araujo, V. (2020). Stress test evaluation of transformer-based models in natural language understanding tasks. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1882–1894, Marseille, France, May. European Language Resources Association.
- Carlson, L., Marcu, D., and Okurovsky, M. E. (2001). Building a discourse-tagged corpus in the framework of Rhetorical Structure Theory. In *Proceedings of the Second SIGdial Workshop on Discourse and Dialogue*.
- Cañete, J., Chaperon, G., Fuentes, R., Ho, J.-H., Kang, H., and Pérez, J. (2020). Spanish pre-trained bert model and evaluation data. *PML4DC at ICLR 2020*.
- Cer, D., Diab, M., Agirre, E., Lopez-Gazpio, I., and Specia, L. (2017). SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada, August. Association for Computational Linguistics.
- Che, W., Liu, Y., Wang, Y., Zheng, B., and Liu, T. (2018). Towards better UD parsing: Deep contextualized word embeddings, ensemble, and treebank concatenation. In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 55–64, Brussels, Belgium, October. Association for Computational Linguistics.
- Chen, M., Chu, Z., and Gimpel, K. (2019). Evaluation benchmarks and learning criteria for discourse-aware sentence representations. 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 649–662, Hong Kong, China, November. Association for Computational Linguistics.
- Clark, K., Luong, M.-T., Le, Q. V., and Manning, C. D. (2020). Electra: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*.
- Conneau, A. and Kiela, D. (2018). SentEval: An evaluation toolkit for universal sentence representations. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May. European Language Resources Association (ELRA).
- Conneau, A., Kruszewski, G., Lample, G., Barrault, L., and Baroni, M. (2018a). What you can cram into a single $\$ \& ! \# *$ vector: Probing sentence embeddings for linguistic properties. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2126–2136, Melbourne, Australia, July. Association for Computational Linguistics.

- Conneau, A., Rinott, R., Lample, G., Williams, A., Bowman, S., Schwenk, H., and Stoyanov, V. (2018b). XNLI: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Cruz Mata, F., Troyano Jiménez, J. A., Enríquez de Salamanca Ros, F., and Ortega Rodríguez, F. J. (2008). Clasificación de documentos basada en la opinión: experimentos con un corpus de críticas de cine en español. *Procesamiento del lenguaje natural*. N. 41 (septiembre 2008); pp. 73–80.
- Csaky, R. and Recski, G. (2021). The Gutenberg dialogue dataset. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 138–159, Online, April. Association for Computational Linguistics.
- da Cunha, I., Torres-Moreno, J.-M., and Sierra, G. (2011). On the development of the RST Spanish treebank. In *Proceedings of the 5th Linguistic Annotation Workshop*, pages 1–10, Portland, Oregon, USA, June. Association for Computational Linguistics.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Etcheverry, M. and Woneser, D. (2016). Spanish word vectors from Wikipedia. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3681–3685, Portorož, Slovenia, May. European Language Resources Association (ELRA).
- Ettinger, A. (2020). What BERT Is Not: Lessons from a New Suite of Psycholinguistic Diagnostics for Language Models. *Transactions of the Association for Computational Linguistics*, 8:34–48, 01.
- Ferracane, E., Durrett, G., Li, J. J., and Erk, K. (2019). Evaluating discourse in structured text representations. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 646–653, Florence, Italy, July. Association for Computational Linguistics.
- Gutiérrez-Fandiño, A., Armengol-Estapé, J., Pàmies, M., Llop-Palao, J., Silveira-Ocampo, J., Carrino, C. P., Gonzalez-Agirre, A., Armentano-Oller, C., Rodriguez-Penagos, C., and Villegas, M. (2021). Spanish language models. *arXiv preprint arXiv:2107.07253*.
- Honnibal, M. and Montani, I. (2017). spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing.
- Huertas-Tato, J., Martín, A., and Camacho, D. (2022). Silt: Efficient transformer training for inter-lingual inference. *Expert Systems with Applications*, page 116923, March.
- Iter, D., Guu, K., Lansing, L., and Jurafsky, D. (2020). Pretraining with contrastive sentence objectives improves discourse performance of language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4859–4870, Online, July. Association for Computational Linguistics.
- Koto, F., Lau, J. H., and Baldwin, T. (2021). Discourse probing of pretrained language models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3849–3864, Online, June. Association for Computational Linguistics.
- Lee, H., Hudson, D. A., Lee, K., and Manning, C. D. (2020). SLM: Learning a discourse language representation with sentence unshuffling. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1551–1562, Online, November. Association for Computational Linguistics.
- Lison, P. and Tiedemann, J. (2016). OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 923–929, Portorož, Slovenia, May. European Language Resources Association (ELRA).
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Mann, W. C. and Thompson, S. A. (1988). Rhetorical structure theory: Toward a functional theory of text organization. *Text - Interdisciplinary Journal for the Study of Discourse*, 8(3).
- Mikhailov, V., Taktasheva, E., Sigdel, E., and Artemova, E. (2021). RuSentEval: Linguistic source, encoder force! In *Proceedings of the 8th Workshop on Balto-Slavic Natural Language Processing*, pages 43–65, Kiyv, Ukraine, April. Association for Computational Linguistics.
- Naik, A., Ravichander, A., Sadeh, N., Rose, C., and Neubig, G. (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, August. Association for Computational Linguistics.
- Pagliardini, M., Gupta, P., and Jaggi, M. (2018). Unsupervised learning of sentence embeddings using

- compositional n-gram features. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 528–540, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. (2018). 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)*, pages 2227–2237, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Plaza del Arco, F. M., Strapparava, C., Urena Lopez, L. A., and Martin, M. (2020). EmoEvent: A multilingual emotion corpus based on different events. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1492–1498, Marseille, France, May. European Language Resources Association.
- Puccetti, G., Miaschi, A., and Dell’Orletta, F. (2021). How do BERT embeddings organize linguistic knowledge? In *Proceedings of Deep Learning Inside Out (DeeLIO): The 2nd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 48–57, Online, June. Association for Computational Linguistics.
- Ravishankar, V., Øvrelid, L., and Veldal, E. (2019). Probing multilingual sentence representations with X-probe. In *Proceedings of the 4th Workshop on Representation Learning for NLP (ReplANLP-2019)*, pages 156–168, Florence, Italy, August. Association for Computational Linguistics.
- Scialom, T., Dray, P.-A., Lamprier, S., Piwowarski, B., and Staiano, J. (2020). MLSUM: The multilingual summarization corpus. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8051–8067. ACL.
- Sobrevilla Cabezudo, M. A., La Serna Palomino, N., and Perez, R. M. (2015). Improving subjectivity detection for spanish texts using subjectivity word sense disambiguation based on knowledge. In *2015 Latin American Computing Conference (CLEI)*, pages 1–7.
- Van Dijk, T. A. (1983). Discourse analysis: Its development and application to the structure of news. *Journal of Communication*, 33(2):20–43.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. (2017). Attention is all you need. In I. Guyon, et al., editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Vega, M. G., Díaz-Galiano, M. C., Cumbleras, M. Á. G., del Arco, F. M. P., Montejo-Ráez, A., Zafra, S. M. J., Cámara, E. M., Aguilar, C. A., Cabezudo, M. A. S., Chiruzzo, L., and Moctezuma, D. (2020). Overview of TASS 2020: Introducing emotion detection. In Miguel Ángel García Cumbleras, et al., editors, *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2020) co-located with 36th Conference of the Spanish Society for Natural Language Processing (SEPLN 2020), Málaga, Spain, September 23th, 2020*, volume 2664 of *CEUR Workshop Proceedings*, pages 163–170.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. (2018). 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*, pages 353–355, Brussels, Belgium, November. Association for Computational Linguistics.
- Wang, A., Pruksachatkun, Y., Nangia, N., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. (2019). Superglue: A stickier benchmark for general-purpose language understanding systems. In H. Wallach, et al., editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Warstadt, A., Cao, Y., Grosu, I., Peng, W., Blix, H., Nie, Y., Alsop, A., Bordia, S., Liu, H., Parrish, A., Wang, S.-F., Phang, J., Mohanane, A., Htut, P. M., Jeretic, P., and Bowman, S. R. (2019). Investigating BERT’s knowledge of language: Five analysis methods with NPIs. 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 2877–2887, Hong Kong, China, November. Association for Computational Linguistics.
- Yang, Y., Zhang, Y., Tar, C., and Baldrige, J. (2019). PAWS-X: A cross-lingual adversarial dataset for paraphrase identification. 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 3687–3692, Hong Kong, China, November. Association for Computational Linguistics.
- Şahin, G. G., Vania, C., Kuznetsov, I., and Gurevych, I. (2020). LINSPECTOR: Multilingual Probing Tasks for Word Representations. *Computational Linguistics*, 46(2):335–385, 06.

Appendix A: Details of Datasets

Name	N	Description	C
TASS1	14.4k	Polarity (2020 joint task 1)	3
TASS2	8.4k	Emotion (2020 task 2)	7
MC	2.6k	Sentiment (Movie reviews)	2
FAC	5k	Subjectivity/Objectivity	2
SQC	6.7k	Question-type	6
PAWS-es	51.4k	Paraphrase	2
NLI-es	9.9k	Entailment	2
SICK-es	9.8k	Entailment/Relatedness	2
STS14	804	Semantic Similarity	score
STS15	751	Semantic Similarity	score
STS17	250	Semantic Similarity	score

Name	N	Description	C
SP wiki/mlsum	18k	Sentence position	5
SP thesis	14k	Sentence position	5
BSO wiki/mlsum	18k	Sentence ordering	2
BSO thesis	14k	Sentence ordering	2
DC wiki/opus	18k	Coherence	2
DC gdd	6.6k	Coherence	2
SSP mlsum	18k	Section prediction	2
DR rst	3.3k	Rhetorical structure theory	29

Table 2: Details of the SentEval and DiscoEval datasets. N shows the number of instances and C the number of classes.

Appendix B: Full Results of Spanish Models

Models	SC					SPC			SS			
	TASS 1	TASS 2	MC	FAC	SQC	PAWS -es	NLI -es	SICK -es-E	SICK -es-R	STS 14	STS 15	STS 17
Sent2Vec	60.91	65.47	75.94	95.46	77.78	56.15	48.35	74.03	78.98	83.12	60.02	82.10
ELMo	56.07	62.40	66.91	94.66	77.48	58.8	49.28	76.78	75.56	69.76	45.88	57.02
ELECTRA	58.04	54.83	65.61	94.18	41.33	55.90	36.39	61.90	56.40	54.29	30.85	26.74
RoBERTa-BNE	64.10	64.62	75.11	96.66	62.07	56.30	40.60	66.82	67.79	57.11	34.13	6.31
BERTIN	59.20	61.86	74.46	95.90	76.30	56.75	42.45	67.20	51.83	31.24	18.67	28.38
BETO	62.95	64.44	77.01	97.60	79.70	57.65	46.39	70.46	63.24	65.84	47.36	45.01
mBERT	56.28	58.02	70.00	96.92	71.11	57.50	48.63	74.01	67.77	62.14	43.96	36.71

Table 3: Results of the SentEval tasks for each dataset (SC, SPC, SS).

Models	LPT								
	SentLen	WC	Tree Depth	BiShift	Tense	SubjNum	ObjNum	SOMO	CoordInv
Sent2Vec	73.99	4.98	41.25	81.76	87.72	83.49	78.50	51.64	74.44
ELMo	44.22	66.43	49.84	78.47	94.65	87.27	80.21	53.76	74.25
ELECTRA	73.99	4.98	41.25	81.76	87.72	83.49	78.50	51.64	74.44
RoBERTa-BNE	77.27	22.60	44.80	77.07	93.63	87.15	79.65	51.98	79.83
BERTIN	75.33	20.84	45.02	73.40	93.36	89.71	73.92	53.80	84.06
BETO	68.85	41.60	45.42	75.72	93.26	87.30	77.99	51.03	83.27
mBERT	75.41	41.27	46.50	76.34	95.30	89.57	78.22	54.10	86.02

Table 4: Results of the SentEval tasks for each dataset (LPT).

Models	SP			BSO			DC			SSP	DR
	wiki	mlsum	thesis	wiki	mlsum	thesis	wiki	opus	gdd	mlsum	rst
Sent2Vec	42.77	30.40	36.29	51.56	50.46	62.75	66.67	49.85	50.78	70.88	36.69
ELMo	44.50	29.88	37.01	51.42	50.89	63.07	68.73	52.62	54.69	72.60	45.14
ELECTRA	45.65	31.45	38.57	52.44	52.83	65.29	57.88	50.62	57.03	76.22	37.59
RoBERTa-BNE	47.58	35.10	42.78	52.99	51.04	67.04	68.12	50.80	50.00	76.83	39.21
BERTIN	47.90	33.38	43.71	52.09	49.90	67.99	60.75	53.52	52.34	78.42	45.86
BETO	48.90	32.35	43.03	53.50	51.94	67.15	75.62	53.15	53.91	75.33	47.84
mBERT	51.95	33.23	44.46	53.79	51.75	68.36	80.00	53.33	57.03	77.80	51.08

Table 5: Results of the DiscoEval tasks for each dataset.