

# Sequence-to-Sequence Spanish Pre-trained Language Models

Vladimir Araujo<sup>1</sup>, Maria Mihaela Trusca<sup>1</sup>, Rodrigo Tufiño<sup>2</sup>, Marie-Francine Moens<sup>1</sup>

<sup>1</sup>LIIR Lab, KU Leuven, Leuven, Belgium

<sup>2</sup>IDEIAGEOCA, Universidad Politécnica Salesiana, Quito, Ecuador

[vladimir.araujo@kuleuven.be](mailto:vladimir.araujo@kuleuven.be)

## Abstract

In recent years, significant advancements in pre-trained language models have driven the creation of numerous non-English language variants, with a particular emphasis on encoder-only and decoder-only architectures. While Spanish language models based on BERT and GPT have demonstrated proficiency in natural language understanding and generation, there remains a noticeable scarcity of encoder-decoder models explicitly designed for sequence-to-sequence tasks, which aim to map input sequences to generate output sequences conditionally. This paper breaks new ground by introducing the implementation and evaluation of renowned encoder-decoder architectures exclusively pre-trained on Spanish corpora. Specifically, we present Spanish versions of BART, T5, and BERT2BERT-style models and subject them to a comprehensive assessment across various sequence-to-sequence tasks, including summarization, question answering, split-and-rephrase, dialogue, and translation. Our findings underscore the competitive performance of all models, with the BART- and T5-based models emerging as top performers across all tasks. We have made all models publicly available to the research community to foster future explorations and advancements in Spanish NLP: <https://github.com/vgaraujov/Seq2Seq-Spanish-PLMs>.

**Keywords:** generative models, pre-trained language models, sequence-to-sequence models, transformer

## 1. Introduction

Spanish ranks among the most extensively used languages globally. This fact has captured the interest of the NLP community, prompting efforts towards resource development for this NLP domain. Consequently, a number of pre-trained language models tailored for Spanish have emerged in recent years, predominantly employing encoder-only (Cañete et al., 2020; De la Rosa et al., 2022; Araujo et al., 2023) and decoder-only (Gutiérrez-Fandiño et al., 2022) architectures. These models have demonstrated exemplary performance in natural language understanding across several downstream tasks (Cañete et al., 2020; Araujo et al., 2022). Nevertheless, there has been limited advancement in addressing tasks revolving around generating new sentences depending on a given input, such as summarization, generative question answering, dialogue, or machine translation.

Encoder-decoder models primarily serve for addressing sequence-to-sequence tasks, and over recent years, numerous architectures have emerged. The pre-training of these models is often based on the whole transformer architecture (Vaswani et al., 2017) and entails more intricate learning objectives than those of encoder or decoder-only models individually. For instance, BART (Lewis et al., 2020a) is specifically trained to reconstruct text that has been intentionally corrupted, while T5 (Raffel et al., 2020) is designed to adeptly fill in missing sections of text, simulating a scenario where text spans have been omitted. These models have been developed predominantly for the English language, and recent

efforts have been made to pre-train them in languages other than English (Kamal Eddine et al., 2021). Unfortunately, when it comes to the Spanish language, there is a notable scarcity of such models that may be valuable to the community.

In this paper, with the aim of democratizing sequence-to-sequence models for the Spanish NLP community, we introduce BARTO and T5S, which are the Spanish counterparts of the BART and T5 models. These models are exclusively pre-trained on Spanish corpora, aligning with their self-supervised methodology. We also introduce models in the style of BERT2BERT (Rothe et al., 2020), utilizing well-established BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) models for Spanish as baselines. Additionally, we curate a variety of Spanish sequence-to-sequence tasks, such as summarization, question answering, split-and-rephrase, dialogue, and machine translation, to assess our models comprehensively.

Our results demonstrate that all models perform competitively on the proposed benchmark tasks. Particularly, BARTO and T5S achieve the best performance, especially with lengthy sequences, surpassing BERT2BERT-style models and their multilingual counterparts in many generative tasks. Furthermore, we evaluate the models' performance on discriminative tasks, such as sequence and token classification. While their performance is slightly behind encoder-only models in some tasks, they still offer prominent results. To facilitate future research and practical applications in Spanish NLP, we have made these models available.

## 2. Related Work

### 2.1. Language-specific Pre-trained Language Models

Pre-trained language models represent a class of advanced language models trained through self-supervised learning on large text corpora, making them versatile for various applications. Notably, two prominent models are BERT (Devlin et al., 2019), an encoder-only model, and GPT (Radford and Sutskever, 2018; Radford et al., 2019), a decoder-only model. These models have established robust baselines for a wide range of NLP tasks in English. Numerous language-specific BERT-based and GPT-based models have emerged in recent times. Examples include CamemBERT (Martin et al., 2020) tailored for French, RobBERT (Delobelle et al., 2020) designed for Dutch, FinBERT (Virtanen et al., 2019) for Finnish, GePpeTto (Mattei et al., 2020) for Italian, and several others. These models have consistently outperformed their multilingual counterparts, highlighting the value of their existence for language-specific tasks.

In the context of the Spanish language, we find BETO (Cañete et al., 2020) and ALBETO (Cañete et al., 2022), a BERT and ALBERT model, respectively, pre-trained on the SUC corpus (Cañete et al., 2020). BILMA (Tellez et al., 2023), regionalized BERT models for Spanish language variations based on Twitter data. BERTIN (De la Rosa et al., 2022), a RoBERTa base model trained on the Spanish portion sampled from mC4 (Xue et al., 2021). Additionally, MarIA (Gutiérrez-Fandiño et al., 2022) introduces a family of models, including RoBERTa and GPT-2 trained on the corpus crawled by the National Library of Spain. Furthermore, RigoBERTa (Serrano et al., 2022), follows the DeBERTa (He et al., 2021) architecture and was trained with several corpora, including OSCAR (Suárez et al., 2019), SUC, and mC4-es. Nevertheless, a notable gap exists in the availability of encoder-decoder models exclusively trained for Spanish.

### 2.2. Sequence-to-Sequence Pre-trained Language Models

A sequence-to-sequence model aims to map a fixed-length input with a fixed-length output where the length of the input and output may differ (Sutskever et al., 2014). It comprises an encoder, which processes the entire input sequence, and a decoder, which receives the representations computed by the encoder and generates the output sequence in an autoregressive manner. These models have proven to be valuable in addressing tasks including summarization, dialogue, generative question answering, and translation.

Following the paradigm of pre-training with self-supervision, several models have been proposed. One of the first models is MASS (Song et al., 2019),

which uses a transformer to reconstruct an input sequence where a contiguous span of tokens is masked and mapped to a sequence consisting of the missing tokens. Later, T5 (Raffel et al., 2020) proposed a pre-train on a multitask combination of supervised and self-supervised tasks, the latter being a task to complete fill-in dropped-out spans of text from documents. BART (Lewis et al., 2020a) is slightly similar to T5 but only uses a self-supervised objective in which spans are masked from the input, but the complete output is predicted to improve the decoder’s language modeling ability. Moreover, Rothe et al. (2020) proposed the utilization of encoder or decoder-only pre-trained checkpoints for initializing new encoder-decoder models, showcasing competitive performance compared to purely encoder-decoder pre-trained models.

More recently, there has been a notable surge in endeavors to deploy sequence-to-sequence models for languages beyond English. BART, for instance, has been released for many other languages, including French (Kamal Eddine et al., 2021), Greek (Evdaimon et al., 2023), Indic (Dabre et al., 2022), Arabic (Kamal Eddine et al., 2022) and various other languages (Tran et al., 2022; La Quatra and Cagliero, 2023). Furthermore, T5 has been pre-trained in Portuguese (Carmo et al., 2020), Italian (Sarti and Nissim, 2022), Arabic (Nagoudi et al., 2022), Indic (Aralikatte et al., 2023), among others. While the aforementioned recent models encompass a broad range of languages, the availability of Spanish models remains limited.

### 3. Sequence-to-Sequence Spanish Pre-trained Language Models

In this section, we begin by presenting our data collection and preparation procedures for pre-training our models. Subsequently, we provide detailed descriptions of each model and outline the corresponding pre-training processes.

#### 3.1. Pre-training Data

We employ the OSCAR 21.09 corpus (Abadji et al., 2022), which includes a deduplicated Spanish set of approximately 160GB of text. Furthermore, we utilize the mC4-es corpus (Xue et al., 2021), specifically adopting the Gaussian perplexity sampling subset proposed by De la Rosa et al. (2022), which boasts an extensive 500GB text dataset and has demonstrated superior model consistency. Additionally, we incorporate SUC, the corpus utilized for pre-training BETO, comprising around 14GB of raw text from diverse sources. Note that we exclude Wikipedia data from SUC, instead opting for an updated Wikipedia dump<sup>1</sup> of approximately 10GB. As established by prior research (Liu et al., 2019; Raffel et al., 2020), the corpus quality significantly

<sup>1</sup><https://dumps.wikimedia.org/eswiki/latest/>

impacts the outcomes of pre-training models. Consequently, we closely follow the preprocessing methodologies previously established for both English (Raffel et al., 2020) and Spanish models (Gutiérrez-Fandiño et al., 2022; Serrano et al., 2022). Below, we describe the procedure:

1. **Document-level Formatting:** We ensure that all data adheres to a document-level format, which means that each instance is a document containing several contiguous coherent sentences. As demonstrated by Liu et al. (2019), restricting sequences to come from a single document performs slightly better than packing sequences from multiple documents. Furthermore, this causes the models to capture broad contextual dependencies.
2. **Data Filtering:** To enhance data quality, we employ straightforward and cost-effective filtering methods. We eliminate very short documents based on sentence and document length. We filter out text containing repeated characters or special characters not commonly used in Spanish. We exclude pages containing code and sensitive content. Moreover, we utilize the fastText language identification model (Joulin et al., 2017) to exclude documents classified with less than 98% accuracy for Spanish text. We set this threshold to ensure our corpus retains a small yet representative portion of other languages that are frequently interspersed in contemporary Spanish texts.
3. **Deduplication:** We employ a deduplication process across all corpora using the `text-dedup`<sup>2</sup> library. This generates a dataset smaller and faster to train on while potentially enhancing the resulting performance due to avoiding duplicated data (Lee et al., 2022). Due to its computational intensity, this step is performed after formatting and filtering.
4. **Encoding Correction:** Some documents may have inconsistent encodings or exhibit encoding issues. To address this, we utilize the `ftfy`<sup>3</sup> tool to rectify encoding mix-ups, ensuring UTF-8 encoding and NFKC normalization. We carry out this process at the end of the pipeline because, while `ftfy` is highly effective in text correction, it can be computationally expensive for large corpora. Additionally, this guarantees that we have proper encoding after all the preceding steps.

The resulting corpus size after preprocessing exceeds 120GB of uncompressed text, a scale similar to the one used for training RoBERTa and BART models (Liu et al., 2019; Lewis et al., 2020a).

<sup>2</sup><https://github.com/ChenghaoMou/text-dedup>

<sup>3</sup><https://ftfy.readthedocs.io/>

### 3.2. BARTO Model

BARTO follows the BART base architecture, which consists of an encoder and a decoder with 6 layers each. Also, it has 12 attention heads and 768 hidden dimensions in both the encoder and decoder. BARTO is pre-trained by denoising the corrupted input documents. As suggested by Lewis et al. (2020a), we use a combination of text infilling and sentence permutation transformations for robust performance, masking 30% of tokens in each document and permuting all sentences.

We use `sentencepiece` (Kudo and Richardson, 2018) to build a BPE tokenizer of 50,264 tokens. Furthermore, we rely on the `fairseq` library (Ott et al., 2019) to perform the training. BARTO is pre-trained for 100,000 steps on 8 NVIDIA A100 GPUs with input texts of 1024 and a batch size of 2048. We use the Adam optimizer (Kingma and Ba, 2015), a warm-up of 10,000 steps, a dropout of 0.1, and FP16 to speed up training.

### 3.3. T5S Model

T5S follows the T5.1.1<sup>4</sup> base version of the T5 model, which includes some improvements. This model consists of an encoder and decoder with 12 layers, 12 attention heads, and 768 hidden dimensions each. Like T5.1.1, we pre-train only using the denoising objective by filling in dropped-out spans of text from documents. As proposed by Raffel et al. (2020), we use a corruption rate of 15% and an average span length of 3.

We use the `sentencepiece` library to build a unigram tokenizer of 32,000 tokens. Also, we rely on `nanoT5` (Nawrot, 2023), which allows pre-training T5 models on a limited budget. Our T5S is pre-trained for 130,000 steps on 4 GPUs NVIDIA A100 with input texts of 1024 and a batch size of 320. Additionally, we use the AdamW optimizer (Loshchilov and Hutter, 2019), a warm-up of 10,000 steps, a dropout of 0, and BP16 to speed up training.

### 3.4. BERT2BERT-style Models

Our BERT2BERT-style models follow the procedure proposed by Rothe et al. (2020), which consists of initializing encoder-decoder models with pre-trained encoder and/or decoder-only checkpoints. We use two configurations: BERT2BERT, which is an encoder initialized by a BERT-type checkpoint paired with a decoder initialized with the same checkpoint, and BERTShare, which is similar to BERT2BERT but the parameters between the encoder and decoder are shared.

We rely on the `transformers`<sup>5</sup> library (Wolf et al., 2020) to initialize models based on two well-known

<sup>4</sup>[https://github.com/google-research/text-to-text-transfer-transformer/blob/main/released\\_checkpoints.md](https://github.com/google-research/text-to-text-transfer-transformer/blob/main/released_checkpoints.md)

<sup>5</sup>[https://huggingface.co/docs/transformers/model\\_doc/encoder-decoder](https://huggingface.co/docs/transformers/model_doc/encoder-decoder)

architectures. On the one hand, by using the BETO checkpoint, we initialize a BETO2BETO and BETOShare model. On the other hand, by leveraging the RoBERTa-BNE checkpoint from MarIA, we initialize a RoBERTa2RoBERTa and RoBERTaShare. Note that these models do not need to continue pre-training but rather fine-tune them directly in downstream tasks. We will delve into this process in more detail in the following section.

## 4. Evaluation

In this section, we introduce the downstream tasks chosen to evaluate the performance of our sequence-to-sequence models. These datasets primarily consist of generative tasks, where both input and output sequence texts are provided. Furthermore, we describe some discriminative tasks we adopt as part of our assessment. Finally, we provide details on the model fine-tuning process.

### 4.1. Generative Tasks

**Abstractive Summarization** Summarization involves creating a concise version of a document while retaining its key information. We consider ML-SUM (Scialom et al., 2020) and WikiLingua (Ladhak et al., 2020) datasets to evaluate our models. On the one hand, MLSUM is a collection of newspaper articles with an average number of  $\sim 900$  tokens<sup>6</sup> and summaries of approximately  $\sim 24$  tokens. On the other hand, WikiLingua consists of guide-based articles with around  $\sim 500$  tokens on average, and its summaries contain about  $\sim 50$  tokens. Note that both datasets do not contain overly lengthy texts, particularly in the case of the summaries. This factor simplifies the generative process significantly.

**Long-form Abstractive Summarization** A more challenging task of abstractive summarization is the processing and/or generation of long-form text. We use XL-Sum (Hasan et al., 2021), a dataset with lengthy articles, and EUR-Lex-Sum (Aumiller et al., 2022), a dataset with long summaries. XL-Sum (Hasan et al., 2021) contains long news articles with about  $\sim 1200$  tokens on average and short summaries of about  $\sim 40$  tokens. In contrast, EUR-Lex-Sum is a dataset of legal documents with an average length of  $\sim 19000$  tokens and lengthy summaries of approximately  $\sim 1200$  tokens. These datasets present more significant challenges compared to the previous ones. This is due to the more intensive nature of the encoding and generative process, which demands models capable of processing lengthy sequences and capturing extended dependencies to produce high-quality summaries.

**Split and Rephrase** The split-and-rephrase task assumes rewriting the content of a long sentence into shorter and less verbose sentences. We use

the Spanish subset of the BiSECT dataset (Kim et al., 2021) to evaluate this task, which contains about 290,000 instances. The average number of tokens within the input sentences is approximately  $\sim 51$ , while after rephrasing into two sentences, the average increases to  $\sim 75$  tokens across a pair of sentences. Following Kim et al. (2021), we use a special token to separate split sentences during model fine-tuning. Specifically, we use the special tokens  $\langle s \rangle$  or  $[CLS]$  when using BARTO or BERT2BERT-style models, empirically providing the best performance. Although T5S lacks the  $\langle s \rangle$  token, this configuration leads to the best result.

**Generative Question Answering** This task focuses on generating an abstractive answer to a given question from a provided passage. To the best of our knowledge, there is currently no dataset tailored for abstractive question answering in Spanish. In line with prior work (Raffel et al., 2020), we utilize span-based question answering datasets to train the models to generate the correct answers rather than predicting the specific token positions of the answer. We rely on MLQA (Lewis et al., 2020b) and SQAC (Gutiérrez-Fandiño et al., 2022) datasets for this evaluation. MLQA presents a collection of parallel multi-lingual articles extracted from Wikipedia and offers a development set and test set professionally translated into Spanish. Unlike MLQA, SQAC was proposed exclusively for Spanish evaluation and contains articles extracted from Spanish sources. Following the BART fine-tuning procedure (Lewis et al., 2020a), models generate answers conditioned on the concatenation of questions and supporting documents.

**Dialogue** The dialogue response generation task aims to produce an appropriate and coherent response based on the dialogue context. We employ the Spanish partition of the MIAM dataset (Colombo et al., 2021), a benchmark comprising dialogue act corpora. This dataset proposed initially to identify the specific act that a speaker performs is not the task we want to address in this work. Therefore, we adapt the data to suit our evaluation needs. Following previous work (Zhou et al., 2021; Zhang et al., 2020), we focus solely on the utterances exchanged between the speakers to create pairs of dialogue context and responses.

**Machine Translation** The objective of this task is to translate a sentence from a source language to a different target language. We rely on Fapesp-v2 (Aziz and Specia, 2011) and WMT13 (Bojar et al., 2013), well-known machine translation benchmarks that encompass Spanish-language data. On the one hand, Fapesp-v2 is a Portuguese  $\leftrightarrow$  Spanish parallel corpora crawled from a Brazilian magazine with about  $\sim 150,000$  examples of training sets and  $\sim 1,300$  examples of development and test

<sup>6</sup>Token counting is done using the BETO tokenizer.

sets. On the other hand, WMT13 includes an English  $\leftrightarrow$  Spanish subset comprising approximately 15 million training instances, which is extensive for conducting small-scale experiments. Therefore, we opted to work with a randomly sampled subset of 600,000 examples, a size similar to that used for fine-tuning English models (Lewis et al., 2020a). For evaluation, we use the newstest2012 and newstest2013 sets, each containing  $\sim$ 3000 examples.

## 4.2. Discriminative Tasks

**GLUES** We rely on GLUES, a benchmark introduced initially by Cañete et al. (2020). We focus our evaluation on sequence classification tasks: MLDoc, PAWS-X, and XNLI. MLDoc revolves around the classification of long documents into four distinct categories. PAWS-X, on the other hand, centers on the identification of sentence paraphrases. Lastly, XNLI consists of predicting whether a premise logically entails a given hypothesis. Additionally, we use semantic similarity and relatedness tasks, including STS-es (Gutiérrez-Fandiño et al., 2022) and SemRel2024 (Ousidhoum et al., 2024). STS-es entails evaluating the similarity between two text segments by assigning a score from 1 to 5. SemRel2024 assesses the degree of semantic textual relatedness between two sentences, assigning a score between 0 and 1.

**SQAC** In the case of token-level classification, we rely on the SQAC dataset (Gutiérrez-Fandiño et al., 2022). This corpus is a compilation of question-answer pairs extracted from Wikipedia articles. Within SQAC, the primary task involves predicting the specific span within the text that corresponds to the answer to a given question. This is achieved by indicating the start and end positions of this answer span within the text. As a common practice, the input to the model during fine-tuning is concatenating the question and the contextual text.

## 4.3. Fine-tuning

We follow the fine-tuning procedures proposed for the English version models (Lewis et al., 2020a; Raffel et al., 2020). Because BARTO and T5S have an autoregressive decoder, they can be directly fine-tuned for sequence generation tasks. Specifically, their encoders take a complete input, and then their decoders generate a target output autoregressively. For BERT2BERT-style models, we initialize a transformer with a checkpoint in both the encoder and decoder. Note that the decoder has cross-attention layers that are randomly initialized since the checkpoint does not have these parameters. Subsequently, we fine-tune these models following the same procedure as BARTO and T5S.

We fine-tune the models on an RTX 3090 GPU for each task using the `transformers` library implemented in PyTorch. For a fair comparison, we use

the same hyperparameters with the exception of the batch size, learning rate, and the number of training epochs. The optimal settings may depend on the task, therefore we consider a hyperparameter sweep with batch size  $\in$  {4, 8, 16}, learning rate (AdamW)  $\in$  {3e-5, 5e-5}, and epochs  $\in$  {3, 6}.

## 4.4. Multilingual Baselines

To complement our evaluation, we include multilingual sequence-to-sequence baselines for further comparison. We use mT5-base (Xue et al., 2021), a model trained on 101 languages. While there is no base version of mBART directly comparable to our models, we include a mBART-large version for comparison purposes. Specifically, we use mBART-50 (Tang et al., 2020), which includes both Spanish and Portuguese, essential for our experiments. We fine-tune these models following the same procedure as for BARTO and T5S.

## 5. Results

This section presents the results achieved after fine-tuning all the models on the downstream tasks. We evaluate them using conventional metrics, adhering to established methodologies in previous work (Lewis et al., 2020a; Raffel et al., 2020).

### 5.1. Generative Tasks

**Abstractive Summarization** Table 1 presents a comparison of the results achieved by all the models on the MLSUM and WikiLingua tasks, measured in terms of ROUGE metrics (Lin, 2004). Specifically, the T5S model shows the highest performance with an average of 26.54 ROUGE among the Spanish models across all tasks, while BARTO is second with a difference with an average of 25.49 ROUGE. As for the BERT2BERT-style models, they are all outperformed. BETO2BETO is the one that offers the best performance among its group of models, with an average of 24.39 ROUGE. Notably, T5S and BARTO outperform their multilingual counterparts, with the exception of mBART on WikiLingua, which slightly outperforms T5S by 0.45 ROUGE.

**Long-form Abstractive Summarization** The results for long-form summarization are presented in Table 2. In XLSum, T5S slightly outperforms BARTO, averaging 22.06 and 21.58 ROUGE, respectively. Furthermore, BETOShare achieves the best performance of its group of models with an average of 19.82 ROUGE. Regarding EUR-Lex-Sum, BARTO outperforms T5S, averaging 56.82 and 56.42 ROUGE, respectively. Notably, the difference between BARTO and T5S in comparison to the BERT2BERT-style models is significant, being BETO2BETO the best performer with an average of 26.69 ROUGE. These results reflect the unique requirements of EUR-Lex-Sum, which involve processing lengthy articles and generating

	MLSUM			WikiLingua		
	R1	R2	RL	R1	R2	RL
BETO2BETO	28.46/28.09	10.87/10.34	22.89/22.51	38.02/37.92	17.57/17.43	29.38/29.24
BETOShare	28.51/27.84	10.90/10.19	22.99/22.30	37.74/37.68	17.41/17.29	29.19/29.03
RoBERTa2RoBERTa	27.94/27.69	9.66/9.25	21.92/22.07	35.68/35.58	14.53/14.51	26.49/26.37
RoBERTaShare	28.43/27.86	10.17/9.53	22.54/21.92	35.83/35.70	14.95/14.75	26.85/26.62
mBART-large	29.11/28.56	10.85/10.21	22.62/22.00	41.15/40.98	20.44/20.40	33.24/33.01
mT5	29.33/28.69	11.53/10.89	23.91/23.24	37.27/37.20	18.57/18.48	31.39/31.25
BARTO	29.65/29.12	10.96/10.32	22.71/22.18	39.48/39.37	19.65/19.46	32.33/30.63
T5S	30.14/29.44	12.27/11.56	24.62/23.88	39.63/39.53	20.42/20.27	33.48/33.20

Table 1: Summarization task results on the development / test sets for all models using ROUGE metric.

	XLSum			EUR-Lex-Sum		
	R1	R2	RL	R1	R2	RL
BETO2BETO	28.76/28.88	8.92/9.02	20.85/21.03	42.76/43.46	13.89/14.17	22.77/23.07
BETOShare	28.96/29.24	9.17/9.22	21.08/21.27	41.66/42.76	13.42/13.73	22.66/22.95
RoBERTa2RoBERTa	26.92/27.22	6.98/7.32	19.01/19.23	44.70/45.63	14.58/14.93	22.86/23.06
RoBERTaShare	26.89/27.08	6.99/7.15	18.94/19.11	44.24/44.22	13.84/13.90	22.55/22.65
mBART-large	31.56/31.76	10.94/10.89	22.26/22.27	68.20/67.37	52.89/51.45	58.05/56.35
mT5	28.54/28.59	10.18/10.28	21.42/21.49	65.66/64.20	50.32/48.88	56.07/53.95
BARTO	31.02/31.26	10.68/10.72	21.96/23.81	66.49/65.91	49.99/48.39	56.01/54.15
T5S	30.59/30.80	11.99/12.14	23.38/23.48	64.94/64.61	49.84/49.14	55.45/54.56

Table 2: Long-form summarization task results on the development / test sets using ROUGE metric.

extensive summaries. BART and T5S show better performance in this task, mainly because they leverage document formatting during pre-training. This strategy enhances their ability to handle long sequences, especially when dealing with extended documents. Finally, the multilingual models slightly outperform BARTO and T5S only in the case of EUR-Lex-Sum. This may be due to the diverse range of text types encountered by multilingual models during pre-training, including legal documents. [Appendix A](#) presents additional experiments exploring the potential of BARTO as a Longformer ([Beltagy et al., 2020](#)) for this task.

**Generative Question Answering** Table 3 presents the results of the generative question answering tasks in terms of ROUGE scores. Although SQAC and MLQA were originally designed as discriminative tasks, our results indicate that they serve as a suitable benchmark for generative question answering. BARTO and T5S show the best performance in all tasks, with T5S being the best, reaching 67.92 ROUGE on average. In this task, the performance difference between these two models is notably significant compared to BERT2BERT-style models, with BETOShare as the top performer of the group with an average of 27 ROUGE. This difference may stem from the self-supervised objective of BART and T5, enabling better transfer learning to this task compared to others ([Raffel et al., 2020](#)). Significantly, in these tasks, both BARTO and T5S outperform their multilingual counterparts, including the large-sized mBART. These results highlight the superior reading comprehension capabilities of our models.

**Split and Rephrase** Table 4 presents the comparison of models on the BiSECT dataset, with evaluations based on SARI ([Xu et al., 2016](#)) and BLEU ([Post, 2018](#)) scores. T5S achieves the highest scores for both metrics, averaging 56.37 SARI and 43.27 BLEU. While BARTO secures the second-highest BLEU score, its performance in SARI falls short. In particular, RoBERTaShare has the best performance among BERT2BERT-style models, averaging 51.29 SARI but surpassed by BARTO in BLEU score. T5S’s success could be attributed to its ability to generate sequences that closely resemble the input’s word order, a trait closely related to SARI’s evaluation criteria. Additionally, T5’s span-filling objective may facilitate sentence splitting, contributing to its overall performance boost. Lastly, T5S outshines its multimodal counterpart, while BARTO performs comparably to mBART-large.

**Dialogue** Table 5 compares the models for dialogue generation based on F1 and METEOR scores ([Banerjee and Lavie, 2005](#)). Among the Spanish models, BARTO leads with 33.27 F1 and 43.29 METEOR, closely followed by BETO2BETO with 32.98 F1 and 42.41 METEOR. Interestingly, T5S ranks fourth with averages of 29.88 F1 and 39.93 METEOR. T5S excels in tasks with high overlap between input and output, such as summarization or split-and-rephrase, which may explain its lower performance in dialogue generation where such overlap is reduced. Notably, multilingual models demonstrate comparable behavior to our models, with mBART outperforming BARTO and T5S due to its larger size. However, T5S outperforms mT5, suggesting our models’ superiority.

	SQAC			MLQA		
	R1	R2	RL	R1	R2	RL
BETO2BETO	23.84/24.19	10.89/11.10	22.87/23.14	35.01/33.10	23.91/22.29	34.36/32.37
BETOShare	28.61/29.22	15.08/15.52	27.69/28.17	33.77/33.37	23.61/23.24	33.08/32.63
RoBERTa2RoBERTa	25.42/24.86	12.43/12.16	24.53/23.78	32.21/31.01	20.50/18.90	31.58/30.12
RoBERTaShare	33.21/33.28	20.37/20.52	32.52/32.45	29.21/27.78	16.65/16.48	28.39/26.69
mBART-large	70.97/71.40	52.25/53.08	70.80/71.24	62.55/59.66	37.87/34.73	62.39/59.50
mT5	75.14/74.11	56.03/55.77	75.07/73.98	70.93/69.49	46.32/43.42	70.80/69.36
BARTO	77.92/77.00	58.88/58.45	77.78/76.87	68.69/66.45	44.08/40.87	68.57/66.34
T5S	80.68/78.80	60.39/59.33	80.64/78.64	72.02/70.45	47.53/44.37	71.84/70.32

Table 3: Generative question answering task results on the development / test sets using ROUGE.

	BISECT	
	SARI	BLEU
BETO2BETO	49.45/49.27	37.79/37.14
BETOShare	49.72/49.37	38.38/37.62
RoBERTa2RoBERTa	50.98/50.56	36.00/35.22
RoBERTaShare	51.49/51.09	37.19/36.16
mBART-large	50.20/50.13	39.78/39.24
mT5	55.91/55.74	43.39/42.62
BARTO	50.45/50.13	39.48/38.97
T5S	56.55/56.19	43.73/42.81

Table 4: Split-and-rephrase task results on the development / test sets for all models using SARI and BLEU metrics.

	MIAM	
	F1	METEOR
BETO2BETO	32.88/33.08	42.17/42.64
BETOShare	31.53/32.03	41.23/41.60
RoBERTa2RoBERTa	29.33/28.44	38.12/36.66
RoBERTaShare	30.31/30.83	38.68/39.18
MBART-large	37.11/36.65	43.96/43.26
mT5	29.41/30.90	38.98/39.51
BARTO	33.31/33.23	43.29/43.28
T5S	29.91/29.85	40.32/39.53

Table 5: Dialogue response generation task results on the development / test sets for all models using F1 and METEOR metrics.

**Machine Translation** BARTO and T5S are pre-trained with documents that comprise at least 98% accuracy for Spanish language prediction. We hypothesize that this threshold allows our models to acquire knowledge of other languages often found alongside Spanish, allowing them to perform the translation task. Therefore, we decide to fine-tune our models directly in their original form, according to the procedure of previous tasks. This contrasts with the approach suggested by Lewis et al. (2020a), which introduces new parameters to facilitate the adaptation of BART to a new language.

For better comparison, we employ the multilingual BERT checkpoint to initialize mBERT2mBERT. We evaluate all the modes using the BLEU score. In our experiments, we empirically find that maintaining a substantial batch size (i.e. 384) and limiting the number of fine-tuning steps positively impacts the

resulting performance of both BARTO and T5S, but not mBERT2mBERT.

Table 6 presents the experimental results. Both BARTO and T5S show similar performance in translating from Portuguese to Spanish (Fapesp-v2 PT→ES) and English to Spanish (WMT13 EN→ES). Additionally, the top-performing BERT2BERT-style model is mBERT2mBERT. Note that the BLEU score is higher for Fapesp-v2, and the performance gap of BARTO and T5S versus mBERT2mBERT is also more significant compared to WMT13. Our hypothesis is that this difference arises not only from the models exposed to the Portuguese text but also from the shared linguistic roots with Spanish. In fact, when analyzing the tokenizer, we observe Portuguese diacritics (e.g., ç, ã, ü).

Table 6 also shows the results of the experiments with Spanish as the source language for translation (Fapesp-v2 ES→PT and WMT13 ES→EN). Intuitively, one might expect our models to perform better when tasked with generating Spanish. However, they demonstrate competitive performance even when generating Portuguese and English. Surpassing mBERT2mBERT, the best BERT2BERT-style model, BARTO and T5S exhibit more competitive performance for Fapesp-v2 and WMT13. Interestingly, the BLEU score is slightly higher than when generating Spanish text. We hypothesize that this is due to the effectiveness of the encoding representation of the Spanish input for conditionally generating quality text.

Finally, we find that BARTO and T5S perform similarly to their multilingual versions. These findings underscore that a sequence-to-sequence model, pre-trained from scratch in Spanish with a minimal amount of input from another language, can effectively address translation tasks.

## 5.2. Discriminative Tasks

The results for the discriminative tasks are presented in Table 7. For all sequence classification tasks, we adopt the input formatting conventions established in previous work (Liu et al., 2019; Lewis et al., 2020a). Therefore, for classification, we utilize the [CLS] token representation from BETO

	Fapesp-v2		WMT13	
	PT → ES	ES → PT	EN → ES	ES → EN
mBERT2mBERT	55.14/61.21	54.90/60.01	25.58/22.29	22.30/19.62
mBERTShare	54.21/60.75	54.71/59.86	25.18/21.58	22.47/19.68
mBART-large	65.64/70.89	67.16/71.44	29.33/29.37	32.28/30.71
mT5	66.58/72.29	68.09/71.02	31.05/28.89	29.68/28.46
BARTO	65.84/71.06	66.98/70.25	30.22/27.41	30.22/28.64
T5S	66.37/71.85	68.65/72.40	31.99/29.37	31.10/29.64

Table 6: Machine translation task results on the development / test sets of Fapesp-v2 and newstest2012 / newstest2013 sets of WMT13 using BLEU metric.

	MLDoc Accuracy	PAWS-X F1	XNLI Accuracy	STS-es Combined	SemRel2024 Combined	SQAC F1
BETO	96.20/95.77	86.39/87.98	81.12/80.89	89.50/81.02	74.20/75.47	79.85/79.15
RoBERTa-BNE	95.80/95.92	88.34/89.76	80.68/80.63	89.37/80.46	71.39/75.09	80.16/80.39
BARTO	96.70/96.20	87.39/88.92	81.81/79.78	86.39/81.73	72.14/75.61	79.22/79.09
T5S <sup>†</sup>	96.70/96.63	88.61/88.58	80.92/80.56	86.64/75.65	63.84/66.60	58.46/62.29

Table 7: GLUES and SQAC tasks results on the development / test sets for all models. The combined metric represents the averaged Pearson and Spearman metrics. <sup>†</sup> indicates that T5S uses “text-to-text” format (Raffel et al., 2020) to solve the tasks.

or RoBERTa, and the  $\langle /s \rangle$  token representation from the decoder with BARTO. As T5 is not usually fine-tuned similarly to BART, we fine-tune T5S to generate the class label following the “text-to-text” format proposed by Raffel et al. (2020).

In the case of MLDoc, our findings indicate that BART and T5S exhibit superior performance, which can likely be attributed to their proficiency in handling lengthy sequences. Regarding PAWS-X and XNLI, our results show that BART and T5S, while not securing always the top positions, still exhibit competitive performance. Finally, in the case of STS-es and SemRel2024, we find that encoder-only models perform the best possible because of their specialization on sentence representation. Our results align with the behavior of English models, where BART and T5 tend to slightly lag behind the state-of-the-art in sentence-paired classification tasks (Lewis et al., 2020a; Raffel et al., 2020).

In the token classification task, we follow the standard span-based question answering approach, except for T5S, which remains operating as a generative model, as previously outlined. Consistent with our sequence classification findings, we observe a slight performance advantage for encoder-only models. Interestingly, T5S shows the lowest performance in this scenario. This discrepancy could be attributed to the F1 metric’s suitability for discriminative assessment, potentially overlooking T5S’s generative capabilities. However, it is worth noting that T5S has demonstrated its ability to generate accurate answers in Section 5.1.

## 6. Qualitative Analysis

Quantitative evaluation of BARTO and T5S (Section 5) demonstrates an improvement in performance compared to BERT2BERT-style models. To

gain deeper insights into the behavior of these two models, we conducted a qualitative analysis of their predictions focused on abstractive summarization, generative question answering, translation, and dialogue tasks. Table 8 shows the selected examples generated by BARTO and T5S, highlighting their respective strengths and weaknesses compared to a BERT2BERT-style baseline.

In the context of the summarization task, both BARTO and T5S exhibit the ability to generate coherent summaries that align with the source content. BARTO’s output tends to closely match the target text, maintaining a high textual similarity. Additionally, T5S works similarly, but in this specific example, we see a hallucination in its summary, “Escribe tu ensayo en voz alta”, which is not present in the source or target text. This could explain the superiority of BARTO in this task compared to T5S.

Shifting to the generative question answering task, it becomes evident that T5S and BARTO excel in generating more comprehensive answers. While RoBERTaShare initially generates similar responses, it occasionally introduces superfluous words that are unrelated to the original question.

Regarding the automatic translation task, both BARTO and T5S generate good translations that are not identical to the target but are completely coherent and valid. However, mBERT2mBERT hallucinates generating “mejorando la información”, something that is not in the source text.

Finally, in the case of dialogue generation, both BARTO and T5S create responses that are almost perfectly aligned with the target response. On the other hand, BETOShare hallucinates and mixes information specified in the source text or dialogue. These illustrative examples highlight the skills of BARTO and T5S, showcasing their strong grasp of



<b>Long-form Summarization: WikiLingua</b>	
Source	En caso de que hayas planificado con anticipación, esto te será fácil. Luego, regresa al ensayo cuando hayan pasado uno o dos días y revísalo [...]. Lee tu ensayo en voz alta, leyendo exactamente lo que haya en la hoja [...]. Podría serte de ayuda que imprimas el borrador [...]. Al trabajar con una copia física, te obligas a prestar atención de una forma nueva [...].
Target	Tómate uno o dos días sin mirar al ensayo. Lee el borrador. Revisa en donde sea necesario. Corrige con cuidado.
T5S	Regresa al ensayo al día siguiente. Revisa el ensayo. Edita el ensayo. Revisa el ensayo
BARTO	Revisa tu ensayo antes de enviarlo. Lee tu ensayo de principio a fin. Revisa tu ensayo en voz alta. Imprime y edita tu ensayo. Listo.
BETO2BETO	Regresa el ensayo a un momento conveniente. Revisa tu ensayo de principio a fin. Corrige tu ensayo en voz alta. Haz una copia física de tu ensayo.
<b>Generative Question Answer: SQAC</b>	
Source	Anillo halo El anillo halo es el más interno y grueso de los anillos de Júpiter. Su borde exterior coincide con el interior del anillo principal aproximadamente a un radio de 122 500 km [...].
Question	¿Cuál es el anillo más interno y grueso de los anillos de Júpiter?
Target	halo
T5S	El anillo halo
BARTO	anillo halo
RoBERTaShare	el anillo halo: Iris (Aright (Award (Aver
<b>Dialogue Generation: MIAM</b>	
Source	[...] . Hay varios trenes . El primero sale a la una y 50 de la madrugada y el último a las 7 de la tarde . Desea algo más ? . Sí , me gustaría obtener el tipo de tren para ir a Barcelona el próximo sábado día uno . Lo siento . No he entendido bien . puede repetir ? .
Target	Sí , me gustaría obtener el tipo de tren para ir a Barcelona el próximo sábado día uno .
T5S	Me gustaría obtener el tipo de tren para ir a Barcelona el próximo sábado día uno.
BARTO	Me gustaría obtener el tipo de tren para ir a Barcelona el próximo sábado día uno.
BETOShare	Sí, me gustaría saber el horario de trenes de Barcelona a Zaragoza el próximo sábado.
<b>Machine Translation: WMT13 EN→ES</b>	
Source	Information entry has improved (the system receives more data), but output still needs to get better.
Target	La introducción de información ha mejorado (el sistema recibe más datos), y la salida todavía no es muy buena.
T5S	La entrada de información ha mejorado (el sistema recibe más datos), pero la salida todavía necesita mejorarse.
BARTO	La entrada de información ha mejorado (el sistema recibe más datos), pero todavía es necesario mejorar la salida.
mBART2mBART	La entrada de información ha mejorado ( el sistema recibe más datos ), pero la salida todavía necesita mejorarse mejorando la información.

Table 8: Predictions generated by BARTO, T5S, and BERT2BERT-style models for qualitative comparison in abstractive summarization, generative question answering, dialogue, and machine translation tasks.

the Spanish language and their ability to generate natural and contextually relevant responses.

## 7. Conclusions

This work lays the foundation for future research in encoder-decoder architectures within the Spanish language domain. We present BART, T5, and BERT2BERT-style models, all exclusively pre-trained in Spanish. These models are accompanied by a diverse range of generative tasks intended to facilitate comprehensive evaluation. Our evaluation has demonstrated the effectiveness of

these models in addressing these challenges, with BARTO and T5S emerging as the top performers.

As we look forward, there is potential to fill the gaps in sequence-to-sequence tasks by building specialized datasets. Additionally, a thorough comparative analysis involving monolingual and multilingual models, akin to the study by [Agerri and Agirre \(2023\)](#), could offer valuable insights into their strengths and limitations. Lastly, we envision the pre-training of larger-scale language models, such as Llama ([Touvron et al., 2023](#)), to pave the way for advancements in emerging areas like chatbots.

## 8. Ethics Statement and Limitations

Our work presents new language models pre-trained exclusively in Spanish. The data used to train the models do not imply any violation of privacy. The potential negative social impacts from this work are similar to any other NLP models. Language models could potentially be used to create malicious and biased systems.

In this paper, we introduce base-sized language models, which may not be suitable for tasks requiring advanced capabilities. The current preference for larger models, which exhibit improved performance and emerging capabilities, directs our future efforts toward the release of larger architectures. For the generative question answering evaluation, we utilized span-based datasets. However, using these datasets generatively may lead the model to focus on reproducing exact information from the source text. Future efforts should prioritize creating and evaluating actual abstractive question answering datasets that present a more diverse range of answers beyond the input replication.

## 9. Acknowledgements

This work was funded by the European Research Council Advanced Grant 788506 and supported by the Google Cloud Research Credits program (GCP) and TPU Research Cloud program (TRC).

## 10. Bibliographical References

- Julien Abadji, Pedro Ortiz Suarez, Laurent Romary, and Benoît Sagot. 2022. [Towards a cleaner document-oriented multilingual crawled corpus](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4344–4355, Marseille, France. European Language Resources Association.
- Rodrigo Agerri and Eneko Agirre. 2023. [Lessons learned from the evaluation of spanish language models](#). *Procesamiento del Lenguaje Natural*, 70(0):157–170.
- Rahul Aralikkatte, Ziling Cheng, Sumanth Doddapaneni, and Jackie Chi Kit Cheung. 2023. [Varta: A large-scale headline-generation dataset for Indic languages](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 3468–3492, Toronto, Canada. Association for Computational Linguistics.
- Vladimir Araujo, Andrés Carvallo, Souvik Kundu, José Cañete, Marcelo Mendoza, Robert E. Mercer, Felipe Bravo-Marquez, Marie-Francine Moens, and Alvaro Soto. 2022. [Evaluation benchmarks for Spanish sentence representations](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6024–6034, Marseille, France. European Language Resources Association.
- Vladimir Araujo, Marie-Francine Moens, and Alvaro Soto. 2023. [Learning sentence-level representations with predictive coding](#). *Machine Learning and Knowledge Extraction*, 5(1):59–77.
- Dennis Aumiller, Ashish Chouhan, and Michael Gertz. 2022. [EUR-lex-sum: A multi- and cross-lingual dataset for long-form summarization in the legal domain](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7626–7639, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Wilker Aziz and Lucia Specia. 2011. [Fully automatic compilation of Portuguese-English and Portuguese-Spanish parallel corpora](#). In *Proceedings of the 8th Brazilian Symposium in Information and Human Language Technology*.
- Satanjeev Banerjee and Alon Lavie. 2005. [ME-TEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#).
- Ondřej Bojar, Christian Buck, Chris Callison-Burch, Christian Federmann, Barry Haddow, Philipp Koehn, Christof Monz, Matt Post, Radu Soricut, and Lucia Specia. 2013. [Findings of the 2013 Workshop on Statistical Machine Translation](#). In *Proceedings of the Eighth Workshop on Statistical Machine Translation*, pages 1–44, Sofia, Bulgaria. Association for Computational Linguistics.
- José Cañete, Gabriel Chaperon, Rodrigo Fuentes, Jou-Hui Ho, Hojin Kang, and Jorge Pérez. 2020. [Spanish pre-trained bert model and evaluation data](#). *PML4DC at ICLR 2020*.
- José Cañete, Sebastian Donoso, Felipe Bravo-Marquez, Andrés Carvallo, and Vladimir Araujo. 2022. [ALBETO and DistilBETO: Lightweight Spanish language models](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4291–4298, Marseille, France. European Language Resources Association.
- Diedre Carmo, Marcos Piau, Israel Campiotti, Rodrigo Nogueira, and Roberto Lotufo. 2020. [PtT5](#)

- Pretraining and validating the t5 model on brazilian portuguese data.
- Pierre Colombo, Emile Chapuis, Matthieu Labeau, and Chlo   Clavel. 2021. [Code-switched inspired losses for spoken dialog representations](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8320–8337, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Raj Dabre, Himani Shrotriya, Anoop Kunchukuttan, Ratish Puduppully, Mitesh Khapra, and Pratyush Kumar. 2022. [IndicBART: A pre-trained model for indic natural language generation](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1849–1863, Dublin, Ireland. Association for Computational Linguistics.
- Javier De la Rosa, Eduardo Ponferrada, Manu Romero, Paulo Villegas, Pablo Gonz  lez de Prado Salas, and Mar  a Grandury. 2022. [Bertin: Efficient pre-training of a spanish language model using perplexity sampling](#). *Procesamiento del Lenguaje Natural*, page 13–23.
- Pieter Delobelle, Thomas Winters, and Bettina Berendt. 2020. [RobBERT: a Dutch RoBERTa-based Language Model](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3255–3265, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Iakovos Evdaimon, Hadi Abdine, Christos Xypolopoulos, Stamatis Outsios, Michalis Vazirgiannis, and Giorgos Stamou. 2023. [Greekbart: The first pretrained greek sequence-to-sequence model](#).
- Asier Guti  rrez-Fandi  o, Jordi Armengol-Estap  , Marc P  mies, Joan Llop-Palao, Joaqu  n Silveira-Ocampo, Casimiro Pio Carrino, Aitor Gonzalez-Agirre, Carme Armentano-Oller, Carlos Rodriguez-Penagos, and Marta Villegas. 2022. [Maria: Spanish language models](#). *Procesamiento del Lenguaje Natural*, 68(0):39–60.
- Tahmid Hasan, Abhik Bhattacharjee, Md. Saiful Islam, Kazi Mubasshir, Yuan-Fang Li, Yong-Bin Kang, M. Sohel Rahman, and Rifat Shahriyar. 2021. [XL-sum: Large-scale multilingual abstractive summarization for 44 languages](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4693–4703, Online. Association for Computational Linguistics.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [DeBERTa: Decoding-enhanced bert with disentangled attention](#). In *International Conference on Learning Representations*.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. [Bag of tricks for efficient text classification](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. Association for Computational Linguistics.
- Moussa Kamal Eddine, Antoine Tixier, and Michalis Vazirgiannis. 2021. [BARThez: a skilled pre-trained French sequence-to-sequence model](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9369–9390, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Moussa Kamal Eddine, Nadi Tomeh, Nizar Habash, Joseph Le Roux, and Michalis Vazirgiannis. 2022. [AraBART: a pretrained Arabic sequence-to-sequence model for abstractive summarization](#). In *Proceedings of the The Seventh Arabic Natural Language Processing Workshop (WANLP)*, pages 31–42, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Joongwon Kim, Mounica Maddela, Reno Kriz, Wei Xu, and Chris Callison-Burch. 2021. [BiSECT: Learning to split and rephrase sentences with bitexts](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6193–6209, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Diederick P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.

- Moreno La Quatra and Luca Cagliero. 2023. [Bart-it: An efficient sequence-to-sequence model for Italian text summarization](#). *Future Internet*, 15(1).
- Faisal Ladhak, Esin Durmus, Claire Cardie, and Kathleen McKeown. 2020. [WikiLingua: A new benchmark dataset for cross-lingual abstractive summarization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4034–4048, Online. Association for Computational Linguistics.
- Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. 2022. [Deduplicating training data makes language models better](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8424–8445, Dublin, Ireland. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020a. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Patrick Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2020b. [MLQA: Evaluating cross-lingual extractive question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7315–7330, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pre-training approach](#).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie, Djamel Seddah, and Benoît Sagot. 2020. [Camembert: a tasty French language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online. Association for Computational Linguistics.
- Lorenzo De Mattei, Michele Cafagna, Felice Dell’Orletta, Malvina Nissim, and Marco Guerini. 2020. [Geppetto carves Italian into a language model](#). In *Proceedings of the Seventh Italian Conference on Computational Linguistics, CLiC-it 2020, Bologna, Italy, March 1-3, 2021*, volume 2769 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- El Moatez Billah Nagoudi, AbdelRahim Elmadany, and Muhammad Abdul-Mageed. 2022. [AraT5: Text-to-text transformers for Arabic language generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 628–647, Dublin, Ireland. Association for Computational Linguistics.
- Piotr Nawrot. 2023. [nanoT5: Fast & simple pre-training and fine-tuning of t5 models with limited resources](#). In *Proceedings of the 3rd Workshop for Natural Language Processing Open Source Software (NLP-OSS 2023)*, pages 95–101, Singapore. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nedjma Ousidhoum, Shamsuddeen Hassan Muhammad, Mohamed Abdalla, Idris Abdulmumin, Ibrahim Said Ahmad, Sanchit Ahuja, Alham Fikri Aji, Vladimir Araujo, Abinew Ali Ayele, Pavan Baswani, Meriem Beloucif, Chris Biemann, Sofia Bourhim, Christine De Kock, Genet Shanko Dekebo, Oumaima Hourrane, Gopichand Kanumolu, Lokesh Madasu, Samuel Rutunda, Manish Shrivastava, Tamar Solorio, Nirmal Surange, Hailegnaw Getaneh Tilaye, Krishnapriya Vishnubhotla, Genta Winata, Seid Muhie Yimam, and Saif M. Mohammad. 2024. [Semrel2024: A collection of semantic textual relatedness datasets for 14 languages](#).
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.

- Alec Radford and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67.
- Sascha Rothe, Shashi Narayan, and Aliaksei Severyn. 2020. Leveraging pre-trained checkpoints for sequence generation tasks. *Transactions of the Association for Computational Linguistics*, 8:264–280.
- Gabriele Sarti and Malvina Nissim. 2022. It5: Large-scale text-to-text pretraining for italian language understanding and generation.
- Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. 2020. MLSUM: The multilingual summarization corpus. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8051–8067, Online. Association for Computational Linguistics.
- Alejandro Vaca Serrano, Guillem Garcia Subies, Helena Montoro Zamorano, Nuria Aldama Garcia, Doaa Samy, David Betancur Sanchez, Antonio Moreno Sandoval, Marta Guerrero Nieto, and Alvaro Barbero Jimenez. 2022. Rigoberta: A state-of-the-art language model for spanish.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. MASS: Masked sequence to sequence pre-training for language generation. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5926–5936. PMLR.
- Pedro Javier Ortiz Suárez, Benoît Sagot, and Laurent Romary. 2019. Asynchronous pipelines for processing huge corpora on medium to low resource infrastructures. *Proceedings of the Workshop on Challenges in the Management of Large Corpora (CMLC-7) 2019*. Cardiff, 22nd July 2019, pages 9 – 16, Mannheim. Leibniz-Institut für Deutsche Sprache.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. Multilingual translation with extensible multilingual pretraining and fine-tuning.
- Eric S. Tellez, Daniela Moctezuma, Sabino Miranda, Mario Graff, and Guillermo Ruiz. 2023. Regionalized models for spanish language variations based on twitter. *Language Resources and Evaluation*, 57(4):1697–1727.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. Llama: Open and efficient foundation language models.
- Nguyen Luong Tran, Duong Le, and Dat Quoc Nguyen. 2022. BARTpho: Pre-trained sequence-to-sequence models for vietnamese. In *InterSpeech 2022*. ISCA.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Antti Virtanen, Jenna Kanerva, Rami Ilo, Jouni Luoma, Juhani Luotolahti, Tapio Salakoski, Filip Ginter, and Sampo Pyysalo. 2019. Multilingual is not enough: Bert for finnish.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. 2016. Optimizing statistical machine translation for text simplification. *Transactions of the Association for Computational Linguistics*, 4:401–415.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference*

of the North American Chapter of the Association for Computational Linguistics: *Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. [DIALOGPT : Large-scale generative pre-training for conversational response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.

Meng Zhou, Zechen Li, Bowen Tan, Guangtao Zeng, Wenmian Yang, Xuehai He, Zeqian Ju, Subrato Chakravorty, Shu Chen, Xingyi Yang, Yichen Zhang, Qingyang Wu, Zhou Yu, Kun Xu, Eric Xing, and Pengtao Xie. 2021. [On the generation of medical dialogs for COVID-19](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 886–896, Online. Association for Computational Linguistics.

## Appendix A: Longformer-Encoder-Decoder for Spanish (LEDO)

Transformer-based pre-trained language models face challenges in processing long sequences due to the quadratic scaling of their self-attention mechanism with sequence length. To address this issue, [Beltagy et al. \(2020\)](#) introduced Longformer, which combines windowed local-context self-attention with task-motivated global attention to alleviate the maximum input length limitation. An extension of Longformer, LED, adapts BART for this purpose. In our work, we explore the suitability of our BARTO model within this framework. We present LEDO, a model capable of processing sequences up to 16K tokens. Following the approach of [Beltagy et al. \(2020\)](#), we build LEDO by leveraging the weights of BARTO and initializing its new position embedding matrix by repeatedly copying BARTO’s 1K position embeddings 16 times.

	XLSum			EUR-Lex-Sum		
	R1	R2	RL	R1	R2	RL
BARTO	31.02/31.26	10.68/10.72	21.96/23.81	66.49/65.91	49.99/48.39	56.01/54.15
LEDO	32.23/32.19	12.82/12.72	24.10/24.13	62.12/61.11	40.98/39.48	44.95/43.51

Table 9: Long-form summarization task results on the development / test sets using ROUGE metric.

Table 9 presents the results for long-form summarization using BARTO and LEDO. We maintain consistent setup and hyperparameters for running the LEDO experiments. Interestingly, LEDO demonstrates superior performance over BARTO in XLSum, with an average of 23.03%, showcasing its capacity to handle lengthy inputs more effectively. However, in EUR-Lex-Sum, BARTO outperforms LEDO, which achieved an average performance of 48.69%. This difference may be attributed to the use of the same hyperparameters for LEDO as those for BARTO, which might not be optimal. Further investigation is required to fully explore the potential of LEDO, which will be pursued in future work.