

# Automatic Spanish Translation of the SQuAD Dataset for Multilingual Question Answering

Casimiro Pio Carrino, Marta R. Costa-jussà, José A. R. Fonollosa

TALP Research Center

Universitat Politècnica de Catalunya, Barcelona

{casimiro.pio.carrino, marta.ruiz, jose.fonollosa}@upc.edu

## Abstract

Recently, multilingual question answering became a crucial research topic, and it is receiving increased interest in the NLP community. However, the unavailability of large-scale datasets makes it challenging to train multilingual QA systems with performance comparable to the English ones. In this work, we develop the Translate Align Retrieve (TAR) method to automatically translate the Stanford Question Answering Dataset (SQuAD) v1.1 to Spanish. We then used this dataset to train Spanish QA systems by fine-tuning a Multilingual-BERT model. Finally, we evaluated our QA models with the recently proposed MLQA and XQuAD benchmarks for cross-lingual Extractive QA. Experimental results show that our models outperform the previous Multilingual-BERT baselines achieving the new state-of-the-art values of 68.1 F1 on the Spanish MLQA corpus and 77.6 F1 on the Spanish XQuAD corpus. The resulting, synthetically generated SQuAD-es v1.1 corpora, with almost 100% of data contained in the original English version, to the best of our knowledge, is the first large-scale QA training resource for Spanish.

**Keywords:** Question Answering, Multilinguality, Corpus Creation

## 1. Introduction

Question answering is a crucial and challenging task for machine-reading comprehension and represents a classical probe to assess the ability of a machine to understand natural language (Hermann et al., 2015). In the last years, the field of QA has made enormous progress, primarily by fine-tuning deep pre-trained architectures (Vaswani et al., 2017; Devlin et al., 2018; Liu et al., 2019b; Yang et al., 2019) on large-scale QA datasets. Unfortunately, large and high-quality annotated corpora are usually scarce for languages other than English, hindering advancement in Multilingual QA research.

Several approaches based on cross-lingual learning and synthetic corpora generation have been proposed. Cross-lingual learning refers to zero, and few-shot techniques applied to transfer the knowledge of a QA model trained on many source examples to a given target language with fewer training data. (Artetxe et al., 2019; Lee and Lee, 2019; Liu et al., 2019a) On the other hand, synthetic corpora generation methods are machine-translation (MT) based designed to automatically generate language-specific QA datasets as training resources (Alberti et al., 2019; Lee et al., 2018; Türe and Boschee, 2016). Additionally, a multilingual QA system based on MT at test time has also been explored (Asai et al., 2018)

In this paper, we follow the synthetic corpora generation approach. In particular, we developed the Translate-Align-Retrieve (TAR) method, based on MT and unsupervised alignment algorithm to translate an English QA dataset to Spanish automatically. Indeed, we applied our method to the popular SQuAD v1.1 generating its first Spanish version. We then trained two Spanish QA systems by fine-tuning the pre-trained Multilingual-BERT model. Finally, we evaluated our models on two Spanish QA evaluation set taken from the MLQA (Lewis et al., 2019) and XQuAD (Artetxe et al., 2019) benchmarks.

Our improvements over the current Spanish QA baselines demonstrated the capability of the TAR method, assessing the quality of our synthetically generated translated dataset.

In summary, the contributions we make are the followings: i) We define an automatic method to translate the SQuAD v1.1 training dataset to Spanish that can be generalized to multiple languages. ii) We created SQuAD-es v1.1, the first large-scale Spanish QA. iii) We establish the current state-of-the-art for Spanish QA systems validating our approach. We make both the code and the SQuAD-es v1.1 dataset freely available<sup>1</sup>.

## 2. The Translate-Align-Retrieve (TAR) Method on SQuAD

This section describes the TAR method and its application for the automatic translation of the Stanford Question Answering Dataset (SQuAD) v.1.1 (Rajpurkar et al., 2016) into Spanish. The SQuAD v1.1 is a large-scale machine reading comprehension dataset containing more than 100,000 questions crowd-sourced on Wikipedia articles. It represents a high-quality dataset for extractive question answering tasks where the answer to each question is a span annotated from the context paragraph. It is partitioned into a training and development set with 80% and 10% of the total examples, respectively. Unlike the training set, the development set contains at least two answers for each posed question intended for robust evaluation. Each example in the SQuAD v.1.1 is a  $(c, q, a)$  tuple made of a context paragraph, a question, and the related answers along with its start position in the context  $a_{start}$ .

Generally speaking, the TAR method is designed for the translation of the source  $(c_{src}, q_{src}, a_{src})$  examples into the corresponding target language  $(c_{tgt}, q_{tgt}, a_{tgt})$  examples. It consists of three main components:

- A trained neural machine translation (NMT) model from source to target language

<sup>1</sup><https://github.com/ccasimiro88/TranslateAlignRetrieve>

- A trained unsupervised word alignment model
- A procedure to translate the source  $(c_{src}, q_{src}, a_{src})$  examples and retrieve the corresponding  $(c_{tgt}, q_{tgt}, a_{tgt})$  translations using the previous components

The next sections explain in details how we built each component for the Spanish translation of the SQuAD v1.1 training set.

## 2.1. NMT Training

We created the first TAR component from scratch, by training an NMT model for English to Spanish direction. Our NMT parallel corpus is created by collecting the en-es parallel data from several resources. We first collected data from the WikiMatrix project (Schwenk et al., 2019) which uses state-of-the-art multilingual sentence embeddings techniques from the LASER toolkit<sup>2</sup> (Artetxe and Schwenk, 2018a; Artetxe and Schwenk, 2018b) to extract N-way parallel corpora from Wikipedia. Then, to further increase the size of the parallel data, we gathered additional resources from the open-source OPUS corpus (Tiedemann, 2012). Eventually, we selected data from 5 different resources, such as Wikipedia, TED-2013, News-Commentary, Tatoeba and OpenSubTitles (Lison and Tiedemann, 2016; Wolk and Marasek, 2015; Tiedemann, 2012).

The data pre-processing pipeline consisted of punctuation normalization, tokenisation, true-casing and eventually a joint source-target BPE segmentation (Sennrich et al., 2016) with a maximum of 50k BPE symbols. Then, we filtered out sentences longer than 80 tokens and removed all source-target duplicates. The final corpora consist of almost 6.5M parallel sentences for the training set, 5k sentence for the validation and 1k for the test set. The pre-processing pipeline is performed with the scripts in the Moses repository<sup>3</sup> and the Subword-nmt repository<sup>4</sup>.

We then trained the NMT system with the Transformer model (Vaswani et al., 2017). We used the implementation available in OpenNMT-py toolkit (Klein et al., 2017) in its default configuration for 200k steps with one GeForce GTX TITAN X device. Additionally, we shared the source and target vocabularies and consequently, we also share the corresponding source and target embeddings between the encoder and decoder. After the training, our best model is obtained by averaging across the final three consecutive checkpoints.

Finally, we evaluated the NMT system with the BLEU score (Papineni et al., 2002) on our test set. The model achieved a BLEU score of 45.60 point showing that it is good enough to be used as a pre-trained English-Spanish translator suitable for our purpose.

## 2.2. Source-Translation Context-Alignment

The role of the alignment component is to compute the alignment between the context sentences and their trans-

lations. We relied on an efficient and accurate unsupervised word alignment called *eflomal* (Östling and Tiedemann, 2016) based on a Bayesian model with Markov Chain Monte Carlo inference. We used a fast implementation named *efmaral*<sup>5</sup> released by the same author of *eflomal* model. The implementation also allows the generation of some priors that are used at inference time to align quickly. Therefore, we used tokenized sentences from the NMT training corpus to train a token-level alignment model and generate such priors.

## 2.3. Translate and Retrieve the Answers

The final component defines a strategy for translating and retrieving the answer translations to obtain the translated SQuAD examples. Giving the original SQuAD Dataset textual content, three steps are applied:

1. Translate all the  $(c_{src}, q_{src}, a_{src})$  instances
2. Compute the source-translation context alignment  $align_{src-tran}$
3. Retrieve the answer translations  $a_{tran}$  by using the source-translation context alignment  $align_{src-tran}$

The next sections describe in details the steps below.

**Content Translation and Context-Alignment** The first two steps are quite straightforward and easy to describe. First of all, all the  $(c_{src}, q_{src}, a_{src})$  examples are collected and translated with the trained NMT model. Second, Each source context is split into sentences, and then the alignments between the context sentences and their translations are computed. Before the translation, each context source  $c_{src}$  is split into sentences. Both the final context translation  $c_{tran}$  and context alignment  $align(c_{src}, c_{tran})$  are consequently obtained by merging the sentence translations and sentence alignments, respectively. Furthermore, it is important to mention that, since the context alignment is computed at a token level, we computed an additional map from the token positions to the word positions in the raw context. The resulting alignment maps a word position in the source context to the corresponding word position in the context translation. Eventually, each source context  $c_{src}$  and question  $q_{src}$  is replaced with its translation  $c_{tran}$  and  $q_{tran}$  while the source answer  $a_{src}$  is left unchanged. To obtain the correct answer translations, we designed a specific retrieval mechanism implemented in the last TAR component described next.

**Retrieve Answers with the Alignment** The SQuAD Dataset is designed for the extractive question answering task where, for each question on the paragraph, the corresponding answer is a span of the paragraph. However, when we analyzed the automatic translated  $(c_{tran}, q_{tran}, a_{tran})$  examples, we noticed many cases where the answer translated  $a_{tran}$  does not belong to the translation of the context  $c_{tran}$ . A possible explanation may be that the conditional probability of the target words

<sup>2</sup><https://github.com/facebookresearch/LASER>

<sup>3</sup><https://github.com/moses-smt/mosesdecoder>

<sup>4</sup><https://github.com/rsennrich/subword-nmt>

<sup>5</sup><https://github.com/robertostling/efmaral>

in the translation of the answer is conditioned on different sequences of previous words, depending on whether the source is a paragraph or the answer. To overcome this issue, we leverage on the previously computed context alignment  $align_{src-tran}$ . Therefore, we designed an answer extraction procedure that retrieves answers even when the answer translation is not contained in the paragraph translation. First, we use the  $align_{src-tran}$  to map the word positions of the answer source  $(a_{src}^{start}, \dots, a_{src}^{end})$  to the corresponding word positions of the answer translation  $(a_{tran}^{start}, \dots, a_{tran}^{end})$ . Also, a position reordering is applied to extract the  $a_{tran}^{start}$  and the  $a_{tran}^{end}$  as the minimum and maximum over  $(a_{tran}^{start}, \dots, a_{tran}^{end})$ , respectively. This position reordering accounts for the inversion during the translation. Then, for a given context, we look up the answer translation  $a_{tran}$  as a span of the context translation  $c_{tran}$ . The span it is searched from the corresponding start position  $a_{tran}^{start}$  in the context translation  $c_{tran}$ . It is necessary to detect in which part of the context translation  $c_{tran}$  the answer translation  $a_{tran}$  is mapped, to prevent the extraction of the incorrect answer span when it appears more than one sentence. Furthermore, the  $a_{tran}$  is lower-cased to improve the matching probability on the  $c_{tran}$ . If the answer translated is found in context translation, it is retrieved. In the opposite case, we retrieve the answer translation  $a_{tran}$  as the span between the  $a_{tran}^{start}$  and  $a_{tran}^{end}$ . The pseudo-code in figure 1 shows the algorithm implementation.

---

**Algorithm 1:** Implementation of the answer retrieval with alignment for each  $(c, q, a)$  example. The  $c_{src}$  and  $c_{tran}$  are the context source and translation,  $q_{src}$  is the question source,  $a_{src}^{start}$ ,  $a_{src}^{end}$  and  $a_{tran}^{start}$ ,  $a_{tran}^{end}$  are the start and end positions for the answer source and the answer translation,  $align_{src-tran}$  is the source-translation context alignment, and  $a_{tran}$  is the retrieved answer.

---

**Result:**  $a_{tran}$

```

for  $c_{src}$  in contexts paragraph do
   $c_{tran} \leftarrow$  get context translation ;
   $align(c_{src}, c_{tran}) \leftarrow$  get context alignment;
  for  $q$  in questions do
    for  $a_{src}$  in answers do
       $a_{tran} \leftarrow$  get answer translation ;
      if  $a_{tran}$  in  $c_{tran}$  then
        return  $a_{tran}$ ;
      else
         $(a_{src}^{start}, \dots, a_{src}^{end}) \leftarrow$  get src positions;
        compute tran positions with alignment
         $(a_{tran}^{start}, \dots, a_{tran}^{end}) \leftarrow align_{src-tran}$ ;
        compute start position as minimum
         $a_{tran}^{\prime start} \leftarrow \min(a_{tran}^{start}, a_{tran}^{end})$ ;
        compute end position as maximum
         $a_{tran}^{\prime end} \leftarrow \max(a_{tran}^{start}, a_{tran}^{end})$ ;
        return  $a_{tran} \leftarrow c_{tran}[a_{tran}^{\prime start} : a_{tran}^{\prime end}]$ ;
    end
  end
end

```

---

### 3. The SQuAD-es Dataset

We applied the TAR method to both the SQuAD v1.1 training datasets. The resulting Spanish version is referred as SQuAD-es v1.1. In table 4 we show some  $(c_{es}, q_{es}, a_{es})$  examples taken from the SQuAD-es v1. These examples show some good and bad  $(c, q, a)$  translated examples giving us a first glimpse of the quality of the TAR method for the SQuAD dataset.

#### 3.1. Error Analysis

As follows, we conduct a detailed error analysis in order to better understand the quality of the translated  $(c_{es}, q_{es}, a_{es})$  data.

The quality of the translated  $(c_{es}, q_{es}, a_{es})$  examples in the SQuAD-es v1.1 Dataset depends on both the NMT model and the unsupervised alignment model. The interplay among them in the TAR method that determines the final result. Indeed, while bad NMT translations irretrievably hurt the data quality also an erroneous source-translation alignment is responsible for the retrieval of wrong answer spans.

Therefore, we carried out a qualitative error analysis on SQuAD-es v1.1 in order to detect and characterize the produced errors and the factors behind them. We inferred the quality of the  $(c_{es}, q_{es}, a_{es})$  examples by looking at errors in the answer translations  $\{a_{es}\}$ . Indeed, based on the TAR method, a wrong answer translation provides an easy diagnostic clue for a potential error in both the context translation  $c_{es}$  or the source-translation context alignment  $align_{en-es}$ . We also make use of the question translations  $\{q_{es}\}$  to assess the level of correctness of the answer translations  $\{a_{es}\}$ . We pointed out systematic errors and classified them in two types:

**Type I: Misaligned Span** The answer translation is a span extracted from a wrong part of the context translation and indeed does not represent the correct candidate span. This error is caused by a misalignment in the source-translation alignment, or by a translation error when the NMT model generates a context translation shorter than the source context, thus retrieving to a wrong span.

**Type II: Overlapping Span** The answer translation is a span with a certain degree of overlap with the golden span on the context translation. Indeed, it may contain some missing/additional words or additional characters, such as commas, periods, parentheses or quotations. In particular, the additional characters in the retrieved answer become particularly evident compared to a fine-grained human annotation that excludes them (see Table 5). We also found cases when the answer translation span overlaps part of the next sentence. In general, this error type is generated by a slightly imprecise source-translation alignment or by an NMT translation error. However, quite often, the translation of the answer, with respect to its question, turns out to be acceptable as shown the examples shown in Table 5. Furthermore, based on their morphology, some of the errors of this type are easy to edit with a post-processing phase, unlike type I errors.

Overall, the two factors responsible for these error

types are both the NMT component and the alignment component in our TAR method. In particular, we noticed that the translation errors occur when the source context exceeds the average number of tokens of the source sentences in the training set. Furthermore, to have a rough estimation of the percentage of error types across the  $(c_{es}, q_{es}, a_{es})$  examples, we randomly selected three articles from SQuAD v1.1 and manually counted the error types. Besides, we divided the total examples into two sets, depending on how the answer translations are retrieved. The first set contains all the  $(c_{es}, q_{es}, a_{es})$  instances, while the second, smaller set, contains only the  $(c_{es}, q_{es}, a_{es})$  instances retrieved when the answer translation is matched as span in the context translation. In this way, we isolate the effect of the alignment component in the answer retrieval and evaluate its impact on the distribution of the error types. Table 1 shows the number of occurrence for each error type on the two sets of  $(c_{es}, q_{es}, a_{es})$  examples.

Error # (%)	$\{(c, q, a)\}_{all}$	$\{(c, q, a)\}_{no\_align}$
Type I	15 (7%)	0 (0%)
Type II	98 (43%)	7 (10%)
Total	113 (50%)	7 (10%)
Type II (acc.)	68	4
# of $(c, q, a)$ ex.	227	67

Table 1: The number of error types occurrences and its percentage for two sets of  $(c_{es}, q_{es}, a_{es})$  instances retrieved with and without alignment.

As a result, we found that the alignment is responsible for the introduction of a large number of error occurrences in the translated  $(c_{es}, q_{es}, a_{es})$  instances. As a consequence, when the answer translations are retrieved with the source-translation alignment, we found a significant increase of 40% of the total errors. On the other side, when the alignment is not involved in the retrieval process, the total number of translated examples is drastically reduced by 70% of the total number of examples in the other case. However, the number shows that the relative percentage of acceptable answers increased when the alignment is used. This analysis indicates that the TAR method can produce two types of synthetic data sets, a larger one with noisy examples and a smaller one with high quality. In the next section, we generate two Spanish translation of the SQuAD v1.1 training dataset, by considering or not  $(c_{es}, q_{es}, a_{es})$  retrieved with the alignment, to empirically evaluate their impact on the QA system performance.

### 3.2. Cleaning and Refinements

After the translation, we applied some heuristics to clean and refine the retrieved  $(c_{es}, q_{es}, a_{es})$  examples. Taking into account the error analysis of the previous section, we post-processed the type II errors. In particular, we first filter out words in the answers translations belonging to the next sentences. Then, we removed the extra leading and trailing punctuation. Eventually, we also removed empty answers translations that might be generated during the translation process. As shown in Table 2, almost half of the errors are corrected during post-processing, resulting in a cleaner

dataset.

Error #	$\{(c, q, a)\}_{all}$	$\{(c, q, a)\}_{no\_align}$
Type II	98	7
Type II (post-proc.)	51	4
$\Delta$	47	3

Table 2: The number of error types II for two sets of  $(c_{es}, q_{es}, a_{es})$  instances retrieved with and without alignment after the post-processing. The  $\Delta$  symbol indicates the number of errors that we have corrected during the post-processing, expressed as the difference between the number of errors before and after post-processing.

Moreover, in order to examine the impact on the QA system performance, we produced two versions of the SQuAD-es v1.1 training dataset. A standard one, containing all the translated  $(c_{es}, q_{es}, a_{es})$  examples and referred to as SQuAD-es v1.1 and another that keep only the  $(c_{es}, q_{es}, a_{es})$  examples retrieved without the use of the alignment, named SQuAD-es-small. Table 3 shows the statistics of the final SQuAD-es v1.1 datasets in terms of how many  $(c_{es}, q_{es}, a_{es})$  examples are translated over the total number of examples in its original English version. We also show the average context, question and answer length in terms of token. As a result, the SQuAD-es v1.1 training contains almost the 100% of the SQuAD v1.1 data while the SQuAD-es-small v1.1 is approximately half the size, with a about 53% of the data. In the next section, these two SQuAD-es v1.1 datasets will be employed to train question answering models for Spanish.

	SQuAD-es	SQuAD-es-small
# of ex.	87595/87599	46260/87599
Avg. $c$ len	140	138
Avg. $q$ len	13	12
Avg. $a$ len	4	3

Table 3: Number of  $(c_{es}, q_{es}, a_{es})$  examples in the final SQuAD-es v1.1 Datasets over the total number of the original SQuAD v1.1  $(c_{en}, q_{en}, a_{en})$ . We also computed the average length for the context, question and answers in terms of tokens.

## 4. QA Experiments

We trained two Spanish QA models by fine-tuning a pre-trained Multilingual-BERT (mBERT) model on our SQuAD-es v1.1 datasets following the method used in (Devlin et al., 2018). We employed the implementation in the open-source HuggingFace’s Transformers library (Wolf et al., 2019). Our models have been trained for three epochs a GTX TITAN X GPU device using the default parameter values set in the library scripts.

The goal is to assess the quality of our synthetically generated datasets used as a training resource for Spanish QA models. We performed the Spanish QA evaluation on two recently proposed, freely available, corpus for cross-lingual

	en	es
Context	Fryderyk Chopin was born in <b>Żelazowa Wola</b> , 46 kilometres (29 miles) west of Warsaw, in what was then the Duchy of Warsaw, a Polish state established by Napoleon. The parish baptismal record gives his birthday as <b>22 February 1810</b> , and cites his given names in the Latin form <b>Fridericus Franciscus</b> (in Polish, he was Fryderyk Franciszek). However, the composer and his family used the birthdate 1 March,[n 2] which is now generally accepted as the correct date.	Fryderyk Chopin nació en <b>Żelazowa Wola</b> , 46 kilómetros al oeste de Varsovia, en lo que entonces era el Ducado de Varsovia, un estado polaco establecido por Napoleón. El registro de bautismo de la parroquia da su cumpleaños el <b>22 de febrero de 1810</b> , y cita sus nombres en latín <b>Fridericus Franciscus</b> (en polaco, Fryderyk Franciszek). Sin embargo, el compositor y su familia utilizaron la fecha de nacimiento 1 de marzo, [n 2] que ahora se acepta generalmente como la fecha correcta.
Question	1) In what village was Frédéric born in?	1) ¿En qué pueblo nació Frédéric?
Answer	1) <b>Żelazowa Wola</b>	1) <b>Żelazowa Wola</b> ,
Question	2) When was his birthday recorded as being?	2) ¿Cuándo se registró su cumpleaños?
Answers	2) <b>22 February 1810</b>	2) <b>22 de febrero de 1810</b> ,
Question	3) What is the Latin form of Chopin’s name?	3) ¿Cuál es la forma latina del nombre de Chopin?
Answer	3) <b>Fridericus Franciscus</b>	3) <b>Fridericus Franciscus</b>
Context	During the Five Dynasties and Ten Kingdoms period of China ( <b>907–960</b> ), while the fractured political realm of China saw no threat in a Tibet which was in just as much political disarray, there was little in the way of Sino-Tibetan relations. Few documents involving Sino-Tibetan contacts survive from the Song dynasty (960–1279). The Song were far more concerned with countering northern enemy states of <b>the Khitan</b> -ruled Liao dynasty (907–1125) and <b>Jurchen</b> -ruled Jin dynasty (1115–1234).	Durante el período de las Cinco Dinastías y los Diez Reinos de China ( <b>907-960</b> ), mientras que el reino político fracturado de China no vio ninguna amenaza en un Tíbet que estaba en el mismo desorden político. Pocos documentos que involucran contactos sino-tibetanos sobreviven de la dinastía Song (960-1279). Los Song estaban mucho más preocupados por contrarrestar los estados enemigos del norte de <b>la dinastía Liao gobernada por Kitán</b> (907-1125) y la dinastía Jin <b>gobernada por Jurchen</b> (1115-1234).
Question	1) When did the Five Dynasties and Ten Kingdoms period of China take place?	1) ¿Cuándo tuvo lugar el período de las Cinco Dinastías y los Diez Reinos de China?
Answer	1) <b>907–960</b>	1) ( <b>907-960</b> ),
Question	2) Who ruled the Liao dynasty?	2) ¿Quién gobernó la dinastía Liao?
Answer	2) <b>the Khitan</b>	2) <b>la dinastía Liao gobernada por Kitán</b>
Question	3) Who ruled the Jin dynasty?	3) ¿Quién gobernó la dinastía Jin?
Answer	3) <b>Jurchen</b>	3) <b>gobernada por Jurchen</b>

Table 4: Examples of  $(c_{es}, q_{es}, a_{es})$  examples selected in two of the first articles in the SQuAD v1.1 training dataset, *Frédéric Chopin* and *Sino-Tibetan relations during the Ming dynasty*. The corresponding answer spans in the contexts are highlighted in bold.

QA evaluation, the MLQA and XQuAD corpora (Lewis et al., 2019; Artetxe et al., 2019). The MLQA corpus is an N-way parallel dataset consisting of thousands of QA examples crowd-sourced from Wikipedia in 7 languages, among which Spanish. It is split into development and test sets, and it represents an evaluation benchmark for cross-lingual QA systems.

Similarly, the XQuAD corpus is another cross-lingual QA

benchmark consisting of question-answers pairs from the development set of SQuAD v1.1 translated by professional translators in 10 languages, including Spanish. Therefore, we used both the MLQA and XQuAD benchmark to evaluate the performance of our trained models. We measured the QA performance with the F1, and Exact Match (EM) score computed using the official evaluation script available in the MLQA repository, that represents a multilingual gen-

	en	es	Error
Context	The Medill School of Journalism has produced notable journalists and political activists including <b>38</b> Pulitzer Prize laureates. National correspondents, reporters and columnists such as The New York Times’s Elisabeth Bumiller, David Barstow, Dean Murphy, and Vincent Laforet, USA Today’s Gary Levin, Susan Page and Christine Brennan, NBC correspondent Kelly O’Donnell, CBS correspondent <b>Richard Threlkeld</b> , CNN correspondent Nicole Lapin and former CNN and current Al Jazeera America anchor Joie Chen, and ESPN personalities Rachel Nichols, Michael Wilbon, Mike Greenberg, Steve Weissman, J. A. Adande, and Kevin Blackstone. The bestselling author of the A Song of Ice and Fire series, George R. R. Martin, earned a B.S. and M.S. from Medill. Elisabeth Leamy is the recipient of 13 Emmy awards and 4 Edward R. Murrow Awards.	La Escuela Medill de Periodismo ha producido notables periodistas y activistas políticos incluyendo <b>38 premios Pulitzer</b> . Los corresponsales nacionales, reporteros y columnistas como Elisabeth Bumiller de The New York Times, David Barstow, Dean Murphy, y Vincent Laforet, Gary Levin de USA Today, Susan Page y Christine Brennan de <b>NBC</b> . A. Adande, y Kevin Blackstone. El autor más vendido de la serie Canción de Hielo y Fuego, George R. R. Martin, obtuvo un B.S. Y M.S. De Medill. Elisabeth Leamy ha recibido 13 premios Emmy y 4 premios Edward R. Murrow.	-
Question	1) Which CBS correspondant graduated from The Medill School of Journalism?	1) ¿Qué corresponsal de CBS se graduó de la Escuela Medill de Periodismo?	-
Answer	1) <b>Richard Threlkeld</b>	1) <b>NBC</b> .	Type I
Question	2) How many Pulitzer Prize laureates attended the Medill School of Journalism?	1) ¿Cuántos premios Pulitzer asistieron a la Escuela Medill de Periodismo?	-
Answers	2) <b>38</b>	2) <b>38 premios Pulitzer. Los</b>	Type II
Context	Admissions are characterized as <b>”most selective”</b> by U.S. News & World Report. There were 35,099 applications for the undergraduate class of 2020 (entering 2016), and 3,751 ( <b>10.7%</b> ) were admitted, making Northwestern one of the most selective schools in the United States. For freshmen enrolling in the class of 2019, the interquartile range (middle 50%) on the SAT was 690–760 for critical reading and 710-800 for math, ACT composite scores for the middle 50% ranged from 31–34, and <b>91%</b> ranked in the top ten percent of their high school class.	Las admisiones se caracterizan como <b>”más selectivas”</b> por U.S. News & World Report. Hubo 35.099 solicitudes para la clase de pregrado de 2020 (ingresando en 2016), y 3.751 ( <b>10.7%</b> ) fueron admitidos, haciendo de Northwestern una de las escuelas más selectivas en los Estados Unidos. Para los estudiantes de primer año de matriculación en la clase de 2019, el rango intermedio (50% medio) en el SAT fue de 690-760 para lectura crítica y 710-800 para matemáticas, compuesto ACT para el <b>31%</b> .	-
Question	1) What percentage of freshman students enrolling in the class of 2019 ranked in the top 10% of their high school class?	1) ¿Qué porcentaje de estudiantes de primer año matriculados en la clase de 2019 se ubicó en el 10% superior de su clase de secundaria?	-
Answer	1) <b>91%</b>	1) <b>31%</b>	Type I
Question	2) What percentage of applications were admitted for the undergraduate class entering in 2016?	2) ¿Qué porcentaje de solicitudes fueron admitidas para la clase de pregrado en 2016?	-
Answer	2) <b>10.7%</b>	2) ( <b>10.7%</b> )	Type II (acc.)
Question	3) How selective are admissions at Northwestern characterized by U.S. News	3) ¿Cuán selectivas son las admisiones en Northwestern caracterizadas por U.S. News	-
Answer	3) <b>most selective</b>	3) <b>”más selectivas”</b>	Type II (acc.)

Table 5: Examples of error types in the  $(c_{es}, q_{es}, a_{es})$  examples in the randomly selected article *Northwestern University* from SQuAD v1.1. Dataset. Wrong answers are highlighted in red while acceptable answers in green.

eralization of the commonly-used SQuAD evaluation script<sup>6</sup>.

Results of the evaluation are shown in Table 6 and Table 7. On the MLQA corpus, our best model beat the state-of-the-art F1 score of the XLM (MLM + TLM, 15 languages)(Lample and Conneau, 2019) baseline for Spanish QA. Equally, on the XQuAD corpus, we set the state-of-the-art for F1 and EM score. Overall, the TAR-train+mBERT models perform better than the current mBERT-based baselines, showing significant improvements of 1.2%–3.3% increase in F1 score and 0.6%–6.5% raise in EM points across the MLQA and XQuAD benchmarks. Interestingly enough, when compared to a model trained on synthetic data, such as the translate-train-mBERT, our TAR-train-mBERT models show even more substantial improvements, setting out a notable increase of 11.6–14.2% F1 score and 9.8–10.9% EM score. These results indicate that the quality of the SQuAD-es v1.1 dataset is good enough to train a Spanish QA model able to reach state-of-the-art accuracy, therefore proving the efficacy of the TAR method for synthetic corpora generation.

The QA evaluation demonstrates that the performance on the Spanish MLQA and XQuAD benchmarks of the mBERT increased by 2.6 – 4.2% F1 score and 1.1 – 2.3% EM score when the SQuAD-es v1.1 dataset is used compared the SQuAD-es-small v1.1 dataset. Based on the error analysis in section 3, we can assume that the SQuAD-es v1.1 is a bigger but noisy dataset, compared to the SQuAD-es-small that is the smaller but more accurate. Therefore, we observe that the mBERT model may be robust enough to tolerate noisy data giving more importance to the quantity. This observation connects to the problem of quality versus quantity in synthetic corpora generation and its application to multilingual reading comprehension (Lee et al., 2019)

## 5. Conclusions

In this paper, we designed the TAR method to automatically translate the SQuAD-es v1.1 training dataset into Spanish. Hence, we applied the TAR method to generated the SQuAD-es v1.1 dataset, the first large-scale training resources for Spanish QA. Finally, we employed the SQuAD-es v1.1 dataset to train QA systems that achieved state-of-the-art performance on the Spanish QA task, demonstrating the efficacy of the TAR approach for synthetic corpora generation. Therefore, we make the SQuAD-es dataset freely available and encourage its usage for multilingual QA.

The results achieved so far encourage us to look forward and extend our approach in future works. First of all, we will apply the TAR method to translated the SQuAD v2.0 dataset (Rajpurkar et al., 2018) and other large-scale extractive QA such as Natural Questions(Kwiatkowski et al., 2019). Moreover, we will also exploit the modularity of the TAR method to support languages other than Spanish to prove the validity of our approach for synthetic corpora generation.

## Acknowledgements

This work is supported in part by the Spanish Ministerio de Economía y Competitividad, the European Regional Development Fund and the Agencia Estatal de Investigación, through the postdoctoral senior grant Ramón y Cajal (FEDER/MINECO) and the project PCIN-2017-079 (AEI/MINECO).

## 6. Bibliographical References

- Alberti, C., Andor, D., Pitler, E., Devlin, J., and Collins, M. (2019). Synthetic QA corpora generation with roundtrip consistency. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6168–6173, Florence, Italy, July. Association for Computational Linguistics.
- Artetxe, M. and Schwenk, H. (2018a). Margin-based parallel corpus mining with multilingual sentence embeddings. *CoRR*, abs/1811.01136.
- Artetxe, M. and Schwenk, H. (2018b). Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *CoRR*, abs/1812.10464.
- Artetxe, M., Ruder, S., and Yogatama, D. (2019). On the cross-lingual transferability of monolingual representations. *CoRR*, abs/1910.11856.
- Asai, A., Eriguchi, A., Hashimoto, K., and Tsuruoka, Y. (2018). Multilingual extractive reading comprehension by runtime machine translation. *CoRR*, abs/1809.03275.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. (2018). BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.
- Hermann, K. M., Kociský, T., Grefenstette, E., Espeholt, L., Kay, W., Suleyman, M., and Blunsom, P. (2015). Teaching machines to read and comprehend. *CoRR*, abs/1506.03340.
- Klein, G., Kim, Y., Deng, Y., Crego, J., Senellart, J., and Rush, A. (2017). Opennmt: Open-source toolkit for neural machine translation. 09.
- Kwiatkowski, T., Palomaki, J., Redfield, O., Collins, M., Parikh, A., Alberti, C., Epstein, D., Polosukhin, I., Kelleey, M., Devlin, J., Lee, K., Toutanova, K. N., Jones, L., Chang, M.-W., Dai, A., Uszkoreit, J., Le, Q., and Petrov, S. (2019). Natural questions: a benchmark for question answering research. *Transactions of the Association of Computational Linguistics*.
- Lample, G. and Conneau, A. (2019). Cross-lingual language model pretraining. *CoRR*, abs/1901.07291.
- Lee, C. and Lee, H. (2019). Cross-lingual transfer learning for question answering. *CoRR*, abs/1907.06042.
- Lee, K., Yoon, K., Park, S., and Hwang, S.-w. (2018). Semi-supervised training data generation for multilingual question answering. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May. European Language Resources Association (ELRA).
- Lee, K., Park, S., Han, H., Yeo, J., Hwang, S.-w., and Lee, J. (2019). Learning with limited data for multilingual reading comprehension. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on*

<sup>6</sup><https://rajpurkar.github.io/SQuAD-explorer/>

	<b>F1 / EM</b>	<b>es</b>
Our models	TAR-train + mBERT (SQuAD-es)	<b>68.1</b> / 48.3
	TAR-train + mBERT (SQuAD-es-small)	65.5 / 47.2
MLQA	mBERT	64.3 / 46.6
mBERT baselines	Translate-train + mBERT	53.9 / 37.4
state-of-the-art	XLM (MLM + TLM, 15 languages)	68.0 / <b>49.8</b>

Table 6: Evaluation results in terms of F1 and EM scores on the MLQA corpus for our Multilingual-BERT models trained with two versions of SQuAD-es v1.1 and the current Spanish Multilingual-BERT baselines

	<b>F1 / EM</b>	<b>es</b>
Our models	TAR-train + mBERT (SQuAD-es)	<b>77.6 / 61.8</b>
	TAR-train + mBERT (SQuAD-es-small)	73.8 / 59.5
XQuAD mBERT baselines	JointMulti 32k voc	59.5 / 41.3
	JointPair with Joint voc	68.3 / 47.8
	JointPair with Disjoint voc	72.5 / 52.5
	JointMulti 200k voc (state-of-the-art)	74.3 / 55.3

Table 7: Evaluation results in terms of F1 and EM scores on the XQuAD corpus for our Multilingual-BERT models trained with two versions of SQuAD-es v1.1 and the current Spanish Multilingual-BERT baselines

- Natural Language Processing (EMNLP-IJCNLP)*, pages 2840–2850, Hong Kong, China, November. Association for Computational Linguistics.
- Lewis, P., Oguz, B., Rinott, R., Riedel, S., and Schwenk, H. (2019). Mlqa: Evaluating cross-lingual extractive question answering. *arXiv preprint arXiv:1910.07475*.
- Liu, J., Lin, Y., Liu, Z., and Sun, M. (2019a). XQA: A cross-lingual open-domain question answering dataset. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2358–2368, Florence, Italy, July. Association for Computational Linguistics.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019b). Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.
- Östling, R. and Tiedemann, J. (2016). Efficient word alignment with Markov Chain Monte Carlo. *Prague Bulletin of Mathematical Linguistics*, 106:125–146, October.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (2002). Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA, July. Association for Computational Linguistics.
- Rajpurkar, P., Zhang, J., Lopyrev, K., and Liang, P. (2016). Squad: 100, 000+ questions for machine comprehension of text. *CoRR*, abs/1606.05250.
- Rajpurkar, P., Jia, R., and Liang, P. (2018). Know what you don’t know: Unanswerable questions for squad. *CoRR*, abs/1806.03822.
- Sennrich, R., Haddow, B., and Birch, A. (2016). Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany, August. Association for Computational Linguistics.
- Türe, F. and Boschee, E. (2016). Learning to translate for multilingual question answering. *CoRR*, abs/1609.08210.
- 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 30*, pages 5998–6008. Curran Associates, Inc.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., and Brew, J. (2019). Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Yang, Z., Dai, Z., Yang, Y., Carbonell, J. G., Salakhutdinov, R., and Le, Q. V. (2019). Xlnet: Generalized autoregressive pretraining for language understanding. *CoRR*, abs/1906.08237.

## 7. Language Resource References

- 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).
- Schwenk, H., Chaudhary, V., Sun, S., Gong, H., and Guzmán, F. (2019). Wikimatrix: Mining 135m parallel sentences in 1620 language pairs from wikipedia. *CoRR*, abs/1907.05791.
- Tiedemann, J. (2012). Parallel data, tools and interfaces in opus. In Nicoletta Calzolari (Conference Chair), et al., editors, *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC’12)*, Istanbul, Turkey, may. European Language Resources Association (ELRA).



Wolk, K. and Marasek, K. (2015). Building subject-aligned comparable corpora and mining it for truly parallel sentence pairs. *CoRR*, abs/1509.08881.