

INFERES : A Natural Language Inference Corpus for Spanish Featuring Negation-Based Contrastive and Adversarial Examples

Venelin Kovatchev

School of Information
The University of Texas at Austin
venelin@utexas.edu

Mariona Taulé

Centre de Llenguatge i Computació
Institut de Recerca en Sistemes Complexos
Universitat de Barcelona
mtaule@ub.edu

Abstract

In this paper, we present INFERES - an original corpus for Natural Language Inference (NLI) in European Spanish. We propose, implement, and analyze a variety of corpus-creating strategies utilizing expert linguists and crowd workers. The objectives behind INFERES are to provide high-quality data, and, at the same time to facilitate the systematic evaluation of automated systems. Specifically, we focus on measuring and improving the performance of machine learning systems on negation-based adversarial examples and their ability to generalize across out-of-distribution topics.

We train two transformer models on INFERES (8,055 gold examples) in a variety of scenarios. Our best model obtains 72.8% accuracy, leaving a lot of room for improvement. The “hypothesis-only” baseline performs only 2%-5% higher than majority, indicating much fewer annotation artifacts than prior work. We find that models trained on INFERES generalize very well across topics (both in- and out-of-distribution) and perform moderately well on negation-based adversarial examples.

1 Introduction

In the task of Natural Language Inference (NLI), an automated system has to determine the meaning relation that holds between two texts. The model has to make a three-way choice between *entailment*: a hypothesis (h) is true given a premise (p) (e.g. **1.**); *contradiction*: a hypothesis (h) is false given a premise (p) (e.g. **2.**); or *neutral*: the truth value of the hypothesis (h) cannot be determined solely based on the premise (p) (e.g.: **3.**).

- 1.** *p*) John goes to work every day with a car.
h) John has a job.
- 2.** *p*) John goes to work every day with a car.
h) John takes the bus to go to work.
- 3.** *p*) John goes to work every day with a car.
h) John has a Porsche.

NLI (formerly known as Recognizing Textual Entailment (RTE)) is one of the core tasks in the popular benchmarks for Natural Language Understanding GLUE (Wang et al., 2018) and Super GLUE (Wang et al., 2019). Hundreds of machine learning systems compete on these benchmarks, improving the state of NLU.

One key limitation of NLI research is that most of the existing corpora are only for English. Limited research has been done on multilingual and non-English corpora (Peñas et al., 2006; Conneau et al., 2018; Amirkhani et al., 2020; Ham et al., 2020; Hu et al., 2020; Mahendra et al., 2021).

Another well-known issue with NLI is the quality of the existing datasets and the limitations of the models trained on them. On most NLI corpora, state-of-the-art transformer based models can obtain quantitative results (Accuracy and F1) that equal or exceed human performance. Despite this high performance, researchers have identified numerous limitations and potential problems. Poliak et al. (2018) found that annotation artifacts in the datasets enable the models to predict the label by only looking at the hypothesis. NLI models are often prone to adversarial attacks (Williams et al., 2018) and may fail on instances that require specific linguistic capabilities (Hossain et al., 2020; Saha et al., 2020).

In this paper we address both of these shortcomings in NLI research. We present INFERES - to the best of our knowledge, the first original NLI corpus for Spanish, not adapted from another language or task. We study prior work for strategies that can reduce annotation artifacts and increase the linguistic variety of the corpus, resulting in a dataset that is more challenging for automated systems to solve. We also design the corpus in a way that facilitates systematic evaluation of automated systems on: 1) negation-based adversarial examples; 2) out-of-distribution examples.

We propose, implement, and analyze three dif-

ferent strategies for the generation and annotation of text pairs. In the *generation* strategy, expert linguists write original hypotheses given a premise. In the *rewrite* strategy, expert linguists create contrastive and adversarial examples by rewriting and re-annotating “generated” pairs. In the *annotation* strategy, we first generate text pairs in a semi-automated manner and then use crowd annotators to determine the meaning relation. The final INFERES corpus contains 8,055 gold standard premise-hypothesis pairs. The core part of the corpus is expert-generated and we make an additional effort to ensure the quality of the data and the linguistic diversity of the examples.

We provide two baselines for INFERES by fine-tuning multilingual BERT and BETO (Spanish BERT) transformer models. On the full dataset, BETO obtains 72.8% accuracy, indicating that the classification task is non-trivial. Both mBERT and BETO perform poorly in the “hypothesis-only” condition, indicating fewer annotation artifacts in the corpus compared to prior work. Both systems generalize well across the different topics in INFERES both “in-distribution” and “out-of-distribution”. We notice a substantial drop in performance when evaluating negation-based adversarial examples, however the systems still outperform majority and “hypothesis-only”.

INFERES expands the scope of the NLI research in Spanish, provides new set of naturally occurring contrastive and adversarial examples, and facilitates the study of negation and coreference in the context of NLI. As part of the corpus creation, we also present and analyze three unique strategies for creating examples. All our data and baseline models are being released to the community¹.

The rest of this article is organized as follows. Section 2 discusses the related work. Section 3 formulates our objectives and introduces the different corpus-creation strategies. Section 4 describes the final corpus and basic statistical data regarding it. Section 5 presents the machine learning experimental setup and results. Section 6 is devoted to a discussion of the results and their implications. Finally, Section 7 concludes the article.

2 Related Work

The task of Recognizing Textual Entailment (RTE) was proposed in Dagan et al. (2006) as a binary clas-

sification (“entailment” / “non-entailment”). The RTE competition ran for seven editions (Bar Haim et al., 2006; Giampiccolo et al., 2007, 2008; Ben-tivogli et al., 2009, 2010, 2011). RTE was later reformulated as a three-way decision and ultimately renamed Natural Language Inference in the SNLI (Bowman et al., 2015) and the MNLI (Williams et al., 2018) corpora. Both the RTE and the NLI tasks form part of the Natural Language Understanding benchmarks GLUE (Wang et al., 2018) and Super-GLUE (Wang et al., 2019). The NLU benchmarks attracted a lot of attention from the community and by 2020 the state-of-the-art systems reported human level performance. Parrish et al. (2021) proposed a “linguist-in-the-loop” corpus creation to improve the quality of the data.

The “super-human” performance of NLI systems has been questioned by a number of researchers. Poliak et al. (2018) found that annotation artifacts in the datasets enable the models to predict the label by only looking at the hypothesis. McCoy et al. (2019) and Gururangan et al. (2018) demonstrate that state-of-the-art NLI systems often rely on heuristics and annotation artifacts.

Systematic approaches to evaluation propose different sets of stress-tests for NLI and NLU systems (Kovatchev et al., 2018a; Naik et al., 2018; Wallace et al., 2019; Kovatchev et al., 2019; Ribeiro et al., 2020; Kovatchev et al., 2020). The attacks can be inspired by linguistic phenomena or empirical use cases. Systematic evaluations show that NLI and other NLU systems often underperform on complex linguistic phenomena such as conjunction (Saha et al., 2020), negation (Hossain et al., 2020), and coreference (Kovatchev et al., 2022). Researchers also experimented with creating contrastive examples that differ only slightly from training examples, but have a different label (Glockner et al., 2018; Kaushik et al., 2020; Gardner et al., 2020). Adversarially created datasets such as Adversarial NLI (Nie et al., 2020) and Dynabench NLI (Kiela et al., 2021) demonstrate that there is a lot of room for improvement regarding NLI datasets and models.

Most of the available resources for NLI research are in English. Conneau et al. (2018) present XNLI, a multilingual dataset created by translating English NLI examples into other languages. The interest in multilingual NLI has resulted in the creation of some novel non-English resources such as the Korean NLI corpus (Ham et al., 2020), Chinese NLI corpus (Hu et al., 2020), Persian NLI corpus

¹At <https://github.com/venelink/inferes> InferES is also added as a HuggingFace dataset

(Amirkhani et al., 2020), Indonesian NLI corpus (Mahendra et al., 2021), and indigenous languages of the Americas NLI corpus (Ebrahimi et al., 2022). For Spanish, the only available resources are the Spanish portion of XNLI and the SPARTE corpus for RTE (Peñas et al., 2006) which was adapted from Question Answering data.

3 Objectives and Corpus Creation

When creating INFERES we experimented with different strategies for obtaining gold examples. To the best of our knowledge, this is the first time various annotation strategies are combined and compared in a single NLI corpus. We adopt three different approaches, used in prior work: our **generation** strategy is similar to the original RTE and NLI corpus creation; our **rewrite** strategy is inspired from work in generating adversarial and contrastive examples; our **annotation** strategy scales well with data and allows us to compare expert- and crowd-rated datasets. Our aim was to provide interesting and diverse examples that cover a large range of use cases and linguistic phenomena. We hope that INFERES can be used not only to train automated systems, but also to better understand the nature of inference. We formulated three main **objectives**:

- O1 To create a native NLI dataset for the Spanish language. The existing resources are either an adaptation from a different task or a translation from English.
- O2 To promote better data quality and corpus creation practices. We aim to create a more challenging dataset and simultaneously reduce the number of annotation artifacts.
- O3 To facilitate the research on negation and coreference in the context of NLI. More specifically, we focus on contrastive and adversarial examples.

3.1 Premise Extraction

In the first step of the process, we extracted a set of candidate premises. We decided to use a single sentence premise, similar to SNLI and MNLI datasets. We defined two requirements for our premise sentences: 1) that they cover a range of different topics; and 2) that they be complex enough to entail or contradict multiple possible hypotheses.

Choice of topics As a source for premises we used the Spanish version of Wikipedia from October 2019. We chose six topics, covering five different domains: history, art, sports, technology, and politics. We also selected the topics in pairs hypothesizing that this selection might facilitate the creation of contrastive examples, specifically in the context of coreference.

- famous historical figures:
Pablo Picasso (ES: Pablo Picasso)
Christopher Columbus (ES: Cristobal Colón)
- types of “games”:
Olympic games (ES: Juegos Olímpicos)
Video games (ES: videojuegos)
- types of multinational “unions”:
The European Union (ES: Unión Europea)
The Union of Soviet Socialist Republics (ES: Unión Soviética)

Extraction process We extracted the main Wikipedia article for each topic and preprocessed it (sentence segmentation and tokenization) using Spacy (Honnibal and Montani, 2017). We split the text by paragraphs and discarded paragraphs that contained only one sentence or more than five sentences. Then, from each paragraph, we selected a single sentence, prioritizing sentences containing negation² where possible, otherwise selecting a sentence at random. We ensured that each selected sentence had a length between 15 and 45 tokens.

Post-processing At the end of the extraction process, we had 471 candidate-premise sentences as follows: 82 for Picasso, 60 for Columbus, 68 for the Olympic games, 73 for video games, 107 for the EU, and 81 for the USSR. For each sentence, we also kept the corresponding paragraph to enable experimental setup where we provide an additional context to the machine learning models at train and test time. We also used the “context paragraphs” when generating “neutral” pairs. One of the authors manually inspected all 471 candidate-premise sentences. They manually resolved problems with sentence segmentation, removed URLs and internal wikipedia document references, and explicitly resolved any coreferential and anaphorical ambiguities (i.e.: replaced pronouns and coreferential entities with an unambiguous referent).

²To check for negation, we used a simple keyword based search, using a list of the most common negative particles, adverbs, and verbs in Spanish. The list is available at <https://github.com/venelink/inferes>

3.2 Expert “Generation” Strategy

Task formulation We formulated two separate *generation* tasks: the generation of entailment pairs and generation of contradiction pairs. We defined the tasks as follows:

Entailment: *Given a premise, write two different sentences that are true.*

Contradiction: *Given a premise, write two different sentences that are false.*

Our guidelines enforced a strict definition of contradiction and required our generators to write sentences that explicitly contradict the premise, rather than implicitly rely on event and actor coreference³. We asked the generators to provide multiple examples, requiring the use of different strategies. We further instructed the corpus generators to: 1) generate hypotheses that have a low lexical overlap with the premise; 2) generate one affirmative and one negated sentence for each relation; 3) where possible, replace named entities with pronouns or other instances of coreference. Our instructions aim to encourage generators to come up with difficult and diverse examples. We also ensure a high frequency of entailment pairs containing negation and of affirmative contradiction pairs. For a reference, the readers can see an example of generated entailment and contradiction hypotheses in 4.

4. (PREMISE) En la década de 1980 el soporte habitual para el software era el cartucho en las videoconsolas, y el disco magnético o la cinta de casete en los ordenadores.⁴

(ENTAILMENT) Es poco probable que en los 1980s las videoconsolas y los ordenadores utilizaran el mismo soporte.⁵

(CONTRADICTION) Aunque el cartucho se había utilizado en el pasado, en los 80 ya se consideraba desfasado.⁶

Sentence “generators” Four graduate students of linguistics were trained for this task by the authors of the paper. They received detailed instructions, examples, and a two-hour interactive training session prior to the start of the corpus creation. The

³For a discussion of the definition of contradiction in the context of NLI, we refer the reader to Gold et al. (2019).

⁴EN: “In the 1980s, the usual medium for software was the cartridge in video consoles, and the magnetic disk or cassette tape in computers.”

⁵EN: “It is unlikely that in the 1980s video game consoles and computers used the same medium.”

⁶EN: “Although the cartridge had been used in the past, by the 1980s it was already considered outdated.”

students met with the authors of the paper on a weekly basis to discuss challenging or interesting examples. To further increase the diversity of the corpus, we recruited 24 undergraduate students for a two-hour hypothesis generation session preceded by a one-hour interactive training session. All generators were native speakers of European Spanish.

“Generated” portion of the corpus We distributed the premises extracted in Section 3.1 between the four graduate students balancing the number of entailment and contradiction pairs per topic and per premise. For any given premise, a single expert would generate only one of the relations, never both. Some premises were used more than once. Our selection strategy aimed to create maximum diversity in the data and reduce the potential bias from a relatively small number of data generators. The four graduate students created 2,284 pairs from the original 471 premises. The 24 undergraduate students generated further 872 pairs. The final corpus from the *generation* strategy contains 3,156 pairs, split equally between entailment and contradiction. We describe the process of obtaining “neutral” text pairs in Section 3.5.

3.3 Expert “Rewrite” Strategy

Task formulation The *rewrite* strategy is based on the pairs from the *generation* strategy. We defined the task as follows:

Given an existing premise–hypothesis pair, modify both the premise and the hypothesis so that:

1. *the resulting sentence has a substantial difference in meaning from original*

2. *where possible, change the negation status of a sentence. That is, an affirmative sentence would become negated, while a negated sentence would become affirmative*

3. *where possible, replace some words in the original sentences with synonyms and/or coreferential entities*

We further instructed the “rewriters” not to resort only to simple negation. An example of the *rewrite* strategy can be seen in 5. and 6.: when rewriting the premise, our expert replaced “descartaba” (EN: “ruled out”) with “aceptaba” (EN: “accepted”); when rewriting the hypothesis, they changed “inaceptable” (EN: “unacceptable”) for “viable” (EN: “feasible”). The resulting adversarial examples include lexical and morphological negation and are more complex than the “simple negation” benchmark of Hossain et al. (2020).

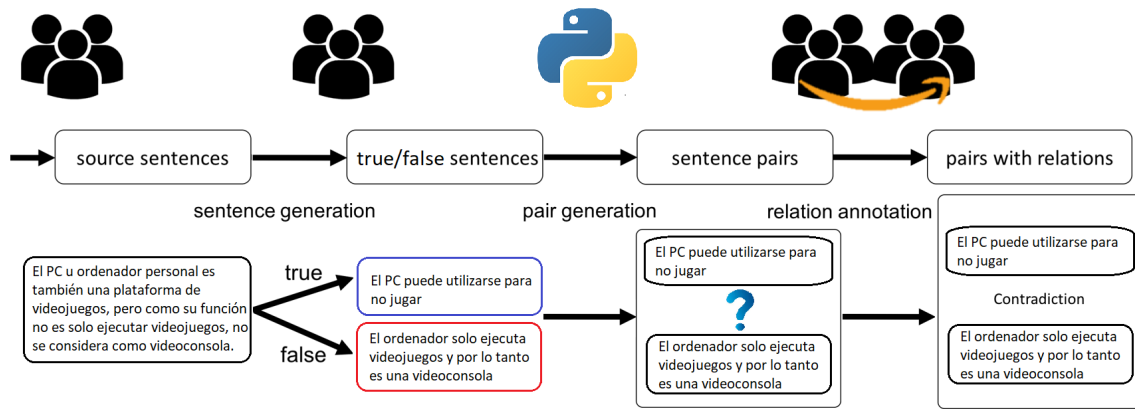


Figure 1: Pipeline for the *annotation* corpus creation strategy.

5. (PR) La reina llamó entonces a Colón, diciéndole que no **descartaba** totalmente su plan.⁷
(HYP) La reina le hizo saber a Colón que su plan no era del todo **inaceptable**.⁸
6. (PR_RW) La reina llamó entonces a Colón, diciéndole que no **aceptaba** totalmente su plan.⁹
(HYP_RW) La reina le hizo saber a Colón que su plan no era del todo **viable**.¹⁰

“Rewritten” portion of the corpus The *rewrite* process was carried out by two graduate students of linguistics. After rewriting both the premise and the hypothesis, we create three new combinations involving an original or rewritten hypothesis. In the provided example, those are the pairs 5.(Pr)–6.(Hyp_rw), 6.(Pr_rw)–5.(Hyp) and 5.(Hyp)–6.(Hyp_rw).¹¹ Our “rewriters” then annotated the relations between the new pairs (in the example, the relations are “contradiction”, “neutral”, and “entailment” respectively). As a source, we selected 20 entailment and 20 contradiction per topic, a total of 240 “generated” pairs. We ensured equal distribution of the original “generators” and created 720 new adversarial “rewrite” pairs.

3.4 Crowd “Annotation” Strategy

Task formulation For the crowd *annotation* strategy we adopted the three-step approach proposed

⁷EN: The queen then called Columbus, telling him that she did not fully **rule out** his plan.

⁸EN: The queen let Columbus know that his plan was not entirely **unacceptable**.

⁹EN: The queen then called Columbus, telling him that she did not fully **accept** his plan.

¹⁰EN: The queen let Columbus know that his plan was not entirely **feasible**.

¹¹We do not to use 5.(Pr)–6.(Pr_rw) due to sentence length.

by Gold et al. (2019). The authors first semi-automatically generated a large pool of premise-hypothesis pairs. Then, they applied stratified sub-sampling. Finally, they recruited crowd workers to annotate the meaning relations between the texts. Figure 1 illustrates the *annotation* strategy. We choose this approach since it’s compatible with our *generation* and *rewrite* strategies. We were interested in comparing and combining expert- and crowd-created corpora, which, to the best of our knowledge has not been done before for NLI.

Creating a sentence pool The first step of the process was identical to the *generation* strategy. We chose 20 of the original premises, five from Picasso, Columbus, Olympic games, and video games. We chose premises that contain multiple predicates and would allow for creativity in generating entailment and contradiction pairs. In Figure 1, these premises are called “source sentences”.¹²

We recruited 26 undergraduate students of linguistics and provided them with one-hour interactive training for the task of generating entailment and contradiction pairs. Each student generated 20 entailment and 20 contradiction pairs from the same 20 “source sentences”. In Figure 1, these hypotheses are called “true/false sentences”.^{13,14} At this step, the students generated a “sentence pool” of of 1,040 “true/false sentences”.

Pair generation In the second step we combined the sentences from the “sentence pool” in pairs using three different selection strategies. The “true-

¹²EN: “The PC or personal computer is also a video game platform, but since its function is not only to play video games, it is not considered a video game console.”

¹³EN: (true) “The PC can be used for non-gaming.”

¹⁴EN: (false) “The computer only runs video games and is therefore a game console.”

“true” strategy combines two “true” sentences derived from the same premise. The “false-false” strategy combines two “false” sentences from the same premise. The “true-false” strategy combines one true and one false sentence derived from the same premise. Unlike Gold et al. (2019), we do not include a random pairing and do not downsample “false-false” and “true-false” strategies. We randomly selected 2,000 of the pairs for annotation, ensuring equal distribution of strategies.

Pair annotation In the third step, we asked crowd workers to annotate the textual relation between pairs. We used the WARP-Text (Kovatchev et al., 2018b) annotation interface for the annotation. Following Gold et al. (2019), we created two separate binary annotation tasks - one for entailment and one for contradiction. For entailment, we included each pair twice, changing the order of P and H to reflect the directional nature of the relation. If a sentence was annotated as not-entailment and not-contradiction, we marked it as “neutral”.

We use three annotators for each example. Following prior work (Marelli et al., 2014; Gold et al., 2019), we calculated the agreement as the average % of annotators that voted for the majority label. We obtained 86.9% agreement for the “entailment” relation and 85.6% agreement for the “contradiction” relation. We also calculated the Fleiss’ kappa score, obtaining a “moderate agreement” of 55.¹⁵ Our agreement and label distribution of labels are consistent with the results reported by Gold et al. (2019) for 10 annotators. Since 56% of the pairs were labeled “neutral”, we kept all “entailment” and “contradiction” pairs and randomly downsampled the “neutral” to obtain a balance between the classes. The *annotate* portion contains 1,290 pairs.

3.5 Generating Neutral Pairs

Using our *generation* strategy, we created “entailment” and “contradiction” pairs. Using our *rewrite* strategy, we created pairs with all three relations. However the “neutral” class was underrepresented compared to the other two. To create a balanced dataset for training automated systems, we needed a separate strategy to introduce more pairs with the “neutral” label. In this subsection, we describe four different rule-based strategies that we used to generate “neutral” pairs in an automated manner.

¹⁵The lower kappa is likely due to the label imbalance in the corpus: only about 22% of the pairs contain entailment and 20% contain contradiction.

Shuffling existing P and H (same topic) We matched each premise to two random hypotheses, generated for different premises on the same topic.

Matching existing P with random contexts (same topic) In Section 3.1, we kept a “context” paragraph for each premise that we extracted. We matched each premise to two random “contexts” on the same topic. We then randomly selected a sentence from each of those contexts.

Matching existing H with random contexts (same topic) Similar to the previous strategy, we randomly matched each hypothesis to a sentence from a “context” paragraph on the same topic.

Shuffling existing P and H (different topics) Typically, the premise and the hypothesis have at least some degree of semantic similarity. We argue that an automated NLI solution should be able to label unrelated pairs. We created a small fixed number of unrelated pairs by matching texts and hypotheses from different topics.

Validating neutral pairs We selected 240 pairs, 60 from each of the four strategies, to manually validate the quality of the “neutral” pairs. One of our “sentence generators” performed a two-stage annotation on each pair. At the first stage they annotated “*whether the premise and hypothesis are semantically related*”. At the second stage they annotated “*whether the meaning relation is neutral, despite a potential semantic relatedness*”. 55% of the “neutral” pairs had some semantic relation (e.g., shared topic or named entities), and 26% had a *strong* relation. 237 out of the 240 examples (98.75%) were annotated as “neutral”. Two pairs were found to be “entailment” and one - “contradiction”.

The “neutral” portion of the corpus Through generation and downsampling, we obtain a total of 1,291 “neutral” sentences for INFERES . We use 298 of them to re-balance the *rewrite* portion of the corpus and the remaining 1,893 to complete the *generation* portion of the corpus.

4 INFERES

We combined the examples from all corpus creation strategies to create INFERES - a corpus of NLI for Spanish containing 8,055 text pairs. Table 1 shows the distribution of pairs and labels based on the creation strategy. Note that for *generation* and *rewrite* strategies, the “neutral” examples were at least in part generated automatically to ensure

label balance. For “annotate” the “neutral” pairs are naturally occurring. More than half of the corpus, 5,029 text pairs, was created using the *generate* strategy. This is the core part of the corpus, in which we have incorporated multiple strategies for ensuring the quality and the linguistic diversity of the examples. 1,716 pairs were created using the *rewrite* strategy and 1,290 pairs were generated using the *annotate* strategy. All six topics are represented roughly equally. In the *generate* and *rewrite* portions, we aimed to ensure that each original premise has the same number of hypotheses, distributed equally across relations.

Strategy	Pairs	Ent	Cnt	Neu
Full	8,055	2,399	2,687	2,969
Generate	5,029*	1,574	1,582	1,893*
Rewrite	1,716*	398	712	606*
Annotate	1,290	427	393	470

Table 1: Distribution of labels INFERES by strategy

We measured the vocabulary size and the lexical overlap between the premise and hypothesis. The full INFERES has a vocabulary size of 12,877 unique types. On average, 22.6% of the tokens from the premise also appear in the hypothesis. 33.4% of the tokens from the hypothesis also appear in the premise. The two numbers differ since we count the number of non-unique tokens, including repetition, and we normalize them using a different denominator (length of premise/hypothesis). The overlap is comparable with prior work for English (20% and 38% for MNLI and 18% and 34% for linguist-in-the-loop NLI).

5 Machine Learning Experiments

To demonstrate the utility of INFERES, we carried out a set of machine learning experiments. The design of INFERES allows us to test NLI models under a variety of conditions: standard train/test split, hypothesis-only condition, performance on negation-based adversarial examples, and performance by topic in- and out-of-distribution.

Machine learning models We used two transformer based models, pretrained for Spanish: the multilingual version of BERT (Devlin et al., 2019) and the Spanish version of BERT, BETO (Cañete et al., 2020). We used the version of the models available on HuggingFace (Wolf et al., 2020) as of May 2022 and finetuned them on INFERES. After

experimenting with different hyperparameter settings, we empirically found the best performance using a PolynomialDecay learning rate scheduler and training the model for five epochs. We kept the rest of the hyperparameters at their default values and used ADAM optimizer¹⁶. All reported results are the average of five different random initializations.

Condition	MB	mBERT	BETO
Full Dataset	36.8	69.6	72.8
Hyp. Only	36.8	38.8	42.3
Adv. Negation	41.5	52	51.2

Table 2: Performance of multilingual-BERT and Spanish BERT (BETO) on INFERES across different conditions. MB: “majority baseline”. “Full”: standard train/test split. “Hyp. Only”: hypothesis-only. “Adv. Negation”: negation-based adversarial examples.

Full corpus performance In a standard “full corpus” condition, we used 80% of INFERES for training, 10% for validation, and 10% for testing. We use the examples generated by all three strategies for both training and testing. As shown in Table 2, BETO obtained 72.8% accuracy and outperformed mBERT (69.6%). While both models reached a fair performance on the test set, the results were much lower than the “super-human” performance that state-of-the-art transformer models obtain on popular benchmarks for English. For a reference, the official performance for the Spanish portion of XNLI is 82% for BETO-cased, and 78.5% for mBERT¹⁷.

Hypothesis-only performance In the “hypothesis-only” condition, the models do not have access to the “premise” during training or testing. NLI explores the meaning relation between the two texts - the premise and the hypothesis. If a model is exposed only to one of the texts, its performance should not exceed that of a random guess, roughly equal to predicting the most common class.

Prior work has shown that existing datasets for English contain a large number of “annotation artifacts” and models are able to obtain much higher performance than chance. Poliak et al. (2018) report 55% accuracy for the “hypothesis-only” condi-

¹⁶The code for the experiments and all hyperparameters are available at <https://github.com/venelink/inferes>

¹⁷See <https://github.com/dccuchile/beto>

tion on the MNLI corpus using non-transformer models, an increase of 20% over the majority baseline. Parrish et al. (2021) show that their “linguist-in-the-loop” approach is less biased in the hypothesis-only condition, but they also report accuracy over 50% on the full dataset using a ROBERTA transformer.

In the “hypothesis-only” condition on INFERES, BETO obtained 42.3%, and mBERT – 38.8%, which is respectively 5.5% and 2% higher than the majority baseline. The relatively small improvement over the majority baseline indicates that the hypothesis-only artifacts in INFERES are substantially fewer than in previous work.

Performance on negation-based adversarial examples

For this experiment, we trained the models on the *generation* and *annotation* portion of the corpus and evaluated them on *rewrite*. The setup is similar to the one from Hossain et al. (2020). The performance of mBERT and BETO drops significantly when facing adversarial examples (See Table 2). The models obtain 52% and 51.1% accuracy, about 10% higher than the majority baseline. Hossain et al. (2020) report that on two of the three datasets they use, BERT performs worse than the majority baseline. Our experimental setup is arguably more difficult, since INFERES contains multiple negation strategies rather than just negating the main verb. The *rewrite* portion of the corpus can provide further insight into the use of negation in NLI and we hope it would facilitate further research and improvement in the area.

In- and Out-of-distribution generalization by topic

We also carried out a set of experiments to determine the ability of models to generalize across the six different topics. We evaluated the models in two different conditions.

The in-distribution (ID) condition is an extension of the “full corpus” condition. We split the full corpus containing all six topics in an 80/10/10 ratio. Then, when evaluating the performance of the models, we split the test set in six sub-sets, based on their topic and we measured the model accuracy on each sub-set. To ensure that the variation of the model performance is not due to a sampling bias, we re-trained each model five times, using a different 80/10/10 random split each time. In Table 3, we report the average accuracy across the five different splits. Both models obtained the highest ID performance on the topics of “Picasso” and “The

Top	mBERT		BETO	
	IND	OOD	ID	OOD
All	69.6		72.8	
1	73.1±4	67.9±2	78.5±4	73.8±.6
2	68.1±4	69.9±1	74.6±3	72.0±.7
3	70.8±3	70.9±1	74.2±3	73.3±.7
4	71.6±2	68.9±.7	69.4±5	69.7±.7
5	77.1±3	76.1±.5	78.3±4	77.2±.9
6	65.3±4	69.8±.9	68.6±5	69.6±.9

Table 3: In-distribution (ID) and Out-of-distribution (OOD) performance of mBERT and BETO on different topics within INFERES. ID: model trained on data covering all 6 topics. OOD: model trained on 5 topics and evaluated on the unseen 6th. 1 (Picasso), 2 (Columbus), 3 (Olympics), 4 (Videogames), 5 (EU), 6 (USSR).

European Union” and the lowest ID performance on “The USSR”.

For the out-of-distribution (OOD) condition, we designed a transfer learning experiment, in which we trained mBERT and BETO on five of the topics and evaluate on the sixth. The results presented in Table 3 demonstrate that the models are able to generalize well across the topics, even in a transfer learning setup. For both models, the OOD performance on most topics drops between 1% and 5% compared to ID. The performance on “Olympics” for mBERT and on “videogames” for BETO was almost identical between conditions. For “The USSR”, both models obtained higher performance for OOD. We inspected the matter further and noticed that due to the corpus size, the ID random split is not very stable (the ID test set only contains between 100 and 120 instances of each topic) and the average can be affected by outliers. The OOD results are much more stable due to the test set having over 1,000 examples per topic. This can be seen in the difference in standard deviation: between 2% and 5% for ID, and below 1% for OOD. Our experiments demonstrate that the models are able to generalize well even to topics they have never seen during training, a promising finding for the overall generalizability of NLI models.

6 Discussion

In Section 3, we formulated three main objectives behind INFERES. In this section, we want to revisit those objectives and briefly discuss the importance of our work for the NLP and NLI communities.

Our **first objective** was “To create a native NLI dataset for the Spanish language”. To the best of

our knowledge, INFERES is the first native NLI dataset for Spanish which is not adapted or translated. We described the creation process and validated that it can be used to train different machine learning models. We have successfully contributed a new resource to the Spanish NLP community and we hope that INFERES can facilitate the further creation of tools and resources for that language.

Our **second objective** was *“To promote better data quality and corpus creation practices.”*. We proposed, implemented, analyzed, and compared several different strategies for creating text pairs. The resulting dataset proves non-trivial to state-of-the-art NLP models with an overall accuracy in the low 70s. This leaves a lot of room for improvement and future research. The results using a “hypothesis-only” baseline indicate that INFERES contains fewer annotation artifacts than prior work. At the same time, models trained on the dataset are able to generalize well across different topics, even in our “out-of-distribution” condition. Overall, we can conclude that INFERES is of high quality and achieves the objective of promoting better data by design.

Our **third objective** was *“To facilitate the research on negation and coreference in the context of NLI.”*. Our *rewrite* strategy was focused on creating naturally occurring contrastive and adversarial examples based on negation and coreference. We followed prior work on evaluating systems’ performance and demonstrated that those examples are non-trivial to solve. However, the two models are still able to outperform the majority baseline by over 10%. These findings indicate that the problem is not unsolvable and the models are learning something about complex negation from the data. INFERES can facilitate the research of negation in Spanish both in the context of NLI and in isolation. We leave quantifying the importance of coreference in the *rewrite* section for future work.

Overall, we have achieved all our objectives: 1) we created a novel dataset for Spanish; 2) we used different generation and annotation strategies to obtain a challenging corpus with fewer annotation artifacts; 3) we created a set of high-quality contrastive and adversarial examples based on negation and coreference. We believe that INFERES is an important contribution to Spanish NLP, and also to researchers interested in NLI, negation, and coreference. We hope that this dataset can be used to train and evaluate more accurate automated sys-

tems, but also to better understand the nature of those linguistic phenomena.

7 Conclusions

We presented INFERES - a new corpus of Natural Language Inference for Spanish. To the best of our knowledge, this is the first original Spanish NLI corpus that is not a translation or an adaptation of an existing dataset. We explored several different strategies for corpus creation and put the emphasis on creating diverse and non-trivial examples, that are also linguistically interesting. More specifically, we created contrastive and adversarial examples involving complex negation and coreference.

We provided two baseline transformer-based systems finetuned on the dataset. We demonstrated that INFERES is challenging and contains fewer annotation artifacts than prior work. We also evaluated the performance of automated systems on adversarial examples and the ability of the models to generalize across topics in- and out-of-distribution. The results validated the quality and the difficulty of the corpus. INFERES leaves a room for analysis and improvement.

Our work opens several directions for future work: studying and improving the performance of NLI models for Spanish; expanding the research on negation in Spanish, and specifically the complex and lexical negation; evaluating the importance of coreference in the context of NLI. We believe that INFERES will be useful both to Spanish researchers and to the general NLP community.

Acknowledgements

We want to thank M. Ant3nia Mart3 and Montse Nofre for their suggestions and support for this work. We are grateful to Patricia Grau Francitorra, Eugenia Verjovodkina, V3ctor Bargiela Zotes, and Xavier Bonet Casals for the annotations and the discussions about the annotation process. We also want to thank the anonymous reviewers for their feedback and suggestions.

This work was partially funded by the project “FairTransNLP-Language: Analysing Toxicity and Stereotypes in Language for Unbiased, Fair and Transparent Systems (PID2021-124361OB-C33)” funded by Ministerio de Ciencia e Innovaci3n (Spain) and co-funded by the European Regional Development Fund (FEDER). Most of this work was carried out during the first author’s APIF grant at the University of Barcelona.

References

- Hossein Amirkhani, Mohammad AzariJafari, Zohreh Pourjafari, Soroush Faridan-Jahromi, Zeinab Kouhkan, and Azadeh Amirak. 2020. [Farstail: A persian natural language inference dataset](#).
- Roy Bar Haim, Ido Dagan, Bill Dolan, Lisa Ferro, Danilo Giampiccolo, Bernardo Magnini, and Idan Szpektor. 2006. The second pascal recognising textual entailment challenge.
- Luisa Bentivogli, Peter Clark, Ido Dagan, and Danilo Giampiccolo. 2009. The Fifth PASCAL Recognizing Textual Entailment Challenge. In *TAC*.
- Luisa Bentivogli, Peter Clark, Ido Dagan, and Danilo Giampiccolo. 2010. The Sixth PASCAL Recognizing Textual Entailment Challenge. In *TAC*.
- Luisa Bentivogli, Peter Clark, Ido Dagan, and Danilo Giampiccolo. 2011. The Seventh PASCAL Recognizing Textual Entailment Challenge. In *TAC*.
- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642.
- José Cañete, Gabriel Chaperon, Rodrigo Fuentes, Jou-Hui Ho, Hojin Kang, and Jorge Pérez. 2020. Spanish pre-trained bert model and evaluation data. In *PML4DC at ICLR 2020*.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2006. The PASCAL recognising textual entailment challenge. In *Machine learning challenges. evaluating predictive uncertainty, visual object classification, and recognising textual entailment*, pages 177–190. Springer.
- 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.
- Abteem Ebrahimi, Manuel Mager, Arturo Oncevay, Vishrav Chaudhary, Luis Chiruzzo, Angela Fan, John Ortega, Ricardo Ramos, Annette Rios, Ivan Vladimir Meza Ruiz, Gustavo Giménez-Lugo, Elisabeth Mager, Graham Neubig, Alexis Palmer, Rolando Coto-Solano, Thang Vu, and Katharina Kann. 2022. [AmericasNLI: Evaluating zero-shot natural language understanding of pretrained multilingual models in truly low-resource languages](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6279–6299, Dublin, Ireland. Association for Computational Linguistics.
- Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hannaneh Hajishirzi, Gabriel Ilharco, Daniel Khashabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. 2020. [Evaluating models’ local decision boundaries via contrast sets](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323, Online. Association for Computational Linguistics.
- Danilo Giampiccolo, Hoa Trang Dang, Bernardo Magnini, Ido Dagan, Elena Cabrio, and Bill Dolan. 2008. The Fourth PASCAL Recognizing Textual Entailment Challenge. In *TAC*.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. The third pascal recognizing textual entailment challenge. In *Proceedings of the ACL-PASCAL workshop on textual entailment and paraphrasing*, pages 1–9. Association for Computational Linguistics.
- Max Glockner, Vered Shwartz, and Yoav Goldberg. 2018. [Breaking NLI systems with sentences that require simple lexical inferences](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 650–655, Melbourne, Australia. Association for Computational Linguistics.
- Darina Gold, Venelin Kovatchev, and Torsten Zesch. 2019. [Annotating and analyzing the interactions between meaning relations](#). In *Proceedings of the 13th Linguistic Annotation Workshop*, pages 26–36, Florence, Italy. Association for Computational Linguistics.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Jiyeon Ham, Yo Joong Choe, Kyubyong Park, Iiji Choi, and Hyungjoon Soh. 2020. [KorNLI and KorSTS: New benchmark datasets for Korean natural language understanding](#). In *Findings of the Association for*

- Computational Linguistics: EMNLP 2020*, pages 422–430, Online. Association for Computational Linguistics.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.
- Md Mosharaf Hossain, Venelin Kovatchev, Pranoy Dutta, Tiffany Kao, Elizabeth Wei, and Eduardo Blanco. 2020. [An analysis of natural language inference benchmarks through the lens of negation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9106–9118, Online. Association for Computational Linguistics.
- Hai Hu, Kyle Richardson, Liang Xu, Lu Li, Sandra Kübler, and Lawrence Moss. 2020. [OCNLI: Original Chinese Natural Language Inference](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3512–3526, Online. Association for Computational Linguistics.
- Divyansh Kaushik, Eduard Hovy, and Zachary Lipton. 2020. [Learning the difference that makes a difference with counterfactually-augmented data](#). In *International Conference on Learning Representations*.
- Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina Williams. 2021. [Dynabench: Rethinking benchmarking in NLP](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4110–4124, Online. Association for Computational Linguistics.
- Venelin Kovatchev, Trina Chatterjee, Venkata S Govindarajan, Jifan Chen, Eunsol Choi, Gabriella Chromis, Anubrata Das, Katrin Erk, Matthew Lease, Junyi Jessy Li, Yating Wu, and Kyle Mahowald. 2022. [longhorns at DADC 2022: How many linguists does it take to fool a question answering model? a systematic approach to adversarial attacks](#). In *Proceedings of the First Workshop on Dynamic Adversarial Data Collection*, pages 41–52, Seattle, WA. Association for Computational Linguistics.
- Venelin Kovatchev, Darina Gold, M. Antonia Marti, Maria Salamo, and Torsten Zesch. 2020. [Decomposing and comparing meaning relations: Paraphrasing, textual entailment, contradiction, and specificity](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5782–5791, Marseille, France. European Language Resources Association.
- Venelin Kovatchev, M. Antònia Martí, and Maria Salamó. 2018a. [ETPC - a paraphrase identification corpus annotated with extended paraphrase typology and negation](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Venelin Kovatchev, M. Antònia Martí, and Maria Salamó. 2018b. [WARP-text: a web-based tool for annotating relationships between pairs of texts](#). In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 132–136, Santa Fe, New Mexico. Association for Computational Linguistics.
- Venelin Kovatchev, M. Antonia Marti, Maria Salamo, and Javier Beltran. 2019. [A qualitative evaluation framework for paraphrase identification](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 568–577, Varna, Bulgaria. INCOMA Ltd.
- Rahmad Mahendra, Alham Fikri Aji, Samuel Louvan, Fahrurrozi Rahman, and Clara Vania. 2021. [IndoNLI: A natural language inference dataset for Indonesian](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10511–10527, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. [A SICK cure for the evaluation of compositional distributional semantic models](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 216–223, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. [Stress test evaluation for natural language inference](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2340–2353, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Alicia Parrish, William Huang, Omar Agha, Soo-Hwan Lee, Nikita Nangia, Alexia Warstadt, Karmanya Aggarwal, Emily Allaway, Tal Linzen, and Samuel R. Bowman. 2021. [Does putting a linguist in the loop](#)

- improve NLU data collection? In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4886–4901, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Anselmo Peñas, Álvaro Rodrigo, and Felisa Verdejo. 2006. Sparte, a test suite for recognising textual entailment in spanish. In *Computational Linguistics and Intelligent Text Processing*, pages 275–286, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. Hypothesis only baselines in natural language inference. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191, New Orleans, Louisiana. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. Beyond accuracy: Behavioral testing of NLP models with CheckList. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Swarnadeep Saha, Yixin Nie, and Mohit Bansal. 2020. ConjNLI: Natural language inference over conjunctive sentences. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8240–8252, Online. Association for Computational Linguistics.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. Universal adversarial triggers for attacking and analyzing NLP. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2153–2162, Hong Kong, China. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. Superglue: A stickier benchmark for general-purpose language understanding systems. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American*
- Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- 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.