

GREEK-BERT: The Greeks visiting Sesame Street

John Koutsikakis*

Ilias Chalkidis*

Prodromos Malakasiotis

Ion Androutsopoulos

[jkoutsikakis,ihalk,ruller,ion]@aueb.gr

Department of Informatics, Athens University of Economics and Business

ABSTRACT

Transformer-based language models, such as BERT and its variants, have achieved state-of-the-art performance in several downstream natural language processing (NLP) tasks on generic benchmark datasets (e.g., GLUE, SQUAD, RACE). However, these models have mostly been applied to the resource-rich English language. In this paper, we present GREEK-BERT, a monolingual BERT-based language model for modern Greek. We evaluate its performance in three NLP tasks, i.e., part-of-speech tagging, named entity recognition, and natural language inference, obtaining state-of-the-art performance. Interestingly, in two of the benchmarks GREEK-BERT outperforms two multilingual Transformer-based models (M-BERT, XLM-R), as well as shallower neural baselines operating on pre-trained word embeddings, by a large margin (5%-10%). Most importantly, we make both GREEK-BERT and our training code publicly available, along with code illustrating how GREEK-BERT can be fine-tuned for downstream NLP tasks. We expect these resources to boost NLP research and applications for modern Greek.

CCS CONCEPTS

• **Computing methodologies** → **Neural networks; Language resources.**

KEYWORDS

Deep Neural Networks, Natural Language Processing, Pre-trained Language Models, Transformers, Greek NLP Resources

ACM Reference Format:

John Koutsikakis, Ilias Chalkidis, Prodromos Malakasiotis, and Ion Androutsopoulos. 2020. GREEK-BERT: The Greeks visiting Sesame Street. In *11th Hellenic Conference on Artificial Intelligence (SETN 2020), September 2–4, 2020, Athens, Greece*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3411408.3411440>

*Equal contribution.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SETN 2020, September 2–4, 2020, Athens, Greece

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8878-8/20/09...\$15.00

<https://doi.org/10.1145/3411408.3411440>

1 INTRODUCTION

Natural Language Processing (NLP) has entered its ImageNet [9] era as advances in transfer learning have pushed the limits of the field in the last two years [32]. Pre-trained language models based on Transformers [33], such as BERT [10] and its variants [20, 21, 38], have achieved state-of-the-art results in several downstream NLP tasks (e.g., text classification, natural language inference) on generic benchmark datasets, such as GLUE [35], SQUAD [31], and RACE [17]. However, these models have mostly targeted the English language, for which vast amounts of data are readily available. Recently, multilingual language models (e.g., M-BERT, XLM, XLM-R) have been proposed [3, 7, 18] covering multiple languages, including modern Greek. While these models provide surprisingly good performance in zero-shot configurations (e.g., fine-tuning a pre-trained model in one language for a particular downstream task, and using the model in another language for the same task without further training), monolingual models, when available, still outperform them in most downstream tasks, with the exception of machine translation, where multilingualism is crucial. Consequently, Transformer-based language models have recently been adapted for, and applied to other languages [23] or even specialized domains (e.g., biomedical [1, 4], finance [2], etc.) with promising results. To our knowledge, there are no publicly available pre-trained Transformer-based language models for modern Greek (hereafter simply Greek), for which much fewer NLP resources are available, compared to English. Our main contributions are:

- We introduce GREEK-BERT, a new monolingual pre-trained Transformer-based language model for Greek, similar to BERT-BASE [10], trained on 29 GB of Greek text with a 35k sub-word BPE vocabulary created from scratch.
- We compare GREEK-BERT against multilingual language models based on Transformers (M-BERT, XLM-R) and other strong neural baselines operating on pre-trained word embeddings in three core downstream NLP tasks, i.e., Part-of-Speech (PoS) tagging, Named Entity Recognition (NER), and Natural Language Inference (NLI). GREEK-BERT achieves state-of-the-art results in all datasets, while outperforming its competitors by a large margin (5-10%) in two of them.
- Most importantly, we make publicly available both the pre-trained GREEK-BERT and our training code, along with code illustrating how GREEK-BERT can be fine-tuned for downstream NLP tasks in Greek.¹ We expect these resources to

¹The pre-trained model is available at <https://huggingface.co/nlpaueb/bert-base-greek-uncased-v1>. The project is available at <https://github.com/nlpaueb/greek-bert>

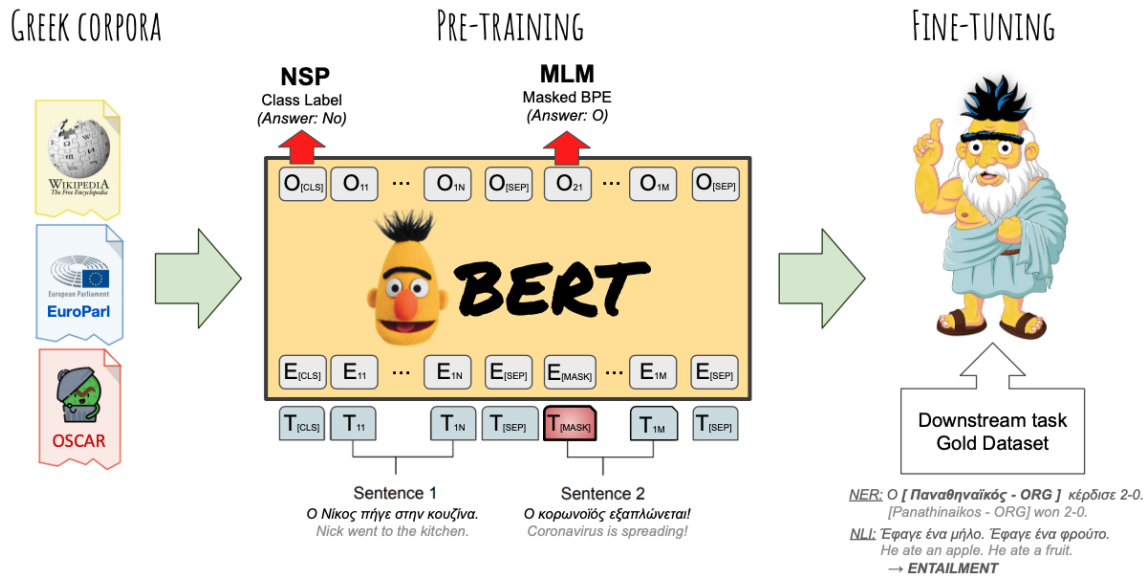


Figure 1: The two stages of employing GREEK-BERT: (a) pre-training BERT with the MLM and NSP objectives on Greek corpora, and (b) fine-tuning the pre-trained BERT model for downstream Greek NLP tasks.

boost NLP research and applications for Greek, since fine-tuning pre-trained Transformer-based language models for particular downstream tasks is the state-of-the-art.

2 RELATED WORK

2.1 BERT: A Transformer language model

Devlin et al. [10] introduced BERT, a deep language model based on Transformers [33], which is pre-trained on pairs of sentences to learn to produce high quality context-aware representations of sub-word tokens and entire sentences (Fig. 1). Each sentence is represented as a sequence of WordPieces, a variation of BPEs [11], while the special tokens [CLS] and [SEP] are used to indicate the start of the sequence and the end of a sentence, respectively. In other words, each pair of sentences has the form: $\langle [CLS], S-1, [SEP], S-2, [SEP] \rangle$, where S_1 is the first and S_2 the second sentence of the pair. BERT is pre-trained in two auxiliary self-supervised tasks: (a) *Masked Language Modelling* (MLM), also called *Denoising Language Modelling* in the literature, where the model tries to predict masked-out (hidden) tokens based on the surrounding context, and (b) *Next Sentence Prediction* (NSP), where the model uses the representation of [CLS] to predict whether S_2 immediately follows S_1 or not, in the corpus they were taken from. The original English BERT was pre-trained on two generic corpora, English Wikipedia and Children’s Books [39] with a vocabulary of 32k sub-words extracted from the same corpora. Devlin et al. [10] originally released four English models. Two of them, BERT-BASE-UNCASED (12 layers of stacked Transformers, each of 768 hidden units, 12 attention heads, 110M parameters) and BERT-LARGE-UNCASED (24 layers, each of 1024 hidden units, 16 attention heads, 340M parameters), convert all text to lower-case. The other two models, BERT-BASE-CASED and BERT-LARGE-CASED, have the exact same architectures as the corresponding previous models, but preserve character casing.

2.2 Multilingual language models

Most work on transfer learning for languages other than English focuses on multilingual language modeling to cover multiple languages at once. Towards that direction, M-BERT [10], a multilingual version of BERT, supports 100 languages, including Greek. M-BERT was pre-trained with the same auxiliary tasks as BERT (MLM, NSP), on the Wikipedias of the supported languages. Each pre-training sentence pair contains sentences from the same language, but now pairs from all 100 languages are used. To cope with multiple languages, M-BERT relies on an extended shared vocabulary of 110k sub-words. Only a small portion of the vocabulary (1,208 WordPieces or approx. 1%) applies to the Greek language, mainly because of the Greek alphabet. By contrast, languages that use the Roman alphabet (e.g., English, French) share more sub-words [8], which as shown by Lample et al. [19] greatly improves the alignment of embedding spaces across these languages. M-BERT has been mainly evaluated as a baseline for zero-shot cross-lingual training [18].

More recently, Lample and Conneau [18] introduced XLM, a multilingual language model pre-trained on the Wikipedias of 15 languages, including Greek. They reported state-of-the-art results in supervised and unsupervised machine translation and cross-lingual classification. Similarly to M-BERT, XLM was trained in two auxiliary tasks, MLM and the newly introduced *Translation Language Modeling* task (TLM). TLM is a supervised extension of MLM, where each training pair contains two sentences, S_1 and S_2 , from two different languages, S_1 being the translation of S_2 . When a word of S_1 (or S_2) is masked, the corresponding translation of that word in S_2 (or S_1) is not masked. In effect, the model learns to align representations across languages. Conneau et al. [7] introduced XLM-R, which further improved the results of XLM, without relying on the supervised TLM, by using more training data from Common Crawl, and a larger vocabulary of 250k sub-words, covering 100

languages. Again, a small portion of the vocabulary covers the Greek language (4,862 sub-words, approx. 2%).

2.3 Monolingual language models

Martin et al. [23] released CAMEMBERT, a monolingual language model for French, based on ROBERTA [21]. CAMEMBERT reported state-of-the-art results on four downstream tasks for French (PoS tagging, dependency parsing, NER, NLI), outperforming M-BERT and XLM among other neural methods. FinBERT [34] is another monolingual language model, for Finnish, based on BERT. It achieved state-of-the-art results, in PoS tagging, dependency parsing, NER, and text classification, outperforming M-BERT among other models. Monolingual Transformer-based models have also been released for other languages (e.g., Italian, German, Spanish, etc.), showing strong performance in preliminary experiments. Most of them are still under development with no published work describing them.²

2.4 NLP in Greek

Publicly available resources for Greek NLP continue to be very limited, compared to more widely spoken languages, although there have been several efforts to develop NLP datasets, tools, and infrastructure for Greek NLP.³ Deep learning resources for Greek NLP are even more limited. Recently, Outsios et al. [26] presented an evaluation of Greek word embeddings, comparing Greek Word2Vec [24] models against the publicly available Greek FastText [5] model. Neural NLP models that only pre-train word embeddings, however, have been largely superseded by deep pre-trained Transformer-based language models in the last two years. To the best of our knowledge, no Transformer-based pre-trained language model especially for Greek has been published to date. Thus, we aim to develop, study, and release such an important resource, as well as to provide an extensive evaluation across several NLP tasks, comparing against state-of-the-art models. Part of our study compares against strong neural, in most cases RNN-based, methods relying on pre-trained word embeddings. Such neural models are often not considered in other recent studies, without convincing justification. Instead, *monolingual* pre-trained Transformer-based language models are solely compared to *multilingual* pre-trained Transformer-based ones. We suspect that the latter, being usually biased towards more resource-rich languages (e.g., English, French, Spanish, etc.), may be outperformed by shallower neural models that rely only on pre-trained word embeddings. Hence, models of the latter kind may be stronger baselines for GREEK-BERT.

3 GREEK-BERT

In this work, we present GREEK-BERT, a new monolingual version of BERT for Greek. We use the architecture of BERT-BASE-UNCASED, because the larger BERT-LARGE-UNCASED architecture is computationally heavy for both pre-training and fine-tuning. We pre-trained GREEK-BERT on 29 GB of text from the following corpora: (a) the Greek part of Wikipedia;⁴ (b) the Greek part of the European Parliament Proceedings Parallel Corpus (Europarl) [14];

²An extensive list of monolingual Transformer-based models can be found at <https://huggingface.co/models>, along with preliminary results (when available).

³Consult <https://www.clarin.gr/el/content/nlpel-clarin-knowledge-centre-natural-language-processing-greece> and <http://nlp.cs.aueb.gr/software.html> for examples.

⁴An up-to-date dump can be found at <https://dumps.wikimedia.org/elwiki/>.

and (c) the Greek part of OSCAR [25], a clean version of Common Crawl.⁵ Accents and other diacritics were removed, and all words were lower-cased to provide the widest possible normalization. The same corpora were used to extract a vocabulary of 35k BPEs with the SentencePiece library [15]. Table 1 presents the statistics of the pre-training corpora. To pre-train GREEK-BERT in the MLM and NSP tasks, we used the official code provided by Google.⁶ Similarly to Devlin et al. [10], we used 1M pre-training steps and the Adam optimizer [13] with an initial learning rate of 1e-4 on a single Google Cloud TPU v3-8.⁷ Pre-training took approx. 5 days.

| Corpus | Size (GB) | Training pairs (M) | Tokens (B) |
|-----------|-----------|--------------------|------------|
| Wikipedia | 0.73 | 0.28 | 0.08 |
| Europarl | 0.38 | 0.14 | 0.04 |
| OSCAR | 27.0 | 10.26 | 2.92 |
| Total | 29.21 | 10.68 | 3.04 |

Table 1: Statistics on pre-training corpora for GREEK-BERT.

4 BENCHMARKS

We compare the performance of GREEK-BERT against strong baselines on datasets for three core NLP downstream tasks.

4.1 Part-of-Speech tagging

For the first downstream task, PoS tagging, we use the Greek Universal Dependencies Treebank (GUDT) [28, 29],⁸ which has been derived from the Greek Dependency Treebank,⁹ a resource developed and maintained by the Institute for Language and Speech Processing, Research Center ‘Athena’.¹⁰ The dataset contains 2,521 sentences split in *train* (1,622), *development* (403), and *test* (456) sets. The sentences have been annotated with PoS tags from a collection of 17 universal PoS tags (UPoS).¹¹ We ignore the syntactic dependencies of the dataset, since we consider only PoS tagging.

4.2 Named Entity Recognition

For the second downstream task, named entity recognition (NER), we use two currently unpublished datasets, developed by I. Darras and A. Romanou, during student projects at NTUA and AUEB, respectively.¹² As both datasets, are fairly small, containing 1,798 and 2,521 sentences, we merged them and eliminated duplicate sentences. The merged dataset contains 4,189 unique sentences. We use the *Person*, *Organization*, and *Location* annotations only, since the other entity types are not shared across the two datasets.

⁵<https://commoncrawl.org>

⁶<https://github.com/google-research/bert>

⁷The Google Cloud TPU v3-8 was provided for free by the TensorFlow Research Cloud (TFRC) program, to which we are grateful.

⁸https://github.com/UniversalDependencies/UD_Greek-GDT.

⁹<http://gdt.ilsp.gr>

¹⁰<http://www.ilsp.gr/>

¹¹Additional information on the curation of the dataset can be found at https://universaldependencies.org/treebanks/el_gdt/index.html.

¹²The annotated dataset of I. Darras was part of his project for Google Summer of Code 2018 (<https://github.com/eellak/gsoc2018-spacy>), while A. Romanou annotated documents with named entities for another project (<http://greekner.me/info>). We are grateful to both for sharing their datasets.

4.3 Natural Language Inference

Finally, we experiment with the Cross-lingual Natural Language Inference corpus (XNLI) [8], which contains 5,000 test and 2,500 development pairs from the MultiNLI corpus [36]. Each pair consists of a *premise* and a *hypothesis*, and the task is to decide whether the premise entails (E), contradicts (C), or is neutral (N) to the hypothesis. The test and development pairs, originally in English, were manually classified (as E, C, N) by crowd-workers, and they were then manually translated by professional translators (using the *One Hour Translation* platform) to 14 languages, including Greek. The premises and hypotheses were translated separately, to ensure that no context is added to the hypothesis that was not there originally. Hence, each pair is available in 14 languages, always with the same class label. MultiNLI also has a training set of 340k pairs. In XNLI [8], the training set of 340k pairs was automatically translated from English to the other languages, hence its quality is questionable; we discuss this further below. Although XNLI has been mainly used as a cross-lingual test-bed for multilingual models [3, 18], we only consider its Greek part, i.e., we only use Greek pairs.

5 EXPERIMENTAL SETUP

5.1 Multilingual models

For each task and dataset, we compare GREEK-BERT against XLM-R and both the cased and uncased versions of M-BERT.¹³ Recall that these models cover Greek with just a small portion of their sub-word vocabularies (approx. 2% for XLM-R and 1% for M-BERT), which may cause excessive word fragmentation.

| Model | GUdT (PoS) | NER | XNLI |
|---------------------|-------------|-------------|-------------|
| M-BERT-UNCASED [10] | 2.38 | 2.43 | 2.22 |
| M-BERT-CASED [10] | 2.58 | 2.65 | 2.40 |
| XLM-R [7] | 1.82 | 1.92 | 1.64 |
| GREEK-BERT (ours) | 1.35 | 1.33 | 1.23 |

Table 2: Word fragmentation ratio of BERT-based models calculated on the development data of all datasets. Lower ratios are better. Best results shown in bold.

Table 2 reports the word fragmentation ratio, measured as the average number of sub-word tokens per word, in the three datasets for all Transformer-based language models considered. The multilingual models tend to fragment the words more than GREEK-BERT, especially M-BERT variants, whose fragmentation ratio is approximately twice as large. XLM-R has a lower fragmentation ratio than M-BERT, as it has four times more sub-words covering Greek.

All three multilingual models often over-fragment Greek words. For example, in M-BERT-UNCASED ‘κατηγορουμενος’ becomes [‘κ’, ‘_ατ’, ‘_η’, ‘_γο’, ‘_ρου’, ‘_μενος’], and ‘γνωμη’ becomes [‘γ’, ‘_ν’, ‘_ωμη’], but both words exist in GREEK-BERT’s vocabulary. We suspect that, despite the ability of sub-words to effectively prevent out-of-vocabulary word instances, such long sequences of meaningless sub-words may be difficult to re-assemble into meaningful

¹³The BPE vocabulary of XLM-R retains both character casing and accents. We use the BASE version of XLM-R to be comparable with the rest of BERT models. The models are available at <https://huggingface.co/xlm-roberta-base>, <https://huggingface.co/bert-base-multilingual-cased>, and <https://huggingface.co/bert-base-multilingual-uncased>.

units, even when using deep pre-trained models. By contrast, baselines that operate on embeddings of entire words do not suffer from this problem, which is one more reason to compare against them.

5.2 Baselines operating on word embeddings

For both PoS tagging and NER, we experiment with an established neural sequence tagging model, dubbed BILSTM-CNN-CRF, introduced by Ma and Hovy [22]. This model initially maps each word to the corresponding pre-trained embedding (e_i), as well as to an embedding (c_i) produced from the characters of the word by a Convolutional Neural Network (CNN). The two embeddings of each word are then concatenated ($w_i = [e_i; c_i]$). Each text T is viewed as a sequence of embeddings $\langle w_1, \dots, w_{|T|} \rangle$. A stacked bidirectional LSTM [12] turns the latter to a sequence of context-aware embeddings, which is fed to a linear Conditional Random Field (CRF) [16] layer to produce the final predictions.

For NLI, we re-implemented the Decomposable Attention Model (DAM) [27], which consists of an attention, a comparison, and an aggregation component. The attention component measures the importance of each word of the premise with respect to each hypothesis word and vice versa, as the normalized (via softmax) similarity of all possible pairs of words between the hypothesis and the premise. The similarities are calculated as the dot products of the corresponding word embeddings, which are first projected through a shared MLP layer. Finally, each word of the premise (or hypothesis) is represented by an attended embedding (a_i), which is simply the similarity weighted average of the embeddings of the hypothesis (or premise). The comparison component concatenates each a_i with the corresponding initial word embedding (e_i) and projects the new representation through an MLP. In effect, the premise and the hypothesis are each represented by a set of comparison vectors ($P = \{v_1, \dots, v_p\}$, $H = \{u_1, \dots, u_h\}$, respectively). Finally, the aggregation component uses an MLP classifier for the final prediction, which operates on the concatenation of s_p with s_h , where $s_p = \sum_{v_i \in P} v_i$ and $s_h = \sum_{u_i \in H} u_i$.

5.3 Implementation details and hyper-parameter tuning

All the baselines that require pre-trained word embeddings use the Greek FastText [5] model¹⁴ to obtain 300-dimensional word embeddings. The code for all experiments on downstream tasks is written in Python with the PyTorch¹⁵ framework using the PyTorch-Wrapper library¹⁶. For the BERT-based models, we use the Transformers library from HuggingFace [37].¹⁷ The best architecture for each model is selected with grid search hyper-parameter tuning, minimizing the development loss, using early stopping without a maximum number of training epochs. For BERT models, we tuned the learning rate considering the range $\{2e-5, 3e-5, 5e-5\}$, the dropout rate in the range $\{0, 0.1, 0.2\}$, and the batch size in $\{16, 32\}$. For BILSTM-CNN-CRF, we used 2 stacked bidirectional LSTM layers and tuned the number of hidden units per layer in $\{100, 200, 300\}$, the learning rate in $\{1e-2, 1e-3\}$, the dropout rate in

¹⁴<https://fasttext.cc>

¹⁵<https://pytorch.org>

¹⁶<https://github.com/jkoutsikakis/pytorch-wrapper>

¹⁷<https://github.com/huggingface/transformers>

{0, 0.1, 0.2, 0.3}, and the batch size in {16, 32, 64}. Finally, for DAM we used 1 hidden layer with 200 hidden units for each MLP, and tuned the learning rate in {1e-2, 1e-3, 1e-4}, the dropout rate in {0, 0.1, 0.2, 0.3}, and the batch size in {16, 32, 64}. Given the best hyper-parameter values, we train each model 3 times with different random seeds and report the mean scores and unbiased standard deviation (over the 3 repetitions) on the test set of each dataset.

5.4 Denoising XNLI training data

As already discussed, XNLI includes 2,500 development pairs, 5,000 test pairs, and 340k training pairs, which have been translated from English to 14 languages, including Greek. The training pairs were machine-translated and, unfortunately, many of the resulting training pairs are of very low quality, which may harm performance. Based on this observation, we wanted to assess the effect of using the full training set, including many noisy pairs, against using a subset of the training set containing only high-quality pairs. We estimate the quality of a machine-translated pair as the perplexity of the concatenated sentences of the pair, computed by using GREEK-BERT as a language model, masking one BPE of the two concatenated sentences at a time.¹⁸ We retain the 40k training pairs (approx. 10%) with the lowest (best) perplexity scores as the high-quality XNLI training subset. For comparison, we also train (fine-tune) our models on the full XNLI training set. Unlike all other experiments, when using the full XNLI training set we do not perform three iterations (with different random seeds), because of the size of the full training set, which makes repeating experiments computationally much more expensive.

6 EXPERIMENTAL RESULTS

Table 3 reports the PoS tagging results. All Transformer-based models have comparable performance (97.8-98.2 accuracy), and XLM-R is marginally (+0.1%) better than GREEK-BERT and M-BERT-CASED. By contrast, BILSTM-CNN-CRF performs clearly worse, but the difference from the other models is small (0.8-1.2%). The fact that all models achieve high scores can be explained by the observation that the correct PoS tag of a word in Greek can often be determined by considering mostly the word’s suffix, or for short function words (e.g., determiners, prepositions) the word itself, and to a lesser extent the word’s context. Thus, even multi-lingual models with a high word fragmentation ratio (M-BERT, XLM-R) are often able to guess the PoS tags of words from their sub-words, even if the sub-words correspond to suffixes, other small parts of words, or frequent short words included in the sub-word vocabulary. Hence, there is often no need to consider the context of each word. The latter is difficult if context information gets scattered across too many very short sub-words.¹⁹ The BILSTM-CNN-CRF method, which uses word embeddings, is also able to exploit information from suffixes, because it produces extra word embeddings from the characters of the words (using a CNN).

For a more complete comparison we also report results per PoS tag for the two best models, i.e., GREEK-BERT and XLM-R (Table 4).

¹⁸See Appendix A for a selection of random noisy samples and the best and worst samples according to GREEK-BERT.

¹⁹In all multilingual models, each word is usually split in 2 or more sub-words and the last one resembles a suffix (e.g., ‘ανησυχιες’ becomes [‘ανησυχ’, ‘ιες’], ‘χαρρακτηριζεται’ becomes [‘χαρρακτηρ’, ‘ιζεται’]), which often suffices to identify PoS tags.

| Model | Accuracy |
|---------------------|--------------------|
| BILSTM-CNN-CRF [22] | 97.0 ± 0.14 |
| M-BERT-UNCASED [10] | 97.8 ± 0.03 |
| M-BERT-CASED [10] | 98.1 ± 0.08 |
| XLM-R [7] | 98.2 ± 0.07 |
| GREEK-BERT (ours) | 98.1 ± 0.08 |

Table 3: PoS tagging results (± std) on test data.

We again observe that the two models have almost identical performance. Interestingly, both models have difficulties in predicting the tag *other* (X) and to a lesser extent *proper nouns* (PROPN) and *numerals* (NUM). In fact, both models tend to confuse these PoS tags, which is reasonable considering that the inflectional morphology and context of Greek proper nouns is similar to that of common nouns, numerals (when written as words, e.g., ‘χιλιαδες’) often have similar morphology with nouns, and words tagged as ‘other’ (X) are often foreign proper nouns; for example, ‘καστρο’ could either refer to ‘Fidel Castro’ (X) or be the Greek noun for ‘castle’. Overall, there is still room for improvement in these particular tags.

| Part-of-Speech tag | GREEK-BERT | XLM-R |
|--------------------|--------------|--------------|
| ADJ | 95.6 ± 0.26 | 96.0 ± 0.17 |
| ADP | 99.7 ± 0.07 | 99.8 ± 0.03 |
| ADV | 97.2 ± 0.34 | 97.4 ± 0.12 |
| AUX | 99.9 ± 0.15 | 99.8 ± 0.20 |
| CCONJ | 99.6 ± 0.24 | 99.7 ± 0.14 |
| DET | 99.8 ± 0.08 | 99.9 ± 0.02 |
| NOUN | 97.9 ± 0.28 | 97.9 ± 0.08 |
| NUM | 92.7 ± 1.14 | 93.0 ± 0.85 |
| PART | 100.0 ± 0.00 | 99.7 ± 0.45 |
| PRON | 98.8 ± 0.25 | 98.6 ± 0.21 |
| PROPN | 86.0 ± 1.03 | 87.0 ± 0.37 |
| PUNCT | 100.0 ± 0.00 | 100.0 ± 0.03 |
| SCONJ | 99.4 ± 0.56 | 99.5 ± 0.16 |
| VERB | 99.3 ± 0.13 | 99.4 ± 0.16 |
| X | 77.3 ± 1.32 | 77.4 ± 2.16 |

Table 4: F1 scores (± std) per PoS tag for the two best models (GREEK-BERT, XLM-R). We do not report results for symbols (SYM) as there are no such annotations in the test data.

Table 5 presents the NER results. We observe that GREEK-BERT outperforms the rest of the methods, by a large margin in most cases; it is 9.3% better than BILSTM-CNN-CRF, 3.9% better than the two M-BERT models on average, and 0.9% better than XLM-R. The NER task is clearly more difficult than PoS tagging, as evidenced by the near-perfect performance of all methods in PoS tagging (Table 3), compared to the far from perfect performance of all methods in NER (Table 5). Being more difficult, the NER task leaves more space for better methods to distinguish themselves from weaker methods, and indeed GREEK-BERT clearly performs better than all the other methods, with XLM-R being the second best method.

In Table 6, we conduct a per entity type evaluation of the two best NER models. Both models are more accurate when predicting *persons* and *locations*. GREEK-BERT is better on persons, and XLM-R marginally better on locations. Concerning *organizations*,

| Model | Micro-F1 |
|---------------------|--------------------|
| BILSTM-CNN-CRF [22] | 76.4 ± 2.07 |
| M-BERT-UNCASED [10] | 81.5 ± 1.77 |
| M-BERT-CASED [10] | 82.1 ± 1.35 |
| XLM-R [7] | 84.8 ± 1.50 |
| GREEK-BERT (ours) | 85.7 ± 1.00 |

Table 5: NER results (± std) on test data.

GREEK-BERT is better, but both models struggle, because organizations often contain person names (*‘μπισκότα Παπαδοπούλου’*) or locations (*‘Αθλέτικο Μπιλιμπάο’*), shown in italics.

| Entity type | GREEK-BERT | XLM-R |
|--------------|-------------|-------------|
| PERSON | 88.8 ± 3.06 | 85.2 ± 1.25 |
| LOCATION | 88.4 ± 0.88 | 88.5 ± 0.86 |
| ORGANIZATION | 69.6 ± 4.28 | 68.9 ± 5.62 |

Table 6: F1 scores (± std) per entity type for the two best models (GREEK-BERT, XLM-R) on test data.

Finally, Table 7 shows that GREEK-BERT is again substantially better than the rest of the methods in NLI, outperforming DAM (+10.1%), the two M-BERT models (+4.9% on average) and XLM-R (+1.3%). Interestingly, performance improves for all models when trained on the entire training set, as opposed to training only on the high-quality 10% training subset, contradicting our assumption that noisy data could harm performance. Using a larger training set seems to be better than using a smaller one, even if the larger training set contains more noise. We suspect that noise may, in effect, be acting as a regularizer, improving the generalization ability of the models. A careful error analysis would shed more light on this phenomenon, but we leave this investigation for future work.

| Training Data | 10% high quality | all train data |
|---------------------|--------------------|--------------------|
| Model | Accuracy | Accuracy |
| DAM [27] | 61.5 ± 0.94 | 68.5 ± 1.71 |
| M-BERT-UNCASED [10] | 65.7 ± 1.01 | 73.9 ± 0.64 |
| M-BERT-CASED [10] | 64.6 ± 1.29 | 73.5 ± 0.49 |
| XLM-R [7] | 70.5 ± 0.69 | 77.3 ± 0.41 |
| GREEK-BERT (ours) | 71.6 ± 0.80 | 78.6 ± 0.62 |

Table 7: NLI results (± std) on test data.

As in the previous datasets, we report results per class for the two best models in Table 8. GREEK-BERT is better in all three classes, while both models have difficulties when predicting *neutral* pairs, which tend to be confused with pairs containing *contradiction*.

7 CONCLUSIONS AND FUTURE WORK

We presented GREEK-BERT, a new monolingual BERT-based language model for modern Greek, which has been pre-trained on large modern Greek corpora and can be fine-tuned (further trained) for particular NLP tasks. The new model achieves state-of-the-art performance in Greek PoS tagging, named entity recognition, and natural language inference, outperforming strong baselines in the latter two, more difficult tasks. The baselines we considered included

| Class | GREEK-BERT | XLM-R |
|---------------|-------------|-------------|
| ENTAILMENT | 78.8 ± 1.20 | 78.0 ± 0.70 |
| CONTRADICTION | 81.2 ± 0.15 | 79.7 ± 0.53 |
| NEUTRAL | 75.9 ± 0.74 | 74.1 ± 0.50 |

Table 8: F1 scores (± std) per NLI class for the two best models (GREEK-BERT, XLM-R) on test data. Both models were trained on all training data.

deep multilingual Transformer-based language models (M-BERT, XLM-R), and shallower established neural methods (BILSTM-CNN-CRF, DAM) operating on word embeddings. Most importantly, we release the pre-trained GREEK-BERT model, code to replicate our experiments, and code illustrating how GREEK-BERT can be fine-tuned for NLP tasks. We expect that these resources will boost NLP research and applications for Greek, a language for which public NLP resources, especially for deep learning, are still scarce.

In future work, we plan to pre-train another version of GREEK-BERT using even larger corpora. We plan to use the entire corpus of Greek legislation [6], as published by the National Publication Office,²⁰ and the entire Greek corpus of EU legislation, as published in Eur-Lex.²¹ Both corpora include well-written Greek text describing policies across many different domains (e.g., economy, health, education, agriculture). Following [30], who showed the importance of cleaning data when pre-training language models, we plan to discard noisy parts from all corpora, e.g., by filtering out documents containing tables or other non-natural text. We also plan to investigate the performance of GREEK-BERT in more downstream tasks, including dependency parsing, to the extent that more Greek datasets for downstream tasks will become publicly available. It would also be interesting to pre-train BERT-based models for earlier forms of Greek, especially classical Greek, for which large datasets are available.²² This could potentially lead to improved NLP tools for classical studies.

ACKNOWLEDGMENTS

This project was supported by the TensorFlow Research Cloud (TFRC) program that provided a Google Cloud TPU v3-8 for free, while we also used free Google Cloud Compute (GCP) research credits. We are grateful to both Google programs.

REFERENCES

- [1] Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew McDermott. 2019. Publicly Available Clinical BERT Embeddings. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*. USA, 72–78. <https://doi.org/10.18653/v1/W19-1909>
- [2] Dogu Araci. 2019. FinBERT: Financial Sentiment Analysis with Pre-trained Language Models. arXiv:cs.CL/1908.10063
- [3] Mikel Artetxe and Holger Schwenk. 2018. Massively Multilingual Sentence Embeddings for Zero-Shot Cross-Lingual Transfer and Beyond. *CoRR* abs/1812.10464 (2018). arXiv:1812.10464 <http://arxiv.org/abs/1812.10464>
- [4] Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. SciBERT: A Pretrained Language Model for Scientific Text. 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)*. Hong Kong, China, 3606–3611. <https://www.aclweb.org/anthology/D19-1371>

²⁰ Available at <http://www.et.gr>.

²¹ EU legislation is translated in the 24 official EU languages.

²² For example, <http://stephanus.tlg.uci.edu/>, <https://www.perseus.tufts.edu/hopper/>.

- [5] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics* 5 (2017), 135–146. https://doi.org/10.1162/tacl_a_00051
- [6] Ilias Chalkidis, Charalampos Nikolaou, Panagiotis Soursos, and Manolis Koubarakis. 2017. Modeling and Querying Greek Legislation Using Semantic Web Technologies. In *The Semantic Web*. Springer International Publishing, Cham, 591–606.
- [7] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised Cross-lingual Representation Learning at Scale. *arXiv:cs.CL/1911.02116*
- [8] 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*. Association for Computational Linguistics, Brussels, Belgium, 2475–2485. <https://doi.org/10.18653/v1/D18-1269>
- [9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- [10] 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 Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. abs/1810.04805. <https://arxiv.org/abs/1810.04805>
- [11] Philip Gage. 1994. A New Algorithm for Data Compression. *C Users Journal* 12, 2 (1994), 23–38.
- [12] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Comput.* 9, 8 (1997), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- [13] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015*, Yoshua Bengio and Yann LeCun (Eds.). USA. <http://arxiv.org/abs/1412.6980>
- [14] Philipp Koehn. 2005. Europarl: A Parallel Corpus for Statistical Machine Translation. In *Conference Proceedings: the tenth Machine Translation Summit*. AAMT, AAMT, Phuket, Thailand, 79–86. <http://mt-archive.info/MTS-2005-Koehn.pdf>
- [15] Taku Kudo and John Richardson. 2018. SentencePiece: A simple and language independent subword tokenizer and detokenizer for Neural Text Processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Brussels, Belgium, 66–71. <https://doi.org/10.18653/v1/D18-2012>
- [16] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML '01)*. Morgan Kaufmann Publishers Inc., USA, 282–289.
- [17] Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. RACE: Large-scale Reading Comprehension Dataset From Examinations. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Copenhagen, Denmark, 785–794. <https://doi.org/10.18653/v1/D17-1082>
- [18] Guillaume Lample and Alexis Conneau. 2019. Cross-lingual Language Model Pretraining. *Advances in Neural Information Processing Systems (NeurIPS)* (2019).
- [19] Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc'Aurelio Ranzato. 2018. Unsupervised Machine Translation Using Monolingual Corpora Only. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=rkYTTf-AZ>
- [20] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. *CoRR* abs/1909.11942 (2019). <http://arxiv.org/abs/1909.11942>
- [21] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR* abs/1907.11692 (2019). <http://arxiv.org/abs/1907.11692>
- [22] Xuezhe Ma and Eduard Hovy. 2016. End-to-end Sequence Labeling via Bidirectional LSTM-CNNs-CRF. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, 1064–1074. <https://doi.org/10.18653/v1/P16-1101>
- [23] Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric Villemonte de la Clergerie, Djamel Seddah, and Benoît Sagot. 2020. Camembert: a Tasty French Language Model. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- [24] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed Representations of Words and Phrases and Their Compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2 (USA) (NIPS'13)*. Curran Associates Inc., USA, 3111–3119.
- [25] Pedro Javier Ortiz Suárez, Benoît Sagot, and Laurent Romary. 2019. Asynchronous Pipeline for Processing Huge Corpora on Medium to Low Resource Infrastructures. In *7th Workshop on the Challenges in the Management of Large Corpora (CMLC-7)*. Cardiff, United Kingdom. <https://hal.inria.fr/hal-02148693>
- [26] Stamatis Outsios, Christos Karatsalos, Konstantinos Skianis, and Michalis Vazirgiannis. 2020. Evaluation of Greek Word Embeddings. In *Proceedings of The 12th Language Resources and Evaluation Conference*. France, 2543–2551. <https://www.aclweb.org/anthology/2020.lrec-1.310>
- [27] Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A Decomposable Attention Model for Natural Language Inference. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Austin, Texas, 2249–2255. <https://doi.org/10.18653/v1/D16-1244>
- [28] Prokopis Prokopidis, Elina Desypri, Maria Koutsombogera, Haris Papageorgiou, and Stelios Piperidis. 2005. Theoretical and Practical Issues in the Construction of a Greek Dependency Treebank. In *Proceedings of The Fourth Workshop on Treebanks and Linguistic Theories (TLT 2005)*, Montserrat Civit, Sandra Kubler, and Ma. Antonia Martí (Eds.). Universitat de Barcelona, Barcelona, Spain, 149–160. http://www.ilsp.gr/homepages/prokopidis/documents/gdt_tlt2005.pdf
- [29] Prokopis Prokopidis and Haris Papageorgiou. 2017. Universal Dependencies for Greek. In *Proceedings of the NoDaLiDa 2017 Workshop on Universal Dependencies (UDW 2017)*. Association for Computational Linguistics, Gothenburg, Sweden, 102–106. <http://www.aclweb.org/anthology/W17-0413.pdf>
- [30] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *arXiv e-prints* (2019). [arXiv:1910.10683](https://arxiv.org/abs/1910.10683)
- [31] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. *arXiv preprint arXiv:1606.05250* (2016). <https://nlp.stanford.edu/pubs/rajpurkar2016squad.pdf>
- [32] Ruder Sebastian, Peters Matthew E., Swayamdipta Swabha, and Wolf Thomas. 2019. Transfer learning in natural language processing. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Tutorials*. 15–18.
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *31th Annual Conference on Neural Information Processing Systems*. USA. <https://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>
- [34] Antti Virtanen, Jenna Kanerva, Rami Ilo, Jouni Luoma, Juhani Luotolahti, Tapio Salakoski, Filip Ginter, and Sampo Pyysalo. 2019. Multilingual is not enough: BERT for Finnish. [arXiv:cs.CL/1912.07076](https://arxiv.org/abs/1912.07076)
- [35] 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*. Association for Computational Linguistics, Brussels, Belgium, 353–355. <https://doi.org/10.18653/v1/W18-5446>
- [36] 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)* (New Orleans, Louisiana). Association for Computational Linguistics, 1112–1122. <https://aclweb.org/anthology/N18-1101>
- [37] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. [arXiv:cs.CL/1910.03771](https://arxiv.org/abs/1910.03771)
- [38] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. XLNet: Generalized Autoregressive Pretraining for Language Understanding. *CoRR* abs/1906.08237 (2019). <http://arxiv.org/abs/1906.08237>
- [39] Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*. IEEE Computer Society, USA, 19–27. <https://doi.org/10.1109/ICCV.2015.11>

A EXAMPLES FROM THE GREEK PART OF XNLI CORPUS

Table 9 presents examples from the Greek part of XNLI including: (a) random noisy samples with morphology and syntax errors, (b) a selection from the best samples according to GREEK-BERT language modeling perplexity, and (c) a selection from the worst samples according to GREEK-BERT language modeling perplexity.

| RANDOM NOISY SAMPLES | | |
|---|--|---------------|
| Premise | Hypothesis | Label |
| Η εννοιολογικά κρέμα κρέμα έχει δύο βασικές διαστάσεις - προϊόν και γεωγραφία. | Το προϊόν και η γεωγραφία είναι αυτά που κάνουν την κρέμα να κλέβει. | neutral |
| Ένας από τους αριθμούς μας θα μεταφέρει τις οδηγίες σας λεπτομερώς. | Ένα μέλος της ομάδας μου θα εκτελέσει τις διαταγές σας με τεράστια ακρίβεια. | entailment |
| Γκέι και λεσβίες. | Ετεροφυλόφιλους. | contradiction |
| Η ταχυδρομική υπηρεσία ήταν η μείωση της συχνότητας παράδοσης. | Η ταχυδρομική υπηρεσία θα μπορούσε να είναι λιγότερο συχνή. | entailment |
| Αυτή η ανάλυση συγκεντρωτική εκτιμήσεις από τις δύο αυτές μελέτες για την ανάπτυξη μιας λειτουργίας c-R που συνδέει το pm με τη χρόνια βρογχίτιδα. | Η ανάλυση αποδεικνύει ότι δεν υπάρχει σύνδεση μεταξύ pm και βρογχίτιδα. | contradiction |
| BEST SAMPLES ACCORDING TO GREEK-BERT | | |
| Premise | Hypothesis | Label |
| Τηγανητό κοτόπουλο, τηγανητό κοτόπουλο, τηγανητό κοτόπουλο. | Χάμπουργκερ, χάμπουργκερ, χάμπουργκερ. | contradiction |
| Τα τελευταία χρόνια, το κογκρέσο έχει λάβει μέτρα για να αλλάξει ριζικά τον τρόπο με τον οποίο οι ομοσπονδιακές υπηρεσίες κάνουν τη δουλειά τους. | Το Κογκρέσο έχει λάβει μέτρα για να αλλάξει ριζικά τον τρόπο με τον οποίο οι ομοσπονδιακές υπηρεσίες κάνουν τη δουλειά τους τα τελευταία χρόνια. | entailment |
| Για παράδειγμα, ορισμένες ηλικιακές ομάδες φαίνεται να είναι πιο ευαίσθητες στην ατμοσφαιρική ρύπανση από άλλες. | Η ατμοσφαιρική ρύπανση δεν μπορεί να επηρεάσει όλες τις ηλικιακές ομάδες. | contradiction |
| Οι επισκέπτες μπορούν να δουν τα δελφίνια να εκπαιδεύονται και να τρέφονται κάθε δύο ώρες από τις 10:00 το πρωί έως τις 4:00μ.μ. | Μπορείτε επίσης να δείτε τα δελφίνια να καθαρίζονται στις 6:00μ.μ. | neutral |
| Ναι, δεν ξέρω, όπως είπα, πιστεύω ότι πιστεύω στην θανατική ποινή, αλλά αν καθόμουν στους ενόρκους και έπρεπε να πάρω την απόφαση, δεν θα ήθελα να είμαι αυτός που θα τα καταφέρει. | Πιστεύω στην θανατική ποινή, αλλά δεν θα ήθελα να είμαι σε θέση να κάνω αυτή την επιλογή. | entailment |
| WORST SAMPLES ACCORDING TO GREEK-BERT | | |
| Premise | Hypothesis | Label |
| Τα μάτια του θολή με δάκρυα και αυτός σίδερο στο πίσω μέρος του λαιμού του. | Έτσιλαγε. | contradiction |
| Um-Βουητό πιο εσωτερική. | Συνηθισμένη | contradiction |
| Bonifacio | Επισκεφτείτε το bonifacio δωρεάν. | neutral |
| Μπόστον σέλτικς δεξιά | Ιντιάνα πέισερς, όχι; | entailment |
| Για μένα ο juneteenth ανακάλεσε ειδικά τον αβεσσαλώμ, τον αβεσσαλώμ! | Juneteenth ειδικά υπενθύμισε αβεσσαλώμ | entailment |

Table 9: Examples of pairs in the Greek part of the XNLI corpus.