

ALBETO and DistilBETO: Lightweight Spanish Language Models

José Cañete^{1,3}, Sebastián Donoso^{1,3}, Felipe Bravo-Marquez^{1,3,4},
Andrés Carvallo^{2,3}, Vladimir Araujo^{2,3,4,5}

¹ Department of Computer Science, University of Chile, Santiago, Chile

² Pontificia Universidad Católica de Chile, Santiago, Chile

³ National Center for Artificial Intelligence (CENIA), Santiago, Chile

⁴ Millennium Institute for Foundational Research on Data (IMFD), Santiago, Chile

⁵ KU Leuven, Leuven, Belgium

jcanete@dcc.uchile.cl

Abstract

In recent years there have been considerable advances in pre-trained language models, where non-English language versions have also been made available. Due to their increasing use, many lightweight versions of these models (with reduced parameters) have also been released to speed up training and inference times. However, versions of these lighter models (e.g., ALBERT, DistilBERT) for languages other than English are still scarce. In this paper we present ALBETO and DistilBETO, which are versions of ALBERT and DistilBERT pre-trained exclusively on Spanish corpora. We train several versions of ALBETO ranging from 5M to 223M parameters and one of DistilBETO with 67M parameters. We evaluate our models in the GLUES benchmark that includes various natural language understanding tasks in Spanish. The results show that our lightweight models achieve competitive results to those of BETO (Spanish-BERT) despite having fewer parameters. More specifically, our larger ALBETO model outperforms all other models on the MLDoc, PAWS-X, XNLI, MLQA, SQAC and XQuAD datasets. However, BETO remains unbeaten for POS and NER. As a further contribution, all models are publicly available to the community for future research.

Keywords: ALBERT, BERT, DistilBERT, Efficient Models, Language Models

1. Introduction

The area of Natural Language Processing has gained a lot of interest in recent years due to the great success of deep neural networks in different tasks such as sentiment analysis (Socher et al., 2013; Rosenthal et al., 2017; Mohammad et al., 2018), sentence (Warstadt et al., 2019) and document classification (Schwenk and Li, 2018; Carvallo et al., 2020), natural language inference (Williams et al., 2017; Conneau et al., 2018), among others.

This success can be explained by a large number of trainable parameters of these models, a considerable amount of training data, and the processing capacity power available nowadays, allowing them to learn complex relationships and cast complex functions.

One of these highly successful models is BERT (Devlin et al., 2019). This model, based on a Transformer (Vaswani et al., 2017) encoder, is pre-trained with lots of unlabeled text data and fine-tuned to specific tasks. BERT, in its large version, has 330 million parameters. Moreover, one of the most larger models is GPT-3 (Brown et al., 2020) which has 175 billion parameters. A problem with these large models is the difficulty of getting them into production on real-time applications (web services) or hardware-restricted devices, such as mobile devices, because of their high memory and computation requirements. Although there have been efforts for solving this issue proposing lighter models such as ALBERT (Lan et al., 2020), DistilBERT (Sanh et al., 2020), TinyBERT (Jiao et al., 2020), or PKD-

BERT (Sun et al., 2019), these versions of models have been trained exclusively in English.

On the other hand, there are recent efforts in the community to provide trained models in Spanish, such as Spanish-BERT (BETO) (Cañete et al., 2020) and Spanish-RoBERTa (Gutiérrez-Fandiño et al., 2022; De la Rosa et al., 2022; Pérez et al., 2021). However, there is still a lack of lighter and more efficient models for the Spanish language.

In this paper we present “A Lite version of BETO” (ALBETO) and a “Distill version of BETO” (DistilBETO) in order to democratize the use of these models in Spanish-speaking regions. We follow the same pre-training methodology as ALBERT and DistilBERT but with the difference that our models were trained exclusively with a corpus in Spanish.

We present an evaluation of all models on a variety of tasks and we also release the code¹ used to fine-tune the models on the tasks as well as the fine-tuned models². Our results indicate that in specific tasks, namely Part of Speech (POS) and Named Entity Recognition (NER), having a lighter version of the language models implies a slight reduction in the performance compared to the standard version of BETO. However, the lightweight models’ performance improves despite having fewer parameters in other tasks, such as Natural Language Inference, Paraphrase Identification, Docu-

¹<https://github.com/OpenCENIA/lightweight-spanish-language-models>

²<https://huggingface.co/CenIA>

ment Classification and Question Answering. In general, by fine-tuning our ALBETO models, we achieved 87% of BETO performance while having 22 times fewer parameters in the case of ALBETO *tiny* and 97% of BETO performance at 9 times that size in the case of ALBETO *base*.

The remainder of this article is structured as follows. In Section 2, we provide a review of related work. In Section 3, our models and the data on which they are trained are described. In Section 4, we present the experiments we conducted to evaluate the proposed resources and discuss results. Finally, in Section 5 the main findings and conclusions are discussed.

2. Related Work

2.1. Multilingual and Monolingual BERT

Multilingual models are models that are trained simultaneously using data from several languages. An example of these models is mBERT (Pires et al., 2019a), a version of the BERT model trained in English and 104 languages. Interestingly, some studies (Wu and Dredze, 2019; Pires et al., 2019b) have shown how these multi-language learners set strong baselines for non-English tasks.

Besides, several single language BERT models have been released. For instance, CamemBERT (Martin et al., 2019) and FlauBERT (Le et al., 2019) for French, BERTje (de Vries et al., 2019) and RobBERT (Delobelle et al., 2020) for Dutch, FinBERT (Virtanen et al., 2019) for Finish, among others. Those models have shown better performance than multilingual ones, demonstrating the importance of having them for language-specific tasks.

In the case of the Spanish language, we can find BETO (Cañete et al., 2020), which is a BERT base model, and a family of RoBERTa models. RoBERTa-BNE (Gutiérrez-Fandiño et al., 2022), which is available in base and large versions was trained on the corpus crawled by the National Library of Spain. BERTIN (De la Rosa et al., 2022) is a base sized model that was trained on the Spanish portion of mC4 (Xue et al., 2020). RoBERTuito (Pérez et al., 2021) is a RoBERTa base model trained exclusively on tweets. However, there is a lack of more efficient, smaller, or lighter language models exclusively trained with Spanish data.

2.2. Efficient Transformer-based Language Models

Several architectures have been designed to avoid over-computing in Transformer-based models. One such strategy is to use lightweight architectures that are trained from scratch. As an example, ALBERT (Lan et al., 2020) proposes a factorization of the embedding layer and a cross-layer parameter sharing as a way to improve model efficiency. Similar methodologies have also been previously explored by the Universal Transformer (Bai et al., 2019), and Deep Equilibrium Mod-

els (Dehghani et al., 2019), demonstrating the effectiveness of weight-tied Transformers.

A second strategy is to distill the knowledge of pre-trained models into smaller models. For instance, PKD-BERT (Sun et al., 2019), TinyBERT (Jiao et al., 2020), MobileBERT (Sun et al., 2020) and DistilBERT (Sanh et al., 2020) compress the knowledge of large models (teachers) into lighter ones (students). In this way, a compact model is obtained with a performance usually close to the original one.

While these approaches reduce the computational demand for execution, they are only available for the English language. In the same spirit, as pre-trained language models are increasingly common, it is essential to have these efficient alternatives available in other languages and thus alleviate excessive computational consumption.

3. Data and Models

3.1. Data

The data³ used to train all models was the same as that used to train BETO (Cañete et al., 2020), which is an updated version of the dataset proposed by Cardellino (2016). This dataset has approximately 3 billion words which includes all Spanish Wikipedia and almost all the Spanish portion of the OPUS Project (Tiedemann, 2012).

3.2. Spanish ALBERT (ALBETO)

ALBERT (Lan et al., 2020) is a more efficient BERT-style model in terms of parameters because it uses the weight-tied strategy, which means to share all parameters across layers of the model. We not only train a model comparable to BETO, but we also train larger models. We introduce 5 ALBETO models: *tiny*, *base*, *large*, *xlarge* and *xxlarge*. All except from the *tiny* one are similar in size to those trained by Lan et al. (2020). ALBETO *tiny* has the same size and configuration as Chinese ALBERT *tiny*⁴. The complete configuration of every model can be seen on Table 1.

These five models share a vocabulary of 31K lowercase tokens, that was constructed using SentencePiece (Kudo and Richardson, 2018) over the training dataset. We trained all the ALBETO models using the LAMB optimizer (You et al., 2019) following all the guidelines used by the authors to obtain better results. We trained each model using a single TPU v3-8. The details about the training of each model can be seen in Table 2.

3.3. Spanish DistilBERT (DistilBETO)

We trained the DistilBETO model using the distillation technique to transfer the knowledge of the BETO model to this new model following the work of Sanh et al. (2020). The architecture of DistilBETO is based

³<https://github.com/josecannete/spanish-corpora>

⁴<https://github.com/ckiplab/ckip-transformers>

Model	Parameters	Layers	Hidden	Embedding
ALBETO <i>tiny</i>	5M	4	312	128
ALBETO <i>base</i>	12M	12	768	128
ALBETO <i>large</i>	18M	24	1024	128
ALBETO <i>xlarge</i>	59M	24	2048	128
ALBETO <i>xxlarge</i>	223M	12	4096	128

Table 1: The configurations of each ALBETO model trained on this work.

on BETO but we removed the *token-type embeddings*, the *pooler layer* and decreased the number of layers from 12 to 6. Due to the limitations of time and hardware that do not allow us to test different architectures, we chose the one that obtained the best results according to the different ablation experiments by Sanh et al. (2020). DistilBETO was trained during 90k steps using a single GPU NVIDIA RTX 3090 following the best model training guidelines obtained by Sanh et al. (2020).

4. Evaluation

4.1. Tasks

We evaluated all our models on 6 tasks, which are all part of GLUES (Cañete et al., 2020), and we also closely follow the experimental setup proposed by them. We next describe the tasks.

4.1.1. Document Classification

Document classification is the task of assigning an entire document to an appropriate category. The categories depend on the chosen dataset and can range from topics. For this task we are using the Spanish part of MLDoc corpus (Schwenk and Li, 2018) which uses a subset of the Reuters Corpus (Lewis et al., 2004).

4.1.2. Part of Speech

Part-of-speech tagging (POS tagging) is a sequence labeling task that consists of tagging words in a text with their corresponding syntactic categories or part-of-speech. A part of speech is a category of words with similar grammatical properties. Common parts of speech are noun, verb, adjective, adverb, pronoun, preposition, conjunction, etc. The dataset used for this task is the Spanish subset of Universal Dependencies (v1.4) Treebank (Nivre et al., 2016).

4.1.3. Named Entity Recognition

Named Entity Recognition (NER) is another sequence labeling task, in which one tries to label entities in the text with their corresponding type, which can be names of people, organizations, places and miscellaneous items. Approaches typically use the BIO notation, which differentiates the beginning (B) and the inside (I) of entities. O is used for tokens that are not entities. For Named-Entity Recognition we are using the Spanish part of the Shared Task of CoNLL-2002 (Sang, 2002).

4.1.4. Paraphrase Identification

The task of Paraphrase Identification consists of verifying whether two sentences are semantically equivalent or not. We are using PAWS-X (Yang et al., 2019) dataset, which is a multilingual version of PAWS (Zhang et al., 2019) dataset. The dataset provides a machine translated train set and professionally translated development and test sets for different languages. We used the Spanish portion of it.

4.1.5. Natural Language Inference

Natural language inference is the task of determining whether a “hypothesis” is true (entailment), false (contradiction), or undetermined (neutral) given a “premise”. For Natural Language Inference we are using the Spanish part of XNLI (Conneau et al., 2018) which has a training set that is a machine translation of MultiNLI (Williams et al., 2017) and development and test set that were professional translated.

4.1.6. Question Answering

The task of Question Answering consists of, given a context and a question about that context, highlighting the sequence of words within that context that answers the question. This task uses two metrics to evaluate performance, which are F1 Score and Exact Match. For this task we considered four different datasets: MLQA (Lewis et al., 2019), TAR (Carrino et al., 2019), XQuAD (Artetxe et al., 2019) and SQAC (Gutiérrez-Fandiño et al., 2022). MLQA and TAR datasets offer a different machine translated version of SQuAD v1.1 (Rajpurkar et al., 2016) to Spanish with train and development sets. Also, MLQA offers a development set and test set professionally translated to Spanish. XQuAD offers a different professionally translated test set for Spanish. SQAC is notably the only one which was created exclusively for Spanish and offers train, development and test sets. In our experiments the machine translated development set of MLQA is left out in favor of the professionally translated set. And in the case of XQuAD which only offers test set, we used TAR as train and development sets.

4.2. Fine-tuning

In order to have a fair comparison, we fine-tuned all our models as well as BETO in both versions, *uncased* and *cased*, using the same code for all of them, which uses PyTorch (Paszke et al., 2019) and the HuggingFace’s Transformers library (Wolf et al., 2019). We follow the procedure described next.

We fine-tuned our models using the standard way proposed by (Devlin et al., 2019). The only preprocessing performed is tokenization according to the token vocabulary of each model which converted words into subwords and added the special tokens [CLS], [PAD] and [SEP] to each sentence. We set the maximum length of an input sentence to 512 tokens in all models and sentences were truncated to it when larger.

Model	Learning Rate	Batch Size	Warmup Ratio	Warmup Steps	Total Steps	Training Time (days)
ALBETO <i>tiny</i>	1.25e-3	2,048	1.25e-2	125,000	8,300,000	58.2
ALBETO <i>base</i>	8.83e-4	960	6.25e-3	53,333	3,650,000	70.4
ALBETO <i>large</i>	6.25e-4	512	3.12e-3	12,500	1,450,000	42.0
ALBETO <i>xlarge</i>	3.12e-4	128	7.81e-4	6,250	2,775,000	64.2
ALBETO <i>xxlarge</i>	3.12e-4	128	7.81e-4	3,125	1,650,000	70.7

Table 2: Training details of all ALBETO models, which were trained using a single TPU v3-8 each one.

Model	POS	NER
BETO <i>uncased</i>	97.70	83.76
BETO <i>cased</i>	98.84	88.24
DistilBETO	97.50	81.19
ALBETO <i>tiny</i>	97.04	75.11
ALBETO <i>base</i>	98.08	83.35
ALBETO <i>large</i>	97.87	83.72
ALBETO <i>xlarge</i>	98.06	82.30
ALBETO <i>xxlarge</i>	98.35	84.36

Table 3: Comparison of ALBETO, DistilBETO and BETO models on the test set for the tasks of POS and NER, which are sequence tagging tasks and are evaluated using F1 score as metric.

Model	MLDoc	PAWS-X	XNLI
BETO <i>uncased</i>	96.38	84.25	77.76
BETO <i>cased</i>	96.65	89.80	81.98
DistilBETO	96.35	75.80	76.59
ALBETO <i>tiny</i>	95.82	80.20	73.43
ALBETO <i>base</i>	96.07	87.95	79.88
ALBETO <i>large</i>	92.22	86.05	78.94
ALBETO <i>xlarge</i>	95.70	89.05	81.68
ALBETO <i>xxlarge</i>	96.85	89.85	82.42

Table 4: Comparison of ALBETO, DistilBETO and BETO models on the test set for the tasks of MLDoc, PAWS-X and XNLI. These tasks are treated as sentence classification tasks and use the accuracy as evaluation metric.

We did a hyperparameter search using combinations of batch size {16, 32, 64}, learning rate {1e-5, 2e-5, 3e-5, 5e-5} and number of epochs {2, 3, 4} on BETO, DistilBETO, ALBETO *tiny* and *base*.

For the larger models (ALBETO *large*, *xlarge* and *xxlarge*), due to numerical instability in the training process we used smaller learning rates {1e-6, 2e-6, 3e-6, 5e-6}.

We used between one and two GPUs NVIDIA RTX 3090 for fine-tuning depending on the model and task. To fine-tune the largest models on QA we used two GPUs NVIDIA A100 from the Patagón supercomputer (of Chile, 2021). In some cases the memory of the GPUs was not enough so we used gradient accumulation to reach the desired batch size. We then selected the models with best results on the development set.

4.3. Results

Table 3 shows the results of all models evaluated on POS and NER tasks in terms of F1 score.

It can be seen that for the POS, BETO *cased* surpasses ALBETO *xxlarge*, the second-best, by a 0.4 percentual difference (pd). Similar behavior is observed for the NER task where the BETO *cased* model outperforms ALBETO *xxlarge* by a considerable margin of 4.5 pd. This can be explained by the type of task since identifying entities that usually start with capital letters allows the *cased* model to have additional knowledge to solve this task, compared to other models.

Concerning the results of the performance of the models on MLDoc, PAWS-X and XNLI tasks in terms of accuracy are shown in Table 4. It is observed that, in general, the largest ALBETO models (ALBETO *xxlarge*) outperform other models by a small margin. This behavior can be explained since this model has more capacity than the other presented models. In the MLDoc task, ALBETO *xxlarge* surpasses BETO *cased* by 0.2 pd. Regarding the PAWS-X task, ALBETO *xxlarge* also outperforms the BETO *cased* model by a small margin of 0.05 pd. Finally, for XNLI, ALBETO *xxlarge* achieves the best performance with an F1 score of 82.42, surpassing BETO *cased*, the second-best, by 0.5 pd.

The results on the task of Question Answering are showed in Table 5. We can observe that generally the models on this task follow the same behaviour of that in the sentence classification tasks, with ALBETO *xxlarge* or ALBETO *xlarge* outperforming BETO by a considerable margin. In the MLQA dataset, ALBETO *xxlarge* surpasses all other models, with a difference of 2.5 pd in F1 score and 2.6 pd in exact match, respect to BETO *cased*, the third-best. In the SQAC dataset, ALBETO *xxlarge* outperforms BETO *cased* by 2.8 pd in F1 score and 1.7 pd in exact match. In the case of XQuAD (trained on TAR), ALBETO *xlarge* is the best one, followed by ALBETO *xxlarge* and BETO *cased*, with a difference of 2.3 pd in F1 score and 2.7 pd in exact match between ALBETO *xlarge* and BETO *cased*.

4.4. Model Size and Task Performance

Table 6 shows a comparison of all models in terms of model size and overall task performance. We average the results of each dataset to have a score of each model in the same way of the GLUE (Wang et al., 2018) benchmark does. For the Question Answering task, where we have two different metrics, we first

Model	MLQA	SQAC	TAR, XQuAD
BETO <i>uncased</i>	64.12 / 40.83	72.22 / 53.45	74.81 / 54.62
BETO <i>cased</i>	67.65 / 43.38	78.65 / 60.94	77.81 / 56.97
DistilBETO	57.97 / 35.50	64.41 / 45.34	66.97 / 46.55
ALBETO <i>tiny</i>	51.84 / 28.28	59.28 / 39.16	66.43 / 45.71
ALBETO <i>base</i>	66.12 / 41.10	77.71 / 59.84	77.18 / 57.05
ALBETO <i>large</i>	65.56 / 40.98	76.36 / 56.54	76.72 / 56.21
ALBETO <i>xlarge</i>	68.26 / 43.76	78.64 / 59.26	80.15 / 59.66
ALBETO <i>xxlarge</i>	70.17 / 45.99	81.49 / 62.67	79.13 / 58.40

Table 5: Comparison of ALBETO, DistilBETO and BETO models on the task of QA. We show the results of the test set in each case. The task uses two metrics which are showed as F1 Score / Exact Match.

Model	Parameters	Evaluation Average	Size	Performance
BETO <i>uncased</i>	110M	77.48	1x	0.95x
BETO <i>cased</i>	110M	81.02	1x	1x
DistilBETO	67M	73.22	1.64x	0.90x
ALBETO <i>tiny</i>	5M	70.86	22x	0.87x
ALBETO <i>base</i>	12M	79.35	9.16x	0.97x
ALBETO <i>large</i>	18M	78.12	6.11x	0.96x
ALBETO <i>xlarge</i>	59M	80.20	1.86x	0.98x
ALBETO <i>xxlarge</i>	223M	81.34	0.49x	1x

Table 6: Comparison of different models in terms of task performance and size. On the last two columns of the table we present a direct comparison of BETO (in particular the *cased* model, which is the best of the two versions) with all the other models in terms of size (parameter count) and task performance (evaluation average of all tasks).

average them to have a single score for every dataset. We can notice that in terms of parameters, ALBETO *tiny* is the smaller model on the comparison while ALBETO *xxlarge* being the largest one. In terms of overall task performance ALBETO *xxlarge* and BETO *cased* are the most competitive ones.

The last two columns of the table also compare all the models with the size and performance of BETO *cased*. Some of the most promising results are ALBETO *tiny*, which is able to retain 87% of BETO performance while being 22 times smaller in size and ALBETO *base*, which has 97% of BETO performance while having more than 9 times less parameters.

5. Conclusion

This paper presents DistilBETO and five ALBETO models (*tiny*, *base*, *large*, *xlarge*, and *xxlarge*), comprising six new pre-trained language models based on the ALBERT and DistilBERT architectures for the Spanish language. While the ALBETO models exploit the weight-tied strategy to be more parameter efficient, DistilBETO was built by compressing BETO into a lighter model using a knowledge distillation approach. We also comprehensively evaluated each proposed model fine-tuned on a set of NLP tasks for Spanish (POS, NER, MLDOC, PAWS-X, XNLI, MLQA, SQAC and XQuAD). Our results indicate that the proposed models are competitive with the current models available for Spanish and are much more efficient in

their number of parameters.

We hope this work will expand the availability of pre-trained language models based on the Spanish language to gather a wider NLP community, including researchers, developers, and students.

We envision several avenues of future research. First, we expect to evaluate these models on more tasks to increase the coverage of GLUES, which is our current evaluation benchmark (Cañete et al., 2020). In addition, we want to explore whether the modifications made to the DistilBETO architecture according to (Sanh et al., 2020) are the best for Spanish models. Also, we also want to further analyze the fine-tuning and inference speed of these models. In addition, we plan to release more distilled models fine-tuned explicitly for many NLP tasks. Finally, we intend to study the capabilities of these resources in few-shot and zero-shot learning settings.

6. Acknowledgements

This work was supported in part by the National Center for Artificial Intelligence CENIA FB210017, Basal ANID and by the Patagón supercomputer of Universidad Austral de Chile (FONDEQUIP EQM180042). Felipe Bravo-Marquez was supported by ANID FONDECYT grant 11200290, U-Inicia VID Project UI-004/20 and ANID -Millennium Science Initiative Program - Code ICN17.002.

7. Bibliographical References

- Artetxe, M., Ruder, S., and Yogatama, D. (2019). On the cross-lingual transferability of monolingual representations. *CoRR*, abs/1910.11856.
- Bai, S., Kolter, J. Z., and Koltun, V. (2019). Deep equilibrium models. In H. Wallach, et al., editors, *Advances in Neural Information Processing Systems 32*, pages 690–701. Curran Associates, Inc.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Cardellino, C. (2016). Spanish Billion Words Corpus and Embeddings, March.
- Carrino, C. P., Costa-jussà, M. R., and Fonollosa, J. A. (2019). Automatic spanish translation of the squad dataset for multilingual question answering. *arXiv preprint arXiv:1912.05200*.
- Carvalho, A., Parra, D., Lobel, H., and Soto, A. (2020). Automatic document screening of medical literature using word and text embeddings in an active learning setting. *Scientometrics*, 125(3):3047–3084.
- Cañete, J., Chaperon, G., Fuentes, R., Ho, J.-H., Kang, H., and Pérez, J. (2020). Spanish pre-trained bert model and evaluation data. In *PML4DC at ICLR 2020*.
- Conneau, A., Lample, G., Rinott, R., Williams, A., Bowman, S. R., Schwenk, H., and Stoyanov, V. (2018). Xnli: Evaluating cross-lingual sentence representations. *arXiv preprint arXiv:1809.05053*.
- De la Rosa, J., Ponferrada, E. G., Romero, M., Villegas, P., de Prado Salas, P. G., and Grandury, M. (2022). Bertin: Efficient pre-training of a spanish language model via perplexity sampling. *Procesamiento del Lenguaje Natural*, 68:13–23.
- de Vries, W., van Cranenburgh, A., Bisazza, A., Caselli, T., van Noord, G., and Nissim, M. (2019). Bertje: A dutch bert model. *arXiv preprint arXiv:1912.09582*.
- Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., and Kaiser, L. (2019). Universal transformers. In *International Conference on Learning Representations*.
- Delobelle, P., Winters, T., and Berendt, B. (2020). Robbert: a dutch roberta-based language model. *arXiv preprint arXiv:2001.06286*.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Gutiérrez-Fandiño, A., Armengol-Estapé, J., Pàmies, M., Llop-Palao, J., Silveira-Ocampo, J., Carrino, C. P., Armentano Oller, C., Rodríguez Penagos, C., Gonzalez-Agirre, A., and Villegas Montserrat, M. (2022). Maria: Spanish language models.
- Jiao, X., Yin, Y., Shang, L., Jiang, X., Chen, X., Li, L., Wang, F., and Liu, Q. (2020). Tinybert: Distilling bert for natural language understanding.
- Kudo, T. and Richardson, J. (2018). Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*.
- Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., and Soricut, R. (2020). Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*.
- Le, H., Vial, L., Frej, J., Segonne, V., Coavoux, M., Lecouteux, B., Allauzen, A., Crabbé, B., Besacier, L., and Schwab, D. (2019). Flaubert: Unsupervised language model pre-training for french. *arXiv preprint arXiv:1912.05372*.
- Lewis, D. D., Yang, Y., Russell-Rose, T., and Li, F. (2004). Rcv1: A new benchmark collection for text categorization research. *Journal of machine learning research*, 5(Apr):361–397.
- Lewis, P., Oğuz, B., Rinott, R., Riedel, S., and Schwenk, H. (2019). Mlqa: Evaluating cross-lingual extractive question answering. *arXiv preprint arXiv:1910.07475*.
- Martin, L., Muller, B., Suárez, P. J. O., Dupont, Y., Romary, L., de la Clergerie, É. V., Seddah, D., and Sagot, B. (2019). Camembert: a tasty french language model. *arXiv preprint arXiv:1911.03894*.
- Mohammad, S., Bravo-Marquez, F., Salameh, M., and Kiritchenko, S. (2018). Semeval-2018 task 1: Affect in tweets. In *Proceedings of the 12th international workshop on semantic evaluation*, pages 1–17.
- Nivre, J., Agić, Ž., Ahrenberg, L., Aranzabe, M. J., Asahara, M., Atutxa, A., Ballesteros, M., Bauer, J., Bengoetxea, K., Berzak, Y., Bhat, R. A., Bick, E., Börstell, C., Bosco, C., Bouma, G., Bowman, S., Cebiroğlu Eryiğit, G., Celano, G. G. A., Chalub, F., Çöltekin, Ç., Connor, M., Davidson, E., de Marneffe, M.-C., Diaz de Ilaraza, A., Dobrovoljc, K., Dozat, T., Droganova, K., Dwivedi, P., Eli, M., Erjavec, T., Farkas, R., Foster, J., Freitas, C., Gajdošová, K., Galbraith, D., Garcia, M., Gärdenfors, M., Garza, S., Ginter, F., Goenaga, I., Gojenola, K., Gökırmak, M., Goldberg, Y., Gómez Guinovart, X., Gonzáles Saavedra, B., Grioni, M., Grūzītis, N., Guillaume, B., Hajič, J., Hà Mý, L., Haug, D., Hladká, B., Ion, R., Irimia, E., Johannsen, A., Jørgensen, F., Kaşıkara, H., Kanayama, H., Kanerva, J., Katz, B., Kenney, J., Kotsyba, N., Krek, S., Laippala, V., Lam, L., Lê H`ông, P., Lenci, A., Ljubešić, N., Lya-shhevskaya, O., Lynn, T., Makazhanov, A., Manning,

- C., Mărănduc, C., Mareček, D., Martínez Alonso, H., Martins, A., Mašek, J., Matsumoto, Y., McDonald, R., Missilä, A., Mititelu, V., Miyao, Y., Montemagni, S., Mori, K. S., Mori, S., Moskalevskiy, B., Muischnek, K., Mustafina, N., Müürisep, K., Nguyễn Thị, L., Nguyễn Thị Minh, H., Nikolaev, V., Nurmi, H., Osenova, P., Östling, R., Øvrelid, L., Paiva, V., Pascual, E., Passarotti, M., Perez, C.-A., Petrov, S., Piitulainen, J., Plank, B., Popel, M., Pretkalniņa, L., Prokopidis, P., Puolakainen, T., Pyysalo, S., Rademaker, A., Ramasamy, L., Real, L., Rituma, L., Rosa, R., Saleh, S., Saulite, B., Schuster, S., Seeker, W., Seraji, M., Shakurova, L., Shen, M., Silveira, N., Simi, M., Simionescu, R., Simkó, K., Šimková, M., Simov, K., Smith, A., Spadine, C., Suhr, A., Sulubacak, U., Szántó, Z., Tanaka, T., Tsarfaty, R., Tyers, F., Uematsu, S., Uria, L., van Noord, G., Varga, V., Vincze, V., Wallin, L., Wang, J. X., Washington, J. N., Wirén, M., Žabokrtský, Z., Zeldes, A., Zeman, D., and Zhu, H. (2016). Universal dependencies 1.4. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- of Chile, A. U. (2021). Patagón supercomputer.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Pérez, J. M., Furman, D. A., Alemany, L. A., and Luque, F. (2021). Robertuito: a pre-trained language model for social media text in spanish. *arXiv preprint arXiv:2111.09453*.
- Pires, T., Schlinger, E., and Garrette, D. (2019a). How multilingual is multilingual bert? *arXiv preprint arXiv:1906.01502*.
- Pires, T., Schlinger, E., and Garrette, D. (2019b). How multilingual is multilingual bert? *CoRR*, abs/1906.01502.
- Rajpurkar, P., Zhang, J., Lopyrev, K., and Liang, P. (2016). Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*.
- Rosenthal, S., Farra, N., and Nakov, P. (2017). Semeval-2017 task 4: Sentiment analysis in twitter. In *Proceedings of the 11th international workshop on semantic evaluation (SemEval-2017)*, pages 502–518.
- Sang, E. F. (2002). Introduction to the conll-2002 shared task: Language-independent named entity recognition. *arXiv preprint cs/0209010*.
- Sanh, V., Debut, L., Chaumond, J., and Wolf, T. (2020). Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter.
- Schwenk, H. and Li, X. (2018). A corpus for multilingual document classification in eight languages. In Nicoletta Calzolari (Conference chair), et al., editors, *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Paris, France, may. European Language Resources Association (ELRA).
- Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A. Y., and Potts, C. (2013). Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Sun, S., Cheng, Y., Gan, Z., and Liu, J. (2019). Patient knowledge distillation for BERT model compression. 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 4323–4332, Hong Kong, China, November. Association for Computational Linguistics.
- Sun, Z., Yu, H., Song, X., Liu, R., Yang, Y., and Zhou, D. (2020). MobileBERT: a compact task-agnostic BERT for resource-limited devices. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2158–2170, Online, July. Association for Computational Linguistics.
- Tiedemann, J. (2012). Parallel data, tools and interfaces in opus. In *Lrec*, volume 2012, pages 2214–2218. Citeseer.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Virtanen, A., Kanerva, J., Ilo, R., Luoma, J., Luotolahti, J., Salakoski, T., Ginter, F., and Pyysalo, S. (2019). Multilingual is not enough: Bert for finnish. *arXiv preprint arXiv:1912.07076*.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. (2018). Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Warstadt, A., Singh, A., and Bowman, S. R. (2019). Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Williams, A., Nangia, N., and Bowman, S. R. (2017). A broad-coverage challenge corpus for sentence understanding through inference. *arXiv preprint arXiv:1704.05426*.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., et al. (2019). Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Wu, S. and Dredze, M. (2019). Beto, bentz, becas: The surprising cross-lingual effectiveness of bert. *arXiv preprint arXiv:1904.09077*.
- Xue, L., Constant, N., Roberts, A., Kale, M., Al-Rfou,

- R., Siddhant, A., Barua, A., and Raffel, C. (2020). mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934*.
- Yang, Y., Zhang, Y., Tar, C., and Baldrige, J. (2019). Paws-x: A cross-lingual adversarial dataset for paraphrase identification. *arXiv preprint arXiv:1908.11828*.
- You, Y., Li, J., Reddi, S., Hseu, J., Kumar, S., Bhojanapalli, S., Song, X., Demmel, J., Keutzer, K., and Hsieh, C.-J. (2019). Large batch optimization for deep learning: Training bert in 76 minutes. *arXiv preprint arXiv:1904.00962*.
- Zhang, Y., Baldrige, J., and He, L. (2019). Paws: Paraphrase adversaries from word scrambling. *arXiv preprint arXiv:1904.01130*.