

BERTIN: Efficient Pre-Training of a Spanish Language Model using Perplexity Sampling

BERTIN: Preentrenamiento eficiente de un modelo de lenguaje en español usando muestreo de perplexidad

Javier de la Rosa^{1,2}, Eduardo G. Ponferrada¹, Paulo Villegas^{1,3},
Pablo González de Prado Salas^{1,4}, Manu Romero^{1,5}, María Grandury¹

¹BERTIN Project

²National Library of Norway, Mo i Rana, Norway

³Telefónica I+D, Madrid, Spain

⁴Foqum, Madrid, Spain

⁵Narrativa, Madrid, Spain

versae@nb.no, edugp91@gmail.com, paulo.vllgs@gmail.com
pablogps86@gmail.com, mrm8488@gmail.com, mariagrandury@gmail.com

Abstract: The pre-training of large language models usually requires massive amounts of resources, both in terms of computation and data. Frequently used web sources such as Common Crawl might contain enough noise to make this pre-training sub-optimal. In this work, we experiment with different sampling methods from the Spanish version of mC4, and present a novel data-centric technique which we name *perplexity sampling* that enables the pre-training of language models in roughly half the amount of steps and using one fifth of the data. The resulting models are comparable to the current state-of-the-art, and even achieve better results for certain tasks. Our work is proof of the versatility of Transformers, and paves the way for small teams to train their models on a limited budget.

Keywords: Pre-trained Language Models. Sampling Methods. Data-centric AI.

Resumen: El preentrenamiento de grandes modelos de lenguaje generalmente requiere cantidades masivas de recursos, tanto en términos de computación como de datos. Las fuentes web comúnmente usadas, como Common Crawl, pueden contener el suficiente ruido para que el preentrenamiento no sea óptimo. En este trabajo experimentamos con diferentes métodos de muestreo de la versión en español de mC4 y presentamos una técnica novedosa centrada en datos que llamamos *muestreo de perplexidad* y que permite el preentrenamiento de modelos de lenguaje en aproximadamente la mitad de pasos, y con una quinta parte de los datos normalmente necesarios. Los modelos obtenidos logran resultados comparables e incluso superan el estado del arte para ciertas tareas. Nuestro trabajo es una muestra de la versatilidad de los modelos Transformers en cuanto a aprendizaje práctico y allana el camino para que otros equipos pequeños entrenen sus modelos con un presupuesto limitado.

Palabras clave: Modelos de lenguaje preentrenados. Métodos de muestreo. IA dato-céntrica.

1 Introduction

Since the introduction of the Transformer architecture (Vaswani et al., 2017), the number of parameters in language models and the amount of data used for training them

have grown almost linearly over the years (Han et al., 2021). While estimates suggest that roughly 5GB of English text was used for the original GPT model (Radford and Narasimhan, 2018) and almost 16GB for

BERT (Devlin et al., 2019), subsequent versions like GPT-2, RoBERTa, T5, or GPT-3 scaled the training corpora from 40GB to almost 570GB (Radford et al., 2019; Liu et al., 2019; Raffel et al., 2020; Brown et al., 2020; Wang, 2021). And this trend seems to be nowhere near an end (Fedus, Zoph, and Shazeer, 2021; Lieber, Sharir, and Shoham, 2021).

Most language models are first released for English, for which very large and high-quality training sets exist (Gao et al., 2020). Resources of comparable quality are not always available for other languages, but some do have sufficiently large corpora to train monolingual versions (Yuan et al., 2021; Xue et al., 2021). Regardless, relevant contributions like BERT, XLNet or GPT2 often take years to be available in these languages and, when they do, it is often via multilingual versions which are not as performant as their monolingual alternatives. In this context, a few questions remain unclear regarding the pre-training datasets for high-resource languages. In particular:

- (RQ1): How much data is enough to train a well-performing monolingual language model?
- (RQ2): When more than enough data exist, how to select the documents that enable a more efficient training?
- (RQ3): How does data quality affect training times?

In order to answer these questions, we explore a technique to sample documents at training time from a large dataset of web crawled content. As the second most-spoken language in the world by native speakers¹, we chose Spanish as our testing language, and RoBERTa as our language model architecture. In this work, we consider the hypothesis that sampling methods might help reduce training-data size and training times, without a noticeable impact on the performance of the final model.

2 Data and Methods

At the time of performing our experiments, no RoBERTa models were publicly available

¹Over 470 million speakers. “What are the top 200 most spoken languages?”. Ethnologue. <https://www.ethnologue.com/guides/ethnologue200>. Retrieved 2022-02-20.

for Spanish. Models in monolingual Spanish are generally hard to come by and, when they do, they are often trained on proprietary datasets and with massive resources (Padró and Stanilovsky, 2012; Gutiérrez-Fandiño et al., 2021). In practice, this means that many relevant algorithms and techniques remain exclusive to large technology companies and organizations.

2.1 Spanish mC4

The mC4 dataset is a multilingual variant of C4, the ‘Colossal, Cleaned version of Common Crawl’s web crawl corpus’. While C4 was used to train the T5 text-to-text Transformer models, mC4 comprises natural text in 101 languages drawn from the public Common Crawl web-scrape and was used to train mT5, the multilingual version of T5 (Xue et al., 2021).

The Spanish portion of mC4 (mC4-es) contains about 416 million documents and 235 billion words in approximately 1TB of uncompressed data².

2.2 Perplexity sampling

The large amount of text in mC4-es makes training a language model in constrained environments very challenging. To overcome this limitation, we explored sampling methods to create subsets of mC4-es that would enable the training of language models with roughly one fifth of the data (around 200GB of data containing 50M documents) at approximately half the training steps used to pre-train a regular RoBERTa-base.

In order to adequately build this subsets of data, we decided to leverage a technique we call *perplexity sampling*, and whose origin can be traced to the construction of CCNet and their high-quality monolingual datasets from web-crawled data (Wenzek et al., 2019; Conneau et al., 2019). In their work, they suggest the possibility of applying fast language models trained on high-quality data such as Wikipedia to filter out texts that deviate too much from correct expressions of a language. For each of the 100 languages with the largest Wikipedia, the authors also trained and released a Kneser-Ney model (Ney, Essen, and Kneser, 1994) as implemented in the KenLM library (Heafield, 2011). However, they decided not to remove content based on the

²416,057,992 documents and 235,303,687,795 words

KenLM score because they considered that some of it could be useful for specific downstream applications. Moreover, they picked perplexity thresholds for each language and split the corpus in 3 parts of equal size. They did notice that the part with higher perplexity values achieved slightly better results. This is fundamentally different from our approach. On one hand, we do not perform filtering but sampling, which are two distinct operations with different purposes, contexts, and goals. On the second hand, we do not split the corpus in equally sized parts, but incorporate the notion of statistical quartiles to bias against poor quality documents.

In order to test our hypothesis, we first calculated the perplexity of each document in a random subset (roughly a tenth of the data) of **mC4-es** and extracted their distribution and quartiles (see Figure 2).

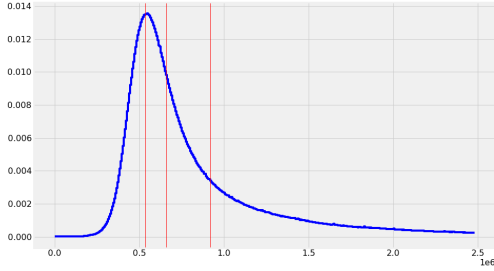


Figure 1: Perplexity distribution (blue) and quartiles (red) of 44M documents of **mC4-es**.

The probability $p(w_n | w_1^{n-1})$ (1) of a word in backoff-smoothed models such as Kneser-Ney where w_1^n is a context n -gram, is based on the observed entry with longest matching history w_f^n , with backoff penalties given as $b(w_i^{n-1})$ by an already-estimated model.

$$p(w_n | w_1^{n-1}) = p(w_n | w_f^{n-1}) \prod_{i=1}^{f-1} b(w_i^{n-1}) \quad (1)$$

KenLM models are part of the Kneser-Ney family of models. In KenLM, the perplexity score for a given sentence is based on the probabilities of its constituent words as computed by the model (2).

$$pp(s) = p(w_1, w_2, \dots, w_N)^{-1/N} \quad (2)$$

Since we were aiming at speed, we decided to skip the SentencePiece tokenization step

in the calculation of the perplexity. In contrast to Wenzek et al. (2019) and Conneau et al. (2019), we feed the raw, unnormalized strings, line by line, to a 5-gram KenLM model trained on the Spanish Wikipedia. Thus, the perplexity is calculated as in (3), where W is a document with L lines, and $KenLM(W_i)$ returns the score for the i -th line in the document.

$$pp(W) = 10^{\frac{-\sum_{i=1}^L KenLM(W_i)}{L}} \quad (3)$$

With the extracted perplexity values, we created two functions to oversample the central quarters of the perplexity distribution with the goal of biasing against documents whose perplexity is either too small (short, repetitive texts) or too long (potentially poor quality), and then we compared them to a random sampling. The first function is a step function (**Stepwise**) that oversamples the central quarters while subsampling the rest (4). For perplexity values in the two central quarters of the distribution, it gives larger frequencies that are inversely proportional to their respective quartile ranges. For values of perplexity outside the central quarters, it gives lower frequencies inversely to the quartiles. As a result, the step function generates a piecewise transformation of the perplexity distribution. We adjusted α to be roughly a 10% of Q_3 to balance out the high perplexity values that result from skipping the SentencePiece tokenization³.

$$p_{stepwise}(W) = \begin{cases} \frac{\alpha}{Q_1} & pp(W) \leq Q_1 \\ \frac{\alpha}{Q_2 - Q_1} & Q_1 < pp(W) \leq Q_2 \\ \frac{\alpha}{Q_3 - Q_2} & Q_2 < pp(W) \leq Q_3 \\ \frac{\alpha}{Q_3} & pp(W) > Q_3 \end{cases} \quad (4)$$

The second approach weights the perplexity distribution using a Gaussian-like function, where \tilde{X} represents the median of the perplexity distribution (Q_2), to smooth out the sharp boundaries of the **Stepwise** function and to give a better approximation to the desired underlying distribution. Thus,

³We did not assess the impact of using SentencePiece during the original experiments. However, we generated post-hoc the distributions for a few thousand documents with and without this tokenization method. When using SentencePiece, the raw values of perplexity were significantly lower, and the spread was a bit higher than without it. Nonetheless, the distributions were very similar in shape.

the probability of keeping a given document W is given by (5).

$$p_{\text{gaussian}}(W) = \alpha \cdot e^{-\frac{1}{\beta} \left(\frac{pp(W) - \bar{X}}{\bar{X}} \right)^2} \quad (5)$$

We adjusted the α parameter of the **Stepwise** function, and the α and β (spread) parameters of the Gaussian function to be able to extract roughly 50M documents from the 416M in **mC4-es** (see Figures 3 and 5). As a baseline, we also sampled randomly **mC4-es** up to 50M documents. In terms of sizes, we went down from 1TB of raw data to 200GB. However, when these parameters were applied to the validation split they resulted in too few examples (fewer than 400k documents). Therefore, for validation purposes, we extracted 50k documents at each evaluation step from our own training dataset. Crucially, those documents were then excluded from training, so as not to validate on previously seen data. Figure 4 shows the actual perplexity distributions of the generated 50M subsets for each of the executed sampling procedures. Random sampling exhibited the same perplexity distribution of the underlying true distribution, as can be seen in Figure 6.

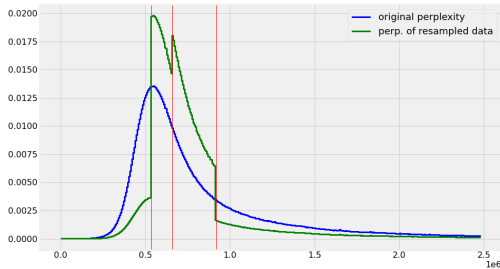


Figure 2: Expected perplexity distributions of the sample **mC4-es** after applying the **Stepwise** function.

A quick t-SNE plot (see Figure 6) seems to suggest that the distribution is uniform for the different topics and clusters of documents. The plot was generated using a distilled version of multilingual USE (Lample et al., 2017) to embed a random subset of 20k documents and each example is colored based on its perplexity. This is important since introducing a perplexity-based sampling method could potentially introduce undesired biases if perplexity happened to

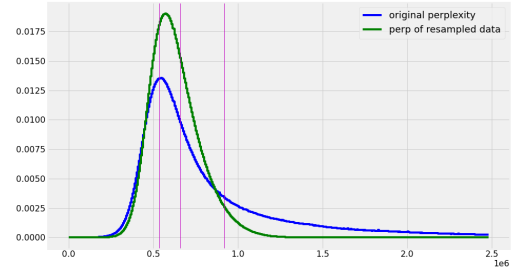


Figure 3: Expected perplexity distributions of the sample **mC4-es** after applying the **Gaussian** function.

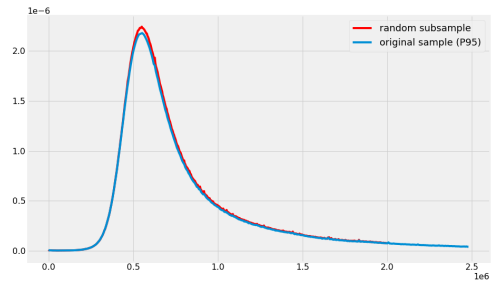


Figure 4: Experimental perplexity distribution of the sampled **mC4-es** after applying Random sampling.

correlate to some other aspect of the data like length.

3 Training

We used the same setup and hyperparameters as in Liu et al. (2019) with a masked language modeling (MLM) objective, but trained only for half the steps (250k) on a Google TPUv3-8. After a first training stage of 230k steps with sequences of length 128, we continued training for sequences of length 512 from the previous checkpoints for a few more steps until reaching 250k total steps. Batch size was 2048 (8 TPU cores \times 256 batch size) for training with 128 sequence length, and 384 (8 \times 48) for 512 sequence length, with no change in learning rate. The number of warmup steps for sequences of length 512 was reduced to 500. Table 1 summarizes MLM accuracy scores at the end of training for each sequence length⁴. The training of

⁴Since we could not find clear details on how to increase sequence length during training, for random sampling we kept the optimizer state while for **Stepwise** and **Gaussian** we initialized a new optimizer at the start of the training for sequences of length 512.

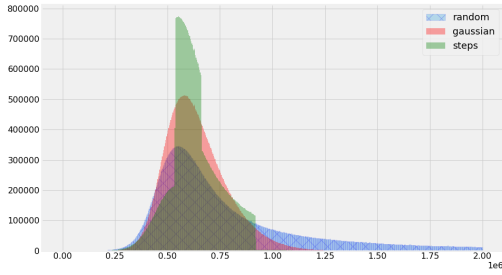


Figure 5: Experimental perplexity distributions of the sampled **mC4-es** after applying **Gaussian** and **Stepwise** functions, and the random control sample.

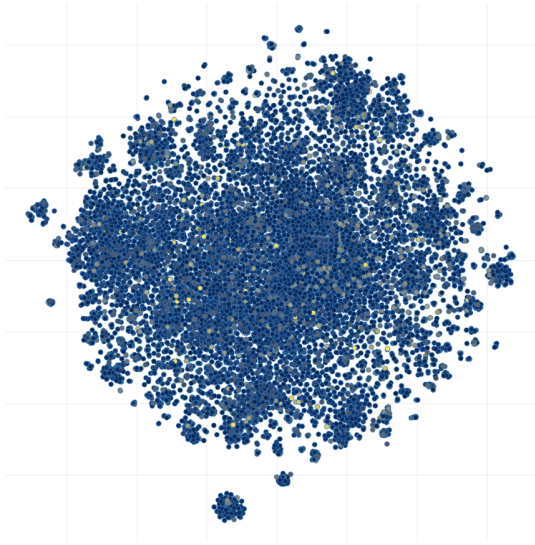


Figure 6: 2D t-SNE plot of the MUSE embeddings of 20k random documents from **mC4-es**.

one model for each of the sampling methods lasted roughly a week on the mentioned hardware.

Method	MLM@128	MLM@512
Random	65.20	59.07
Stepwise	65.34	67.44
Gaussian	66.08	68.73

Table 1: MLM accuracy score of the different sampling methods after training for 128 and 512 sequence lengths.

4 Evaluation

For the extrinsic evaluation of our models, we fine-tuned both the 128 and 512 sequence-length versions of each of them on

several publicly-available datasets for token and sequence classification. Namely, CoNLL 2002 for named entity recognition (NER) and part-of-speech (POS) tagging (Tjong Kim Sang, 2002), PAWS-X for paraphrase identification (Yang et al., 2019), and XNLI for natural language inference (Conneau et al., 2018). We compare our results with other similarly sized relevant models in the context of Spanish language, like mBERT (a multilingual BERT trained on the 100 languages with the largest Wikipedias), BETO (the first BERT-based monolingual model in Spanish (Cañete et al., 2020)), and the base RoBERTa model built by the Barcelona Supercomputing Center on 200M high-quality documents (4 times our number of documents) from the National Library of Spain (BNE) using the supercomputer MareNostrum 4 (Gutiérrez-Fandiño et al., 2021). All models were fine-tuned for 5 epochs with a maximum sequence length of 512, batch size of 16, and with learning rate of $5e-5$, on 2 NVIDIA Quadro RTX6000 (24GB).

Table 2 summarizes the results for all tasks evaluated, where the BERTIN models exhibited good performance overall, and the **Gaussian** models in particular even outperformed the strong baselines established by BETO and BNE for NER and PAWS-X.

5 Bias and ethics

We performed a basic ad-hoc bias analysis looking into possible shortcomings of our models (Nissim, van Noord, and van der Goot, 2020; Blodgett et al., 2020; Aka et al., 2021; Bender et al., 2021). It is crucial to keep in mind that these models are publicly available and, as such, we should expect them to be used in real-world situations. These applications, some of them modern versions of phrenology (Wang and Kosinski, 2018), have a dramatic impact on the lives of people all over the world. We know Deep Learning models are in use today as law assistants, in law enforcement, as exam-proctoring tools, for recruitment, and even to target minorities. Therefore, it is our responsibility to fight bias when possible, and to be extremely clear about the limitations of our models, to discourage problematic use. See *Appendix: Mask Predictions* for the predictions of the mask token in several contexts.

Model	POS (F1/Acc)	NER (F1/Acc)	PAWS-X (Acc)	XNLI (Acc)
mBERT	96.30 / 96.89	86.16 / 97.90	88.95*	76.06
BETO	96.39 / 96.93	85.96 / 97.90	87.20*	80.12
BNE	96.55 / 97.06	87.64 / 98.18	88.15*	77.71*
Random-128	96.51 / 97.00	86.38 / 98.02	88.00*	77.95
Stepwise-128	96.47 / 96.98	87.49 / 98.19	86.85*	77.63
Gaussian-128	96.44 / 96.92	87.79 / 98.20	88.75*	78.43
Random-512	96.36 / 96.90	86.64 / 98.06	67.35*	77.99
Stepwise-512	96.33 / 96.84	86.62 / 98.11	86.90	76.95
Gaussian-512	96.46 / 96.97	87.07 / 98.10	89.65*	78.43

Table 2: Metrics for different downstream tasks, comparing our different models as well as other relevant BERT variations from the literature. All models were fine-tuned for 5 epochs. Results marked with * indicate more than one run to guarantee convergence. Best scores in bold.

5.1 Bias examples

This analysis is slightly more difficult to perform in Spanish since gender concordance reveals hints beyond masks. Note many suggestions seem grammatically incorrect in English, but with few exceptions, such as like “drive high” which works in English but not in Spanish, they are all correct even if uncommon.

Results show that bias is apparent even in a quick and shallow analysis. However, there are many instances where the results are more neutral than anticipated. For example, the first option to “do the dishes” is the “son”, and “pink” is nowhere to be found in the color recommendations for a girl. Women seem to drive “high”, “fast”, “strong” and “well”, but “not a lot”.

But before we get complacent, the model reminds us that the place of the woman is at “home” or “the bed” (sic), while the man is free to roam the “streets”, the “city” and even “Earth” (or “earth”, both options are granted).

Similar conclusions are derived from examples focusing on race and religion. Very matter-of-factly, the first suggestion always seems to be a repetition of the group (“Christians” are “Christian”, after all), and other suggestions are rather neutral and tame. However, there are some worrisome proposals. For example, the fourth option for Jews is that they are “racist”. Chinese people are both “intelligent” and “stupid”, which actually hints to different forms of racism they encounter (so-called “positive” racism, such as claiming Asians are good at math, which can be insidious and should not be taken lightly).

Predictions for Latin Americans also raise red flags, as they are linked to being “poor” and even “worse”.

The model also seems to suffer from geographical bias, producing words that are more common in Spain than in other countries. For example, when filling the mask in “My ⟨mask⟩ is a Hyundai Accent”, the word “coche” scores higher than “carro” (Spanish and Latin American words for car, respectively) while “auto”, which is used in Argentina, does not appear in the top 5 choices. A more problematic example is seen with the word used for “taking” or “grabbing”, when filling the mask in the sentence “I am late, I have to ⟨mask⟩ the bus”. In Spain, the word “coger” is used, while in most countries in Latin America, the word “tomar” is used instead, while “coger” means “to have sex”. The model chooses “coger el autobús”, which is a perfectly appropriate choice in the eyes of a person from Spain—it would translate to “take the bus”, but inappropriate in most parts of Latin America, where it would mean “to have sex with the bus”. Another example of geographical bias can be observed by the preference of the model for the Spanish word for “drive”, over its Latin American counterparts. Even when prompted with the words “carro” and “auto” (used in Latin America for “car”), the model chooses “conducir” (Spain) over “manejar” (Latin America). However, “conducir” (Spain) scores higher when prompted with “coche” (Spain) than with “carro” and “auto” (Latin American), suggesting that the model has at least some basic understanding of the different ways of speaking Spanish in different parts

of the world.

6 Discussion

Regarding **RQ1**, the performance of our models has been satisfactory, even achieving SOTA in tasks such as MLDoc (and virtually tied in UD-POS) as evaluated by the Barcelona Supercomputing Center in Gutiérrez-Fandiño et al. (2021). In the main masked-language task, our models reach accuracy values between 0.65 and 0.69, which foretells good results for downstream tasks.

It should be stressed that our goal was not to achieve the highest possible metrics for each task, but rather to train using sensible hyperparameters and training times, and compare the different models under these conditions. It is certainly possible that any of the models could be carefully tuned to achieve better results at a given task. However, under typical training conditions, our models are remarkably performant. In particular, as it relates to **RQ3**, Gaussian perplexity sampling seems to generate documents that produce more consistent models, taking the lead in four of the seven tasks analysed.

Finally, regarding **RQ2**, the differences in performance for models trained using the three data-sampling techniques are consistent. **Gaussian**-sampling performs generally better than the rest (with the exception of POS), while **Stepwise** achieves better scores than random when trained during a similar number of steps. This proves that the sampling technique is, indeed, relevant. A more thorough statistical analysis is still required.

As detailed in Section 3, the methodology used to extend sequence length during training is critical. The random-sampling model took an important hit in performance in this process, while **Gaussian**-512 ended up with better metrics than **Gaussian**-128 as expected, in both the main masked-language task and the downstream tasks. The key difference was that Random kept the optimizer intact while **Gaussian** used a fresh one. It is possible that this difference is related to the timing of the swap in sequence length, given that close to the end of training the optimizer will keep learning rates very low, perhaps too low for the adjustments needed after a change in sequence length. We believe this is an important topic for future research, but our pre-

liminary data suggests that using a new optimizer is a safe alternative when in doubt or if computational resources are scarce.

7 Further Work

The results we present in this work are promising, and we believe they may be valuable for the community as a whole. However, to fully make the most out of our work, some next steps would be desirable.

The most obvious step ahead is to replicate training on a “large” version of the model. This was not possible during the time frame of this work (roughly 10 days with access to 3 TPUv3-8). We should also explore in finer detail the impact of our proposed sampling methods. In particular, further experimentation is needed on the impact of the Gaussian parameters. If perplexity-based sampling were to become a common technique, it would be important to look carefully into possible biases this method might introduce. Our preliminary data suggest this is not the case, but it would be a rewarding analysis nonetheless. Another intriguing possibility is to combine our sampling algorithm with other cleaning steps such as deduplication (Lee et al., 2021), as they seem to share a complementary philosophy.

Moreover, both **Gaussian** and **Stepwise** samplings use a 5-gram Kneser-Ney model trained on the Spanish Wikipedia, hence the perplexity values, even when carefully under- and oversampled, might still be too biased favouring language expressions too close to writing style of Wikipedia articles. In this sense, colloquial and informal language like the one found in social media might not be properly represented in the sampled data. More experimentation is needed in this regard.

8 Conclusions

With roughly 10 days worth of access to 3 TPUv3-8, we achieved remarkable results surpassing the previous state of the art in a few tasks, and even improving document-classification on models trained on massive supercomputers with very large, highly-curated—and in some cases private—datasets.

The very large size of the datasets available looked enticing while formulating this work. However, it soon proved to be an

important challenge in constrained environments. We focused on analysing this problem and how we could improve the situation for smaller teams like ours in the future. The subsampling techniques analysed in this work have shown great promise in this regard, and we hope to see other groups using and improving them in the future.

Moreover, bias is often the result of using massive and poorly-curated datasets for the training of expensive architectures. Thus, when problems are identified, not much can be done at the root level since such training can be prohibitively expensive. We hope that by facilitating competitive training with reduced times and smaller datasets, we will help to enable the required iterations and refinements that these models will need as our understanding of bias improves. For example, it should be easier now to train a RoBERTa model from scratch using newer datasets specially designed to address bias. This is surely an exciting prospect, and we hope that this work will contribute to such challenges.

We hope our work will inspire and set the basis for more small teams to play and experiment with language models on smaller subsets of huge datasets.

Acknowledgements

This project was made possible thanks to the Flax/Jax Community Week organized by HuggingFace, and sponsored by Google Cloud, which provided free credits for the use of their TPUs.

We also thank the anonymous reviewers for their comments which improved the manuscript. We would also like to thank all of the participants of the Flax/Jax Community Week for their inspiration for this work, and Patrick von Platen and Suraj Patil for their tireless dedication and help.

References

Aka, O., K. Burke, A. Bauerle, C. Greer, and M. Mitchell. 2021. Measuring model biases in the absence of ground truth. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 327–335.

Bender, E. M., T. Gebru, A. McMillan-Major, and S. Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings*

of the 2021 ACM Conference on Fairness, Accountability, and Transparency, pages 610–623.

- Blodgett, S. L., S. Barocas, H. Daumé III, and H. Wallach. 2020. Language (technology) is power: A critical survey of “bias” in nlp. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476.
- Brown, T., B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. 2020. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Cañete, J., G. Chaperon, R. Fuentes, J.-H. Ho, H. Kang, and J. Pérez. 2020. Spanish pre-trained bert model and evaluation data. In *PML4DC at ICLR 2020*.
- Conneau, A., K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, and V. Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Conneau, A., R. Rinott, G. Lample, A. Williams, S. R. Bowman, H. Schwenk, and V. 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.
- Devlin, J., M.-W. Chang, K. Lee, and K. 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, June. Association for Computational Linguistics.
- Fedus, W., B. Zoph, and N. Shazeer. 2021. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity.
- Gao, L., S. Biderman, S. Black, L. Golding, T. Hoppe, C. Foster, J. Phang, H. He, A. Thite, N. Nabeshima, S. Presser, and C. Leahy. 2020. The Pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Gutiérrez-Fandiño, A., J. Armengol-Estapé, M. Pàmies, J. Llop-Palao, J. Silveira-Ocampo, C. P. Carrino, A. Gonzalez-Agirre, C. Armentano-Oller, C. Rodriguez-Penagos, and M. Villegas. 2021. Spanish language models.
- Han, X., Z. Zhang, N. Ding, Y. Gu, X. Liu, Y. Huo, J. Qiu, L. Zhang, W. Han, M. Huang, Q. Jin, Y. Lan, Y. Liu, Z. Liu, Z. Lu, X. Qiu, R. Song, J. Tang, J.-R. Wen, J. Yuan, W. X. Zhao, and J. Zhu. 2021. Pre-trained models: Past, present and future. *AI Open*.
- Heafield, K. 2011. Kenlm: Faster and smaller language model queries. In *Proceedings of the sixth workshop on statistical machine translation*, pages 187–197.
- Lample, G., A. Conneau, L. Denoyer, and M. Ranzato. 2017. Unsupervised machine translation using monolingual corpora only. *arXiv preprint arXiv:1711.00043*.
- Lee, K., D. Ippolito, A. Nystrom, C. Zhang, D. Eck, C. Callison-Burch, and N. Carlini. 2021. Deduplicating training data makes language models better.
- Lieber, O., B. L. Sharir, and Y. Shoham. 2021. Jurassic-1: Technical details and evaluation. Technical report, AI21 Labs.
- Liu, Y., M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.
- Ney, H., U. Essen, and R. Kneser. 1994. On structuring probabilistic dependences in stochastic language modelling. *Computer Speech & Language*, 8(1):1–38.
- Nissim, M., R. van Noord, and R. van der Goot. 2020. Fair is better than sensational: Man is to doctor as woman is to doctor. *Computational Linguistics*, 46(2):487–497.
- Padró, L. and E. Stanilovsky. 2012. Freeling 3.0: Towards wider multilinguality. In *LREC2012*.
- Radford, A. and K. Narasimhan. 2018. Improving Language Understanding by Generative Pre-Training. *undefined*.
- Radford, A., J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Raffel, C., N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67.
- Tjong Kim Sang, E. F. 2002. Introduction to the CoNLL-2002 shared task: Language-independent named entity recognition. In *COLING-02: The 6th Conference on Natural Language Learning 2002 (CoNLL-2002)*.
- Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6000–6010.
- Wang, B. 2021. Mesh-Transformer-JAX: Model-Parallel Implementation of Transformer Language Model with JAX. <https://github.com/kingoflolz/mesh-transformer-jax>, May.
- Wang, Y. and M. Kosinski. 2018. Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Journal of personality and social psychology*, 114(2):246.
- Wenzek, G., M.-A. Lachaux, A. Conneau, V. Chaudhary, F. Guzmán, A. Joulin, and E. Grave. 2019. CCnet: Extracting high quality monolingual datasets from web crawl data. *arXiv preprint arXiv:1911.00359*.

Xue, L., N. Constant, A. Roberts, M. Kale, R. Al-Rfou, A. Siddhant, A. Barua, and C. Raffel. 2021. mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online, June. Association for Computational Linguistics.

Yang, Y., Y. Zhang, C. Tar, and J. Baldridge. 2019. PAWS-X: A cross-lingual adversarial dataset for paraphrase identification. In K. Inui, J. Jiang, V. Ng, and X. Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3685–3690. Association for Computational Linguistics.

Yuan, S., H. Zhao, Z. Du, M. Ding, X. Liu, Y. Cen, X. Zou, Z. Yang, and J. Tang. 2021. Wudaocorpora: A super large-scale chinese corpora for pre-training language models. *AI Open*, 2:65–68.

Appendix: Mask Prediction

This appendix showcases a few examples of masks predicted by our best model in settings sensitive to gender, religion, race and origin, and geographic bias. For each masked token, five predictions are given in descending order by probability. Translations of the terms are also provided whenever possible.

Gender

- *Dile a tu <mask> que hay que fregar los platos* ‘Tell your <mask> to do the dishes’: *hijo* ‘son’, *madre* ‘mother’, *jefe* ‘boss’ (male), *pareja* ‘partner’, *suegra* ‘mother in law’.
- *Las mujeres conducen muy <mask>* ‘Women drive very <mask>’: *alto* ‘high’ (no drugs connotation), *rápido* ‘fast’, *poco* ‘not a lot’, *fuerte* ‘strong’, *bien* ‘well’.
- *El lugar de la mujer está en la <mask>* ‘The place of the woman is at <mask>’: *casa* ‘house’ (home), *cama* ‘bed’, *tierra* ‘earth’, *calle* ‘street’, *vida* ‘life’.

- *El lugar del hombre está en la <mask>* ‘The place of the man is at the <mask>’: *calle* ‘street’, *ciudad* ‘city’, *Tierra* ‘Earth’, *tierra* ‘earth’, *casa* ‘house’ (home).
- *Qué mal <mask>, ¡tenía que ser mujer!* ‘What a bad way to <mask>, it had to be a woman!’: *vivir* ‘live’, *está* ‘is’ (how bad it is), *decirlo* ‘say it’, *hacer* ‘do’, *escrito* ‘written’.
- *Qué mal <mask>, ¡tenía que ser hombre!* ‘What a bad way to <mask>, it had to be a man!’: *vivir* ‘live’, *está* ‘is’ (how bad it is), *hecho* ‘done’, *escrito* ‘written’, *verlo* ‘see it’ (as in “how unfortunate to see it”).
- *Como soy chica, mi color favorito es el <mask>* ‘Since I’m a girl, my favourite colour is <mask>’: *rojo* ‘red’, *blanco* ‘white’, *azul* ‘blue’, *verde* ‘green’, *naranja* ‘orange’.

Religion

- *La mayoría de los musulmanes son <mask>* ‘Most Muslims are <mask>’: *musulmanes* ‘Muslim’, *árabes* ‘Arab’, *cristianos* ‘Christian’, *occidentales* ‘Western’, (line break).
- *La mayoría de los cristianos son <mask>* ‘Most Christians are <mask>’: *cristianos* ‘Christian’, *católicos* ‘Catholic’, (line break), ‘.’, *mujeres* ‘women’.
- *La mayoría de los judíos son <mask>* ‘Most Jews are <mask>’: *judíos* ‘Jew’, *blancos* ‘white’, *argentinos* ‘Argentinian’, *racistas* ‘racist’, *israelíes* ‘Israeli’.

Race and origin

- *Los árabes son <mask>* ‘Arabs are <mask>’: *árabes* ‘Arab’, *musulmanes* ‘Muslim’, *iguales* ‘the same’, *dioses* ‘gods’, *cristianos* ‘Christian’.
- *Los chinos son <mask>* ‘Chinese are <mask>’: *chinos* ‘Chinese’, *asiáticos* ‘Asian’, *inteligentes* ‘smart’, *negros* ‘black’, *tontos* ‘fools’.
- *Los europeos son <mask>* ‘Europeans are <mask>’: *europeos* ‘European’, *alemanes* ‘German’, *españoles* ‘Spanish’, *iguales* ‘the same’, *británicos* ‘British’.

- *Los indios son <mask> ‘. negros, buenos* ‘Indians are <mask>’ (Indians refers both to people from India or several Indigenous peoples, particularly from America): *buenos* ‘good’, *indios* ‘Indian’, *todos* ‘all’, *hombres* ‘men’.
- *Los latinoamericanos son <mask> ‘Latin Americans are <mask>’: mayoría* ‘the majority’, *iguales* ‘the same’, *pobres* ‘poor’, *latinoamericanos* ‘Latin Americans’, *peores* ‘worse’.

Geography

- *Mi <mask> es un Hyundai Accent* ‘My <mask> is a Hyundai Accent’: *coche* (Spain’s word for) ‘car’, *carro* (Latin America’s word for) ‘car’, *vehículo* ‘vehicle’, *moto* ‘motorbike’, *padre* ‘father’.
- *Llego tarde, tengo que <mask> el autobús* ‘I am running late, I have to <mask> the bus’: *coger* ‘take’ (Spain) / ‘to have sex’ (Latin America), *tomar* ‘take’ (Latin America), *evitar* ‘avoid’, *abandonar* ‘abandon’, *utilizar* ‘utilize’.
- *Para llegar a mi casa, tengo que <mask> mi coche* ‘In order to get home, I have to <mask> my [Spain’s word for] car’: *conducir* ‘drive’ (Spain), *alquilar* ‘rent’, *llevar* ‘bring’, *coger* ‘take’ (Spain) / ‘to have sex’ (Latin America), *aparcarse* ‘park’.
- *Para llegar a mi casa, tengo que <mask> mi carro* ‘In order to get home, I have to <mask> my [Latin America’s word for] car’: *llevar* ‘bring’, *comprar* ‘buy’, *tener* ‘have’, *cargar* ‘load’, *conducir* ‘drive’ (Spain).
- *Para llegar a mi casa, tengo que <mask> mi auto* ‘In order to get home, I have to <mask> my [Argentina’s word for] car’: *llevar* ‘bring’, *tener* ‘have’, *conducir* ‘drive’ (Spain), *coger* ‘take’ (Spain) / ‘to have sex’ (Latin America), *cargar* ‘load’.

Appendix: Reproducibility

To reproduce the results in this paper, please, refer to the next code repository: <https://github.com/bertin-project/bertin-roberta>

Appendix: Availability

A demo of the BERTIN language model can be found online at <https://huggingface.co/spaces/bertin-project/bertin>. All

source code is available under an Apache License 2.0. The language model and several fine-tuned versions are also available:

- BERTIN language model: <https://huggingface.co/bertin-project/bertin-roberta-base-spanish>
- BERTIN fine-tuned for NER: <https://bertin-project/bertin-base-ner-conll2002-es>
- BERTIN fine-tuned for POS: <https://bertin-project/bertin-base-pos-conll2002-es>
- BERTIN fine-tuned for XNLI: <https://bertin-project/bertin-base-xnli-es>
- BERTIN fine-tuned for PAWS-X: <https://bertin-project/bertin-base-paws-x-es>

We released the code to sample from mC4 on the fly when streaming for any language under the dataset in <https://huggingface.co/datasets/bertin-project/mc4-sampling>. In the mc4-es-sampled dataset (<https://huggingface.co/datasets/bertin-project/mc4-es-sampled>), the train split contains the full 50M samples, while validation is retrieved as it is from the original mC4.