

BERTiĆ - The Transformer Language Model for Bosnian, Croatian, Montenegrin and Serbian

Nikola Ljubešić

Jožef Stefan Institute

Jamova cesta 39

Ljubljana, Slovenia

nikola.ljubesic@ijs.si

Davor Lauc

Faculty of Humanities and Social Sciences

Ivana Lučića 3

Zagreb, Croatia

davor.lauc@ffzg.hr

Abstract

In this paper we describe a transformer model pre-trained on 8 billion tokens of crawled text from the Croatian, Bosnian, Serbian and Montenegrin web domains. We evaluate the transformer model on the tasks of part-of-speech tagging, named-entity-recognition, geo-location prediction and commonsense causal reasoning, showing improvements on all tasks over state-of-the-art models. For commonsense reasoning evaluation we introduce COPA-HR - a translation of the Choice of Plausible Alternatives (COPA) dataset into Croatian. The BERTiĆ model is made available for free usage and further task-specific fine-tuning through HuggingFace.

1 Introduction

In recent years, pre-trained transformer models have taken the NLP world by storm (Devlin et al., 2018; Liu et al., 2019; Brown et al., 2020), yielding new state-of-the-art results in various tasks and settings. While such models, requiring significant computing power and data quantity, started to emerge for non-English languages (Martin et al., 2019; de Vries et al., 2019), as well as in multilingual flavours (Devlin et al., 2018; Conneau et al., 2019), there is a significant number of languages for which better models can be obtained with the available pre-training techniques.

This paper describes such an effort - training a transformer language model on more than 8 billion tokens of text written in the Bosnian, Croatian, Montenegrin or Serbian language, all these languages being very closely related, mutually intelligible, and classified under the same HBS (Serbo-Croatian) macro-language by the ISO-639-3 standard.¹

¹https://iso639-3.sil.org/code_tables/macrolanguage_mappings/data

The name of the model – BERTiĆ – points at two facts: (1) the language model was trained in Zagreb, Croatia, in whose vernacular diminutives ending in *ić* are frequently used (*fotić* eng. photo camera, *smajlić* eng. smiley, *hengić* eng. hanging together), and (2) in all the countries / languages of this model the patronymic surnames end to a great part with the suffix *ić*, with likely diminutive etymology.

The paper is structured as follows: in the following section we describe the data the model is based on, in the third section we give a short description of the modelling performed, and in the fourth section we present a detailed evaluation of the model.

2 Data

As our data basis we use already existing datasets, namely (1) the hrWaC corpus of the Croatian top-level domain, crawled in 2011 (Ljubešić and Erjavec, 2011) and again in 2014 (Ljubešić and Klubička, 2014), (2) the srWaC corpus of the Serbian top-level domain, crawled in 2014 (Ljubešić and Klubička, 2014), (3) the bsWaC corpus of the Bosnian top-level domain, crawled in 2014 (Ljubešić and Klubička, 2014), (4) the cnrWaC corpus of the Montenegrin top-level domain, crawled in 2019, and (5) the Riznica corpus consisting of Croatian literary works and newspapers (Ćavar and Rončević, 2012).

Given that most of the crawls contain data only up to year 2014, we performed new crawls of the Bosnian, Croatian and Serbian top-level domains. We brand these corpora as CLASSLA web corpora given that CLASSLA is the CLARIN knowledge centre for South Slavic languages² under which we perform most of the described activities. We deduplicate the CLASSLA corpora by removing

²<https://www.clarin.si/info/k-centre/>

dataset	language	# of words
hrWaC	Croatian	1,250,923,836
CLASSLA-hr	Croatian	1,341,494,461
cc100-hr	Croatian	2,880,490,449
Riznica	Croatian	87,724,737
srWaC	Serbian	493,202,149
CLASSLA-sr	Serbian	752,916,260
cc100-sr	Serbian	711,014,370
bsWaC	Bosnian	256,388,597
CLASSLA-bs	Bosnian	534,074,921
cnrWaC	Montenegrin	79,451,738

Table 1: Datasets used for training the BERTiĆ model with their size (in number of words) after deduplication.

identical sentences that were already present in the WaC corpora. The amount of data removed through this deduplication is minor, in all cases in single digit percentages.

We further exploit the recently published cc100 corpora (Conneau et al., 2019) that are based on the CommonCrawl data collection. We perform the same level of deduplication as with the CLASSLA corpora, with every sentence already present in the WaC or CLASSLA corpus being removed from the cc100 corpus. This round of deduplication removed around 15% of the CommonCrawl data.

The resulting sizes of the datasets used for training the BERTiĆ model are presented in Table 1. The overall text collection consists of 8,387,681,518 words.

3 Model training

For training this model we selected the Electra approach to training transformer models (Clark et al., 2020). These models are based on training a smaller generator model and the main, larger, discriminator model whose task is to discriminate whether a specific word is the original word from the text, or a word generated by the generator model. The authors claim that the Electra approach is computationally more efficient than the BERT models (Devlin et al., 2018) based on masked language modelling.

As in BERT and similar transformers models, we constructed a WordPiece vocabulary with a vocabulary size of 32 thousand tokens. A WordPiece model was trained using the HuggingFace tokenizers library³ on the random sample of 10

³<https://huggingface.co/transformers/>

million paragraphs from the whole dataset. Text pre-processing and cleaning differ from the original BERT only in preserving all Unicode characters, while in the original pre-processing diacritics are removed.

Training of the model was performed to the most part with the hyperparameters set for base-sized models (110 million parameters in 12 transformer layers) as defined in the Electra paper (Clark et al., 2020). Training batch size was kept at 1024, the maximum size for the 8 TPUv3 units on which the training was performed. The training was run for 2 million steps (roughly 50 epochs).

4 Evaluation

In this section we present an exhaustive evaluation of the newly trained BERTiĆ model on two token classification tasks – morphosyntactic tagging and named entity recognition, and two sequence classification tasks – geolocation prediction and commonsense causative reasoning.

The reference points in each task are the state-of-the-art transformer models covering the macro-language - multilingual BERT (Devlin et al., 2018) and CroSloEngual BERT (Ulčar and Robnik-Šikonja, 2020). While multilingual BERT (mBERT onwards) was trained on Wikipedia corpora, CroSloEngual BERT (cseBERT onwards) was trained on a similar amount of Croatian data used to train BERTiĆ, but without the data from the remaining languages.

4.1 Morphosyntactic tagging

On the task of morphosyntactic tagging (assigning each word one among multiple hundreds of detailed morphosyntactic classes, e.g. Ncmsay referring to a common masculine noun, in accusative case, singular number, animate) we compare the three transformer models, mBERT, cseBERT and BERTiĆ. We additionally report results, when available, for the current production tagger for the two languages - the CLASSLA tool (Ljubešić and Dobrovoljc, 2019), based on Stanford’s Stanza, exploiting static embedding and BiLSTM technology (Qi et al., 2020).

We perform evaluation of the models on this task on four datasets: the Croatian standard language dataset hr500k (Ljubešić et al., 2018), the Croatian non-standard language dataset ReLDI-hr (Ljubešić et al., 2019a), the Serbian standard language dataset

[main_classes/tokenizer.html](https://huggingface.co/transformers/main_classes/tokenizer.html)

dataset	language	variety	CLASSLA	mBERT	cseBERT	BERTiĆ
hr500k	Croatian	standard	93.87	94.60	95.74	*** 95.81
reldi-hr	Croatian	non-standard	-	88.87	91.63	*** 92.28
SETimes.SR	Serbian	standard	95.00	95.50	96.41	96.31
reldi-sr	Serbian	non-standard	-	91.26	93.54	*** 93.90

Table 2: Average microF1 results on the morphosyntactic annotation task over five training iterations. The highest score per dataset is marked with bold. The statistical significance is tested with the two-sided t-test over the five runs between the two strongest results. Level of significance is labeled with asteriks signs (***) $p \leq 0.001$.

dataset	language	variety	CLASSLA	mBERT	cseBERT	BERTiĆ
hr500k	Croatian	standard	80.13	85.67	88.98	**** 89.21
ReLDI-hr	Croatian	non-standard	-	76.06	81.38	**** 83.05
SETimes.SR	Serbian	standard	84.64	92.41	92.28	92.02
ReLDI-sr	Serbian	non-standard	-	81.29	82.76	*** 87.92

Table 3: Average F1 results on the named entity recognition task over five training iterations. The highest score per dataset is marked with bold. The statistical significance is tested with the two-sided t-test over the five runs between the two strongest results. Level of significance is labeled with asteriks signs (***) $p \leq 0.001$, (****) $p \leq 0.0001$.

SETimes.SR (Batanović et al., 2018) and the Serbian non-standard Twitter language dataset ReLDI-sr (Ljubešić et al., 2019b).

For each dataset and model we perform hyperparameter optimization via Bayesian search on the wandb.ai platform (Biewald, 2020), allowing for 30 iterations. We optimize the initial learning rate (we search between the values of $9e-6$ and $1e-4$) and the epoch number (we search between the values of 3 and 15).

We report average microF1 results of five runs per dataset and model in Table 2. The highest score per dataset is marked with bold. The statistical significance is tested with the two-sided t-test over the five runs between the two highest average results. We can observe that the BERTiĆ model outperforms all the remaining models, cseBERT coming second, on three out of four datasets. Only on the Serbian standard dataset the difference between these two models is insignificant. We argue that this is due to the simplicity of the dataset - it consists of texts from one newspaper only, therefore containing text with little variation even between the training and the testing data.

4.2 Named entity recognition

On the task of named entity recognition we compare the same models on the same datasets as was the case in the previous Section 4.1. We also perform an identical hyperparameter optimisation and experimentation and report the results in Table 3. The results show again that the two best performing

models are cseBERT and BERTiĆ with the latter performing better on three out of four datasets, again, with no significant difference on the standard Serbian task for the same reasons as with the previous task.

4.3 Social media geolocation

In this subsection we compare the three transformer models on the Social Media Geolocation (SMG2020) shared task, which part of the VarDial 2020 Evaluation Campaign (Gaman et al., 2020). The task consists of predicting the exact latitude and longitude of a geo-encoded tweet published in Croatia, Bosnia, Montenegro or Serbia. The shared task winner in 2020 was using the cseBERT model (Scherrer and Ljubešić, 2020) in its approach.

We evaluate the model on the two evaluation metrics of the shared task - median and mean of the distance between gold and predicted geolocations. Given the large size of the training dataset (320,042 instances), we do not perform any additional hyperparameter tuning beyond the one performed during the participation in the shared task and apply the same methodology: we fine-tune the transformer model with batch size of 64 for 40 epochs and retain the model with minimum median distance on development data. The results in Table 4 show that the BERTiĆ model improved the results of the shared task winner – the cseBERT model.

	median	mean
centroid	107.10	145.72
mBERT	42.25	82.05
cseBERT	40.76	81.88
BERTiĆ	37.96	79.30

Table 4: Median distance and mean distance between gold and predicted geolocation (lower is better) on the task of social media geolocation prediction. The best results are marked in bold. No statistical testing was performed due to a large size of the test dataset (39,723 instances).

4.4 Commonsense causal reasoning

The final evaluation round of the new BERTiĆ model is performed on the task of commonsense causal reasoning on a translation of the COPA dataset (Roemmele et al., 2011) into Croatian, the COPA-HR dataset. The translation is performed by following the methodology laid out while preparing the XCOPA dataset (Ponti et al., 2020), a translation of the COPA dataset into 11 typologically balanced languages.

The dataset consists of 400 training, 100 development and 500 examples. Each instance in the dataset consists of a premise (*The man broke his toe*), a question (*What was the cause?*),⁴ and two alternatives, one of them to be chosen by the system as being more plausible (*He got a hole in his sock*, or *He dropped a hammer on his foot*).

While translating the dataset, the translator was also given the task of selecting the more plausible alternative given their translation. The observed agreement between the annotations in the English dataset and the annotations of the Croatian translator was perfect on the training set and the development set, while on the test set one out of 500 choices differed. The problematic example proved to be a rather unclear case – the premise being *I paused to stop talking.*, with the question *What was the cause?*, and the alternatives *I lost my voice.* and *I ran out of breath.*⁵ The dataset is available from the CLARIN.SI repository (Ljubešić, 2021).⁶

The approach taken to benchmarking the three transformer models is that of sentence pair classification, each original instance becoming two sentence pair instances (each sentence pair containing

⁴Roughly half of the instances contain the other question: *What was the effect?*,

⁵The Croatian translator chose the second alternative, while in the original dataset the first alternative is chosen.

⁶<http://hdl.handle.net/11356/1404>

	accuracy
random	50.00
mBERT	54.12
cseBERT	61.80
BERTiĆ	**65.76

Table 5: Average accuracy results on the commonsense causal reasoning task over five training iterations. The highest score per dataset is marked with bold. The statistical significance is tested with the two-sided t-test over the five runs between the two strongest results (** $p < 0.01$).

the premise and one alternative), with different models being trained for *cause* and *effect* questions. During evaluation, separate predictions are made on each of the alternatives, the per-class predictions being fed to a softmax function, and the higher positive-class alternative being chosen as the correct one.

The standard evaluation metric for this dataset is accuracy. Given the balanced nature of the test set, the random baseline is 50%. For hyperparameter optimization the same approach was taken as with the token classification tasks.

The results presented in Table 5 show that both language-specific transformer models outperform mBERT significantly, with BERTiĆ obtaining a significant lead over cseBERT.

5 Conclusion

In this paper we have presented a newly published Electra transformer language model, BERTiĆ, trained on more than 8 billion tokens of previously and newly collected web text written in Bosnian, Croatian, Montenegrin or Serbian. We have applied a very thorough evaluation of the model, comparing it primarily to other state-of-the-art transformer models that support the languages in question. We have obtained significant improvements on all four tasks, with no difference obtained only on one single-source-dataset with little text variation and high training and testing data similarity.

The main conclusions we can draw from our results are the following. (1) Although cseBERT and BERTiĆ use a different approach to building transformer language models, our assumption is that the performance difference between these two lies primarily in the larger amount of data presented to the BERTiĆ model. (2) The improvements on the four tasks with the BERTiĆ model seem to be

smaller on the morphosyntactic tagging task than the remaining three tasks that require more world and commonsense reasoning knowledge. (3) Except for the named entity recognition task on the Serbian non-standard dataset, we fail to observe greater improvements on Serbian tasks than on Croatian ones between `cseBERT` and `BERTiĆ`, regardless the fact that the former has seen none and the latter has seen huge quantities of Serbian text, showing the irrelevance of minor language differences for performance of large transformer models. (4) While BiLSTM models are still close-to-competitive on the morphosyntactic tagging task, they cannot hold up on the named entity recognition task as it requires more common knowledge. Such knowledge transformer models manage to absorb to a much higher level than pre-trained static embeddings used by BiLSTMs.

The `BERTiĆ` model is available from the HuggingFace repository at <https://huggingface.co/CLASSLA/bcms-bertic>.

Acknowledgments

This work has been supported by the Slovenian Research Agency and the Flemish Research Foundation through the bilateral research project ARRS N6-0099 and FWO G070619N “The linguistic landscape of hate speech on social media”, the Slovenian Research Agency research core funding No. P6-0411 “Language resources and technologies for Slovene language”, and the European Union’s Rights, Equality and Citizenship Programme (2014-2020) project IMSyPP (grant no. 875263). We would like to thank the anonymous reviewers and Ivo-Pavao Jazbec for their useful feedback.

References

Vuk Batanović, Nikola Ljubešić, Tanja Samardžić, and Tomaž Erjavec. 2018. [Training corpus SE-Times.SR 1.0](#). Slovenian language resource repository CLARIN.SI.

Lukas Biewald. 2020. [Experiment tracking with weights and biases](#). Software available from wandb.com.

Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.

Damir Čavar and Dunja Brozović Rončević. 2012. Riznica: the Croatian language corpus. *Prace filologiczne*, 63:51–65.

Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*.

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 preprint arXiv:1911.02116*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Mihaela Gaman, Dirk Hovy, Radu Tudor Ionescu, Heidi Jauhiainen, Tommi Jauhiainen, Krister Lindén, Nikola Ljubešić, Niko Partanen, Christoph Purschke, Yves Scherrer, et al. 2020. A report on the VarDial evaluation campaign 2020. In *Proceedings of the 7th Workshop on NLP for Similar Languages, Varieties and Dialects*. International Committee on Computational Linguistics.

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. *arXiv preprint arXiv:1907.11692*.

Nikola Ljubešić. 2021. [Choice of plausible alternatives dataset in Croatian COPA-HR](#). Slovenian language resource repository CLARIN.SI.

Nikola Ljubešić, Željko Agić, Filip Klubička, Vuk Batanović, and Tomaž Erjavec. 2018. [Training corpus hr500k 1.0](#). Slovenian language resource repository CLARIN.SI.

Nikola Ljubešić and Kaja Dobrovoljc. 2019. [What does neural bring? analysing improvements in morphosyntactic annotation and lemmatisation of Slovenian, Croatian and Serbian](#). In *Proceedings of the 7th Workshop on Balto-Slavic Natural Language Processing*, pages 29–34, Florence, Italy. Association for Computational Linguistics.

Nikola Ljubešić and Tomaž Erjavec. 2011. hrWaC and slWaC: Compiling web corpora for Croatian and Slovene. In *International Conference on Text, Speech and Dialogue*, pages 395–402. Springer.

Nikola Ljubešić, Tomaž Erjavec, Vuk Batanović, Maja Miličević, and Tanja Samardžić. 2019a. [Croatian twitter training corpus ReLDI-NormTagNER-hr 2.1](#). Slovenian language resource repository CLARIN.SI.

- Nikola Ljubešić, Tomaž Erjavec, Vuk Batanović, Maja Miličević, and Tanja Samardžić. 2019b. [Serbian twitter training corpus ReLDI-NormTagNER-sr 2.1](#). Slovenian language resource repository CLARIN.SI.
- Nikola Ljubešić and Filip Klubička. 2014. {bs, hr, sr} wac-web corpora of Bosnian, Croatian and Serbian. In *Proceedings of the 9th Web as Corpus Workshop (WaC-9)*, pages 29–35.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric Villamonte de la Clergerie, Djamel Seddah, and Benoît Sagot. 2019. Camembert: a tasty french language model. *arXiv preprint arXiv:1911.03894*.
- Edoardo Maria Ponti, Goran Glavaš, Olga Majewska, Qianchu Liu, Ivan Vulić, and Anna Korhonen. 2020. [XCOPA: A multilingual dataset for causal commonsense reasoning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2362–2376, Online. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D Manning. 2020. Stanza: A Python natural language processing toolkit for many human languages. *arXiv preprint arXiv:2003.07082*.
- Melissa Roemmele, Cosmin Adrian Bejan, and Andrew S Gordon. 2011. Choice of Plausible Alternatives: An Evaluation of Commonsense Causal Reasoning. In *AAAI Spring Symposium: Logical Formalizations of Commonsense Reasoning*, pages 90–95.
- Yves Scherrer and Nikola Ljubešić. 2020. HeLju@VarDial 2020: Social media variety geolocation with BERT models. In *Proceedings of the 7th Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 202–211.
- Matej Ulčar and Marko Robnik-Šikonja. 2020. FinEst BERT and CroSloEngual BERT. In *International Conference on Text, Speech, and Dialogue*, pages 104–111. Springer.
- Wietse de Vries, Andreas van Cranenburgh, Arianna Bisazza, Tommaso Caselli, Gertjan van Noord, and Malvina Nissim. 2019. Bertje: A dutch BERT model. *arXiv preprint arXiv:1912.09582*.