

# HerBERT: Efficiently Pretrained Transformer-based Language Model for Polish

Robert Mroczkowski<sup>1</sup> Piotr Rybak<sup>1</sup> Alina Wróblewska<sup>2</sup> Ireneusz Gawlik<sup>1,3</sup>

<sup>1</sup>ML Research at Allegro.pl

<sup>2</sup>Institute of Computer Science, Polish Academy of Sciences

<sup>3</sup>Department of Computer Science, AGH University of Science and Technology  
{firstname.lastname}@allegro.pl, alina@ipipan.waw.pl

## Abstract

BERT-based models are currently used for solving nearly all Natural Language Processing (NLP) tasks and most often achieve state-of-the-art results. Therefore, the NLP community conducts extensive research on understanding these models, but above all on designing effective and efficient training procedures. Several ablation studies investigating how to train BERT-like models have been carried out, but the vast majority of them concerned only the English language. A training procedure designed for English does not have to be universal and applicable to other especially typologically different languages. Therefore, this paper presents the first ablation study focused on Polish, which, unlike the isolating English language, is a fusional language. We design and thoroughly evaluate a pretraining procedure of transferring knowledge from multilingual to monolingual BERT-based models. In addition to multilingual model initialization, other factors that possibly influence pretraining are also explored, i.e. training objective, corpus size, BPE-Dropout, and pretraining length. Based on the proposed procedure, a Polish BERT-based language model – HerBERT – is trained. This model achieves state-of-the-art results on multiple downstream tasks.

## 1 Introduction

Recent advancements in self-supervised pretraining techniques drastically changed the way we design Natural Language Processing (NLP) systems. Even though, pretraining has been present in NLP for many years (Mikolov et al., 2013; Pennington et al., 2014; Bojanowski et al., 2017), only recently we observed a shift from task-specific to general-purpose models. In particular, the BERT model (Devlin et al., 2019) proved to be a dominant architecture and obtained state-of-the-art results for a variety of NLP tasks.

While most of the research related to analyzing and improving BERT-based models focuses on English, there is an increasing body of work aimed at training and evaluation of models for other languages, including Polish. Thus far, a handful of models specific for Polish has been released, e.g. Polbert<sup>1</sup>, first version of HerBERT (Rybak et al., 2020), and Polish RoBERTa (Dadas et al., 2020).

Aforementioned works lack ablation studies, making it difficult to attribute hyperparameters choices to models performance. In this work, we fill this gap by conducting an extensive set of experiments and developing an efficient BERT training procedure. As a result, we were able to train and release a new BERT-based model for Polish language understanding. Our model establishes a new state-of-the-art on the variety of downstream tasks including semantic relatedness, question answering, sentiment analysis and part-of-speech tagging.

To summarize, our contributions are:

1. development and evaluation of an efficient pretraining procedure for transferring knowledge from multilingual to monolingual language models based on work by Arkhipov et al. (2019),
2. detailed analysis and an ablation study challenging the effectiveness of Sentence Structural Objective (SSO, Wang et al., 2020), and Byte Pair Encoding Dropout (BPE-Dropout, Provilkov et al., 2020),
3. release of HerBERT<sup>2</sup> – a BERT-based model for Polish language understanding, which achieves state-of-the-art results on KLEJ Benchmark (Rybak et al., 2020) and POS tagging task (Wróblewska, 2020).

<sup>1</sup><https://github.com/kldarek/polbert>

<sup>2</sup><https://huggingface.co/allegro/herbert-large-cased>

The rest of the paper is organized as follows. In Section 2, we provide an overview of related work. After that, Section 3 introduces the BERT-based language model and experimental setup used in this work. In Section 4, we conduct a thorough ablation study to investigate the impact of several design choices on the performance of downstream tasks. Next, in Section 5 we apply drawn conclusions and describe the training of HerBERT model. In Section 6, we evaluate HerBERT on a set of eleven tasks and compare its performance to other state-of-the-art models. Finally, we conclude our work in Section 7.

## 2 Related Work

The first significant ablation study of BERT-based language pretraining was described by Liu et al. (2019). Authors demonstrated the ineffectiveness of Next Sentence Prediction (NSP) objective, the importance of dynamic token masking, and gains from using both large batch size and large training dataset. Further large-scale studies analyzed the relation between the model and the training dataset sizes (Kaplan et al., 2020), the amount of compute used for training (Brown et al., 2020) and training strategies and objectives (Raffel et al., 2019).

Other work focused on studying and improving BERT training objectives. As mentioned before, the NSP objective was either removed (Liu et al., 2019) or enhanced either by predicting the correct order of sentences (Sentence Order Prediction (SOP), Lan et al., 2020) or discriminating between previous, next and random sentence (Sentence Structural Objective (SSO), Wang et al., 2020). Similarly, the Masked Language Modelling (MLM) objective was extended to either predict spans of tokens (Joshi et al., 2019), re-order shuffled tokens (Word Structural Objective (WSO), Wang et al., 2020) or replaced altogether with a binary classification problem using mask generation (Clark et al., 2020).

For tokenization, the Byte Pair Encoding algorithm (BPE, Sennrich et al., 2016) is commonly used. The original BERT model used WordPiece implementation (Schuster and Nakajima, 2012), which was later replaced by SentencePiece (Kudo and Richardson, 2018). Gong et al. (2018) discovered that rare words lack semantic meaning. Provilkov et al. (2020) proposed a BPE-Dropout technique to solve this issue.

All of the above work was conducted for En-

glish language understanding. There was little research into understanding how different pretraining techniques affect BERT-based models for other languages. The main research focus was to train BERT-based models and report their performance on downstream tasks. The first such models were released for German<sup>3</sup> and Chinese (Devlin et al., 2019), recently followed by Finnish (Virtanen et al., 2019), French (Martin et al., 2020; Le et al., 2020), Polish (Rybak et al., 2020; Dadas et al., 2020), Russian (Kuratov and Arkhipov, 2019), and many other languages<sup>4</sup>. Research on developing and investigating an efficient procedure of pretraining BERT-based models was rather neglected in these languages.

Language understanding for low-resource languages has also been addressed by training jointly for several languages at the same time. That approach improves performance for moderate and low-resource languages as showed by Conneau and Lample (2019). The first model of this kind was the multilingual BERT trained for 104 languages (Devlin et al., 2019) followed by Conneau and Lample (2019) and Conneau et al. (2020).

## 3 Experimental Setup

In this section, we describe the experimental setup used in the ablation study. First, we introduce the corpora we used to train models. Then, we give an overview of the language model architecture and training procedure. In particular, we describe the method of transferring knowledge from multilingual to monolingual BERT-based models. Finally, we present the evaluation tasks.

### 3.1 Training Data

We gathered six corpora to create two datasets on which we trained HerBERT. The first dataset (henceforth called *Small*) consists of corpora of the highest quality, i.e. NKJP, Wikipedia, and Wolne Lektury. The second dataset (*Large*) is over five times larger as it additionally contains texts of lower quality (CCNet and Open Subtitles). Below, we present a short description of each corpus. Additionally, we include the basic corpora statistics in Table 1.

**NKJP** (Narodowy Korpus Języka Polskiego, eng. *National Corpus of Polish*) (Przepiórkowski, 2012) is a well balanced collection of Polish texts. It

<sup>3</sup><https://deepset.ai/german-bert>

<sup>4</sup><https://huggingface.co/models>

Corpus	Tokens	Documents	Avg len
<b>Source Corpora</b>			
NKJP	1357M	3.9M	347
Wikipedia	260M	1.4M	190
Wolne Lektury	41M	5.5k	7447
CCNet Head	2641M	7.0M	379
CCNet Middle	3243M	7.9M	409
Open Subtitles	1056M	1.1M	961
<b>Final Corpora</b>			
Small	1658M	5.3M	313
Large	8599M	21.3M	404

Table 1: Overview of all data sources used to train HerBERT. We combine them into two corpora. The *Small* corpus consists of the highest quality text resources: NKJP, Wikipedia, and Wolne Lektury. The *Large* corpus consists of all sources. *Avg len* is the average number of tokens per document in each corpus.

consists of texts from many different sources, such as classic literature, books, newspapers, journals, transcripts of conversations, and texts crawled from the internet.

**Wikipedia** is an online encyclopedia created by the community of Internet users. Even though it is crowd-sourced, it is recognized as a high-quality collection of articles.

**Wolne Lektury** (eng. *Free Readings*)<sup>5</sup> is a collection of over five thousand books and poems, mostly from 19th and 20th century, which have already fallen in the public domain.

**CCNet** (Wenzek et al., 2020) is a clean monolingual corpus extracted from Common Crawl<sup>6</sup> dataset of crawled websites.

**Open Subtitles** is a popular website offering movie and TV subtitles, which was used by Lison and Tiedemann (2016) to curate and release a multilingual parallel corpus from which we extracted its monolingual Polish part.

### 3.2 Language Model

**Tokenizer** We used Byte-Pair Encoding (BPE) tokenizer (Sennrich et al., 2016) with the vocabulary size of 50k tokens and trained it on the most

<sup>5</sup><https://wolnelektury.pl>

<sup>6</sup><http://commoncrawl.org/>

representative parts of our corpus, i.e annotated subset of the NKJP, and the Wikipedia.

Subword regularization is supposed to emphasize the semantic meaning of tokens (Gong et al., 2018; Provilkov et al., 2020). To verify its impact on training language model we used a BPE-Dropout (Provilkov et al., 2020) with a probability of dropping a merge equal to 10%.

**Architecture** We followed the original BERT (Devlin et al., 2019) architectures for both BASE (12 layers, 12 attention heads and hidden dimension of 768) and LARGE (24 layers, 16 attention heads and hidden dimension of 1024) variants.

**Initialization** We initialized models either randomly or by using weights from XLM-RoBERTa (Conneau et al., 2020). In the latter case, the parameters for all layers except word embeddings and token type embeddings were copied directly from the source model. Since XLM-RoBERTa does not use the NSP objective and does not have the token type embeddings, we took them from the original BERT model.

To overcome the difference in tokenizers vocabularies we used a method similar to (Arkipov et al., 2019). If a token from the target model vocabulary was present in the source model vocabulary then we directly copied its weights. Otherwise, it was split into smaller units and the embedding was obtained by averaging sub-tokens embeddings.

**Training Objectives** We trained all models with an updated version of the MLM objective (Joshi et al., 2019; Martin et al., 2020), masking ranges of subsequent tokens belonging to single words instead of individual (possibly subword) tokens. We replaced the NSP objective with SSO. The other parameters were kept the same as in the original BERT paper. Training objective is defined in Equation 1.

$$\mathcal{L} = \mathcal{L}_{\text{MLM}}(\theta) + \alpha \cdot \mathcal{L}_{\text{SSO}}(\theta) \quad (1)$$

where  $\alpha$  is the SSO weight.

### 3.3 Tasks

**KLEJ Benchmark** The standard method for evaluating pretrained language models is to use a diverse collection of tasks grouped into a single benchmark. Such benchmarks exist in many languages, e.g. English (GLUE, Wang et al., 2019), Chinese (CLUE, Xu et al., 2020), and Polish (KLEJ, Rybak et al., 2020).

Following this paradigm we first verified the quality of assessed models with KLEJ. It consists of nine tasks: name entity classification (NKJP-NER, Przepiórkowski, 2012), semantic relatedness (CDSC-R, Wróblewska and Krasnowska-Kieraś, 2017), natural language inference (CDSC-E, Wróblewska and Krasnowska-Kieraś, 2017), cyberbullying detection (CBD, Ptaszynski et al., 2019), sentiment analysis (PolEmo2.0-IN, PolEmo2.0-OUT, Kocoń et al., 2019, AR, Rybak et al., 2020), question answering (Czy wiesz?, Marcinczuk et al., 2013), and text similarity (PSC, Ogrodniczuk and Kopeć, 2014).

**POS Tagging and Dependency Parsing** All of the KLEJ Benchmark tasks belong to the classification or regression type. It is therefore difficult to assess the quality of individual token embeddings. To address this issue, we further evaluated HerBERT on part-of-speech tagging and dependency parsing tasks.

For tagging, we used the manually annotated subset of NKJP (Degórski and Przepiórkowski, 2012), converted to the CoNLL-U format by Wróblewska (2020). We evaluated models performance on a test set using accuracy and F1-Score.

For dependency parsing, we applied Polish Dependency Bank (Wróblewska, 2018) from the Universal Dependencies repository (release 2.5, Zeman et al., 2019).

In addition to three Transformer-based models, we also included models trained with static embeddings. The first one did not use pretrained embeddings while the latter utilized fastText (Bojanowski et al., 2017) embeddings trained on Common Crawl.

The models are evaluated with the standard metrics: UAS (unlabeled attachment score) and LAS (labelled attachment score). The gold-standard segmentation was preserved. We report the results on the test set.

## 4 Ablation Study

In this section, we analyze the impact of several design choices on downstream task performance of Polish BERT-based models. In particular, we focus on initialization, corpus size, training objective, BPE-Dropout, and the length of pretraining.

### 4.1 Experimental Design

**Hyperparameters** Unless stated otherwise, in all experiments we trained BERT<sub>BASE</sub> model ini-

tialized with XLM-RoBERTa weights for 10k iterations using a linear decay schedule of the learning rate with a peak value of  $7 \cdot 10^{-4}$  and a warm-up of 500 iterations. We used a batch size of 2560.

**Evaluation** Different experimental setups were compared using the average score on the KLEJ Benchmark. The validation sets are used for evaluation and we report the results corresponding to the median values of the five runs. Since only six tasks in KLEJ Benchmark have validation sets the scores are not directly comparable to those reported in Section 6. We used Welch’s t-test (Welch, 1947) with a p-value of 0.01 to test for statistical differences between experimental variants.

## 4.2 Results

**Initialization** One of the main goals of this work is to propose an efficient strategy to train a monolingual language model. We began with investigating the impact of pretraining the language model itself. For this purpose, the following experiments were designed.

Init	Pretraining	BPE	Score
<b>Ablation Models</b>			
Random	No	-	58.15 ± 0.33
XLM-R	No	-	<u>83.15 ± 1.22</u>
Random	Yes	No	85.65 ± 0.43
XLM-R	Yes	No	<u>88.80 ± 0.15</u>
Random	Yes	Yes	85.78 ± 0.23
XLM-R	Yes	Yes	<b><u>89.10 ± 0.19</u></b>
<b>Original Models</b>			
XLM-R	-	-	88.82 ± 0.15

Table 2: Average scores on KLEJ Benchmark depending on the initialization scheme: Random – initialization with random weights, XLM-R – initialization with XLM-RoBERTa weights. We used BERT<sub>BASE</sub> model trained for 10k iterations with the SSO weight equal to 1.0 on the *Large* corpus. The best score within each group is underlined, the best overall is bold.

First, we fine-tuned randomly initialized BERT model on KLEJ Benchmark tasks. Note that this model is not pretrained in any way. As expected, the results on the KLEJ Benchmark are really poor with the average score equal to 58.15.

Next, we evaluated the BERT model initialized

with XLM-RoBERTa weights (see Table 2). It achieved much better average score than the randomly initialized model (83.15 vs 58.15), but it was still not as good as the original XLM-RoBERTa model (88.82%). The difference in the performance can be explained by the transfer efficiency. The method of transferring token embeddings between different tokenizers proves to retain most information, but not all of it.

To measure the impact of initialization on pre-training optimization, we trained the aforementioned models for 10k iterations. Beside MLM objective, we used SSO loss with  $\alpha = 1.0$  and conducted experiments with both enabled and disabled BPE-Dropout. Models initialized with XLM-RoBERTa achieve significantly higher results than models initialized randomly, 89.10 vs 85.78 and 88.80 vs 85.65 for pretraining with and without BPE-Dropout respectively.

Models initialized with XLM-RoBERTa achieved similar results to the original XLM-RoBERTa (the differences are not statistically significant). It proves that it is possible to quickly recover from the performance drop caused by a tokenizer conversion procedure and obtain a much better model than the one initialized randomly.

**Corpus Size** As mentioned in Section 2, previous research show that pretraining on a larger corpus is beneficial for downstream task performance (Kaplan et al., 2020; Brown et al., 2020). We investigated this by pretraining BERT<sub>BASE</sub> model on both *Small* and *Large* corpora (see Section 3.1). To mitigate a possible impact of confounding variable, we also vary the weight of SSO loss and usage of BPE-Dropout (see Table 3).

As expected, the model pretrained on a *Large* corpus performs better on downstream tasks. However, the difference is statistically significant only for the experiment with SSO weight equal to 1.0 and BPE-Dropout enabled. Therefore it’s not obvious whether a larger corpus is actually beneficial.

**Sentence Structural Objective** Subsequently, we tested SSO, i.e. the recently introduced replacement for the NSP objective, which proved to be ineffective. We compared models trained with three values of SSO weight  $\alpha$  (see Section 3.2): 0.0 (no SSO), 0.1 (small impact of SSO), and 1.0 (SSO equally important as MLM objective) (see Table 4).

The experiment showed, that SSO actually hurts

Corpus	SSO	BPE	Score
Small	1.0	Yes	88.73 $\pm$ 0.08
Large	1.0	Yes	<u>89.10 <math>\pm</math> 0.19</u>
Small	0.1	Yes	88.90 $\pm$ 0.24
Large	0.1	Yes	<b>89.37 <math>\pm</math> 0.25</b>
Small	0.0	Yes	89.18 $\pm$ 0.15
Large	0.0	Yes	<u>89.25 <math>\pm</math> 0.21</u>
Small	0.0	No	89.12 $\pm$ 0.29
Large	0.0	No	<u>89.28 <math>\pm</math> 0.26</u>

Table 3: Average scores on KLEJ Benchmark depending on a corpus size. We used BERT<sub>BASE</sub> model trained for 10k iterations with or without BPE-Dropout and with various SSO weights. The best score within each group is underlined, the best overall is bold.

downstream task performance. The differences between enabled and disabled SSO are statistically significant for two out of three experimental setups. The only scenario for which the negative effect of SSO is not statistically significant is using the Large corpus and BPE-dropout. Overall, the best results are achieved using a small SSO weight but the differences are not significantly different from disabling SSO.

SSO	Corpus	BPE	Score
0.0	Small	Yes	<u>89.18 <math>\pm</math> 0.15</u>
0.1	Small	Yes	88.90 $\pm$ 0.24
1.0	Small	Yes	88.73 $\pm$ 0.08
0.0	Large	Yes	89.25 $\pm$ 0.21
0.1	Large	Yes	<u>89.37 <math>\pm</math> 0.25</u>
1.0	Large	Yes	89.10 $\pm$ 0.19
0.0	Large	No	89.28 $\pm$ 0.26
0.1	Large	No	<b>89.45 <math>\pm</math> 0.18</b>
1.0	Large	No	88.80 $\pm$ 0.15

Table 4: Average scores on KLEJ Benchmark depending on a SSO weight. We used BERT<sub>BASE</sub> model trained for 10k iterations with BPE-Dropout. The best score within each group is underlined, the best overall is bold.

**BPE-Dropout** The BPE-Dropout could be beneficial for downstream task performance, but its impact is difficult to assess due to many confound-

ing variables.

The model initialization with XLM-RoBERTa weights means that token embedding is already semantically meaningful even without additional pretraining. However, for both random and XLM-RoBERTa initialization the BPE-Dropout is beneficial.

BPE	Init	SSO	Corpus	Score
No	Random	1.0	Large	85.65 ± 0.43
Yes	Random	1.0	Large	<u>85.78 ± 0.23</u>
No	XLM-R	1.0	Large	88.80 ± 0.15
Yes	XLM-R	1.0	Large	<u>89.10 ± 0.19</u>
No	XLM-R	0.0	Large	<b>89.28 ± 0.26</b>
Yes	XLM-R	0.0	Large	89.25 ± 0.21
No	XLM-R	0.0	Small	89.12 ± 0.29
Yes	XLM-R	0.0	Small	<u>89.18 ± 0.15</u>

Table 5: Average scores on KLEJ Benchmark depending on usage of BPE-Dropout. We used BERT<sub>BASE</sub> model trained for 10k iterations on a large corpus. The best score within each group is underlined, the best overall is bold.

According to the results (see Table 5), none of the reported differences is statistically significant and we can only conclude that BPE-Dropout does not influence the model performance.

**Length of Pretraining** The length of pretraining in terms of the number of iterations is commonly considered an important factor of the final quality of the model (Kaplan et al., 2020). Even though it seems straightforward to validate this hypothesis in practice it is not so trivial.

When pretraining Transformer-based models linear decaying learning rate is typically used. Therefore, increasing the number of training iterations changes the learning rate schedule and impacts the training. In our initial experiments usage of the same learning rate caused the longer training to collapse. Instead, we chose the learning rate for which the value of loss function after 10k steps was similar. We found that the learning rate equal to  $3 \cdot 10^{-4}$  worked best for training in 50k steps.

Using the presented experiment setup, we tested the impact of pretraining length for two values of SSO weight: 1.0 and 0.1. In both cases, the model pretrained with more iterations achieves

# Iter	LR	SSO	Score
10k	$7 \cdot 10^{-4}$	1.0	89.10 ± 0.19
50k	$3 \cdot 10^{-4}$	1.0	<u>89.43 ± 0.10</u>
10k	$7 \cdot 10^{-4}$	0.1	89.37 ± 0.25
50k	$3 \cdot 10^{-4}$	0.1	<b>89.87 ± 0.22</b>

Table 6: Average scores on KLEJ Benchmark depending training length. We used BERT<sub>BASE</sub> model trained on a large corpus with BPE-Dropout. The best score within each group is underlined, the best overall is bold.

only slightly better but statistically significant results (see Table 6).

## 5 HerBERT

In this section, we apply conclusions drawn from the ablation study (see Section 4) and describe the final pretraining procedure used to train HerBERT model.

**Pretraining Procedure** HerBERT was trained on the *Large* corpus. We used Dropout-BPE in tokenizer with a probability of a drop equals to 10%. Finally, HerBERT models were initialized with weights from XLM-RoBERTa and were trained with the objective defined in Equation 1 with SSO weight equal to 0.1.

**Optimization** We trained HerBERT<sub>BASE</sub> using Adam optimizer (Kingma and Ba, 2014) with parameters:  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 10^{-8}$  and a linear decay learning rate schedule with a peak value of  $3 \cdot 10^{-4}$ . Due to the initial transfer of weights from the already trained model, the warm-up stage was set to a relatively small number of 500 iterations. The whole training took 50k iterations.

Training of HerBERT<sub>LARGE</sub> was longer (60k iterations) and had a more complex learning rate schedule. For the first 15k we linearly decayed the learning rate from  $3 \cdot 10^{-4}$  to  $2.5 \cdot 10^{-4}$ . We observed the saturation of evaluation metrics and decided to drop the learning rate to  $1 \cdot 10^{-4}$ . After training for another 25k steps and reaching the learning rate of  $7 \cdot 10^{-5}$  we again reached the plateau of evaluation metrics. In the last phase of training, we dropped the learning rate to  $3 \cdot 10^{-5}$  and trained for 20k steps until it reached zero. Additionally, during the last phase of training, we disabled both BPE-Dropout and dropout within the Transformer itself as suggested by Lan et al. (2020).

Model	AVG	NKJP-NER	CDSC-E	CDSC-R	CBD	PolEmo2.0-IN	PolEmo2.0-OUT	Czy wiesz?	PSC	AR
<b>Base Models</b>										
XLM-RoBERTa	84.7 ± 0.29	91.7	93.3	93.4	66.4	<u>90.9</u>	77.1	64.3	97.6	87.3
Polish RoBERTa	85.6 ± 0.29	94.0	94.2	<u>94.2</u>	63.6	90.3	76.9	71.6	98.6	87.4
HerBERT	<u>86.3 ± 0.36</u>	<u>94.5</u>	<u>94.5</u>	94.0	<u>67.4</u>	<u>90.9</u>	<u>80.4</u>	<u>68.1</u>	<u>98.9</u>	<u>87.7</u>
<b>Large Models</b>										
XLM-RoBERTa	86.8 ± 0.30	94.2	<b>94.7</b>	93.9	67.6	92.1	81.6	70.0	98.3	88.5
Polish RoBERTa	87.5 ± 0.29	94.9	93.4	94.7	69.3	<b>92.2</b>	81.4	74.1	<b>99.1</b>	88.6
HerBERT	<b>88.4 ± 0.19</b>	<b>96.4</b>	94.1	<b>94.9</b>	<b>72.0</b>	<b>92.2</b>	<b>81.8</b>	<b>75.8</b>	98.9	<b>89.1</b>

Table 7: Evaluation results on KLEJ Benchmark. AVG is the average score across all tasks. Scores are reported for test set and correspond to median values across five runs. The best scores within each group are underlined, the best overall are in bold.

Both HerBERT<sub>BASE</sub> and HerBERT<sub>LARGE</sub> models were trained with a batch size of 2560.

## 6 Evaluation

In this section, we introduce other top-performing models for Polish language understanding and compare their performance on evaluation tasks (see Section 3.3) to HerBERT.

### 6.1 Models

According to the KLEJ Benchmark leaderboard<sup>7</sup> the three top-performing models of Polish language understanding are XLM-RoBERTa-NKJP<sup>8</sup>, Polish RoBERTa, and XLM-RoBERTa. These are also the only three models available in LARGE architecture variant.

Unfortunately, the XLM-RoBERTa-NKJP model is not publicly available, so we cannot use it for our evaluation. However, it has the same average score as the runner-up (Polish RoBERTa) which we compare HerBERT with.

### 6.2 Results

**KLEJ Benchmark** Both variants of HerBERT achieved the best average performance, signifi-

<sup>7</sup><https://klejbenchmark.com/leaderboard/>

<sup>8</sup>XLM-RoBERTa-NKJP is XLM-RoBERTa model additionally fine-tuned on NKJP corpus.

cantly outperforming Polish RoBERTa and XLM-RoBERTa (see Table 7). Regarding BASE models, HerBERT<sub>BASE</sub> improves the state-of-the-art result by 0.7pp and for HerBERT<sub>LARGE</sub> the improvement is even bigger (0.9pp). In particular, HerBERT<sub>BASE</sub> scores best on eight out of nine tasks (tying on PolEmo2.0-OUT, performing slightly worse on CDSC-R) and HerBERT<sub>LARGE</sub> in seven out of nine tasks (tying in PolEmo2.0-IN, performing worse in CDSC-E and PSC). Surprisingly, HerBERT<sub>BASE</sub> is better than HerBERT<sub>LARGE</sub> in CDSC-E. The same behaviour is noticeable for Polish RoBERTa, but not for XLM-RoBERTa. For other tasks, the LARGE models are always better. Summing up, HerBERT<sub>LARGE</sub> is the new state-of-the-art Polish language model based on the results of the KLEJ Benchmark.

It is worth emphasizing that the proposed procedure for efficient pretraining by transferring knowledge from multilingual to monolingual language allowed HerBERT to achieve better results than Polish RoBERTa even though it was optimized with around ten times shorter training.

**POS Tagging** HerBERT achieves overall better results in terms of both accuracy and F1-Score. HerBERT<sub>BASE</sub> beats the second-best model (i.e. Polish RoBERTa) by a margin of 0.33pp (F1-Score by 0.35pp) and HerBERT<sub>LARGE</sub> by a margin of 0.09pp (F1-Score by 0.12pp). It should be em-

Model	Accuracy	F1-Score
<b>Base Models</b>		
XLM-RoBERTa	$95.97 \pm 0.04$	$95.79 \pm 0.05$
Polish RoBERTa	$96.13 \pm 0.03$	$95.92 \pm 0.03$
HerBERT	<u><math>96.46 \pm 0.04</math></u>	<u><math>96.27 \pm 0.04</math></u>
<b>Large Models</b>		
XLM-RoBERTa	$97.07 \pm 0.05$	$96.93 \pm 0.05$
Polish RoBERTa	$97.21 \pm 0.02$	$97.05 \pm 0.03$
HerBERT	<b><u><math>97.30 \pm 0.02</math></u></b>	<b><u><math>97.17 \pm 0.02</math></u></b>

Table 8: Part-of-speech tagging results on NKJP dataset. Scores are reported for the test set and are median values across five runs. Best scores within each group are underlined, best overall are bold.

phasized that while the improvements may appear to be minor, they are statistically significant. All results are presented in Table 8.

**Dependency Parsing** The dependency parsing results are much more ambiguous than in other tasks. As expected, the models with static FastText embeddings performed much worse than Transformer-based models (around 3pp difference for UAS, and 4pp for LAS). In the case of Transformer-based models, the differences are less noticeable. As expected, the LARGE models outperform the BASE models. The best performing model is Polish RoBERTa. HerBERT models performance is the worst across Transformer-based models except for the UAS score which is slightly better than XLM-RoBERTa for BASE models. All results are presented in Table 9.

## 7 Conclusion

In this work, we conducted a thorough ablation study regarding training BERT-based models for Polish language. We evaluated several design choices for pretraining BERT outside of English language. Contrary to Wang et al. (2020), our experiments demonstrated that SSO is not beneficial for the downstream task performance. It also turned out that BPE-Dropout does not increase the quality of a pretrained language model.

As a result of our studies we developed and evaluated an efficient pretraining procedure for transferring knowledge from multilingual to monolingual BERT-based models. We used it to train and release HerBERT – a Transformer-based language model

Model	UAS	LAS
<b>Static Embeddings</b>		
Plain	$90.58 \pm 0.07$	$87.35 \pm 0.12$
FastText	<u><math>92.20 \pm 0.14</math></u>	<u><math>89.57 \pm 0.13</math></u>
<b>Base Models</b>		
XLM-RoBERTa	$95.14 \pm 0.07$	$93.25 \pm 0.12$
Polish RoBERTa	<u><math>95.41 \pm 0.24</math></u>	<u><math>93.65 \pm 0.34</math></u>
HerBERT	$95.18 \pm 0.22$	$93.24 \pm 0.23$
<b>Large Models</b>		
XLM-RoBERTa	$95.38 \pm 0.02$	$93.66 \pm 0.07$
Polish RoBERTa	<b><u><math>95.60 \pm 0.18</math></u></b>	<b><u><math>93.90 \pm 0.21</math></u></b>
HerBERT	$95.11 \pm 0.04$	$93.32 \pm 0.02$

Table 9: Dependency parsing results on Polish Dependency Bank dataset. Scores are reported for the test set and are median values across three runs. Best scores within each group are underlined, best overall are bold.

for Polish. It was trained on a diverse multi-source corpus. The conducted experiments confirmed its high performance on a set of eleven diverse linguistic tasks, as HerBERT turned out to be the best on eight of them. In particular, it is the best model for Polish language understanding according to the KLEJ Benchmark.

It is worth emphasizing that the quality of the obtained language model was even more impressive considering its short training time. Due to multilingual initialization, HerBERT<sub>BASE</sub> outperformed Polish RoBERTa<sub>BASE</sub> even though it was trained with a smaller batch size (2560 vs 8000) for a fewer number of steps (50k vs 125k). The same behaviour is also visible for HerBERT<sub>LARGE</sub>. Additionally, we conducted a separate ablation study to confirm that the success of HerBERT is caused by the described initialization scheme. It showed that in fact, it was the most important factor to improved the quality of HerBERT.

We believe that the proposed training procedure and detailed experiments will encourage NLP researchers to cost-effectively train language models for other languages.

## References

- Mikhail Arhipov, Maria Trofimova, Yuri Kuratov, and Alexey Sorokin. 2019. [Tuning multilingual transformers for language-specific named entity recognition](#). In *Proceedings of the 7th Workshop on Balto-Slavic Natural Language Processing*, pages 89–93, Florence, Italy. Association for Computational Linguistics.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal et al. 2020. [Language models are few-shot learners](#). *CoRR*, abs/2005.14165.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: pre-training text encoders as discriminators rather than generators](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán et al. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 7059–7069. Curran Associates, Inc.
- Sławomir Dadas, Michał Perełkiewicz, and Rafał Poświata. 2020. [Pre-training polish transformer-based language models at scale](#).
- Łukasz Degórski and Adam Przepiórkowski. 2012. *Recznie znakowany milionowy podkorpus NKJP*, pages 51–58.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Chengyue Gong, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tie-Yan Liu. 2018. [Frage: Frequency-agnostic word representation](#). In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 1334–1345. Curran Associates, Inc.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2019. Spanbert: Improving pre-training by representing and predicting spans. *arXiv preprint arXiv:1907.10529*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child et al. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). Cite arxiv:1412.6980Comment: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015.
- Jan Kocoń, Piotr Miłkowski, and Monika Zaśko-Zielińska. 2019. [Multi-level sentiment analysis of PolEmo 2.0: Extended corpus of multi-domain consumer reviews](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 980–991, Hong Kong, China. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Yuri Kuratov and Mikhail Arhipov. 2019. [Adaptation of deep bidirectional multilingual transformers for russian language](#).
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [ALBERT: A lite BERT for self-supervised learning of language representations](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Hang Le, Loïc Vial, Jibril Frej, Vincent Segonne, Maximin Coavoux, Benjamin Lecouteux et al. 2020. [FlauBERT: Unsupervised language model pre-training for French](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2479–2490, Marseille, France. European Language Resources Association.
- Pierre Lison and Jörg Tiedemann. 2016. [Opensubtitles2016: Extracting large parallel corpora from movie and tv subtitles](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France. European Language Resources Association (ELRA).
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen et al. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

- Michał Marcinczuk, Marcin Ptak, Adam Radziszewski, and Maciej Piasecki. 2013. Open dataset for development of polish question answering systems. In *Proceedings of the 6th Language & Technology Conference: Human Language Technologies as a Challenge for Computer Science and Linguistics*, Wydawnictwo Poznańskie, Fundacja Uniwersytetu im. Adama Mickiewicza.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie et al. 2020. [CamemBERT: a tasty French language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online. Association for Computational Linguistics.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Maciej Ogrodniczuk and Mateusz Kopeć. 2014. The Polish Summaries Corpus. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Ivan Provilkov, Dmitrii Emelianenko, and Elena Voita. 2020. BPE-dropout: Simple and effective subword regularization. pages 1882–1892.
- Adam Przepiórkowski. 2012. *Narodowy korpus języka polskiego*. Naukowe PWN.
- Michał Ptaszynski, Agata Pieciukiewicz, and Paweł Dybała. 2019. Results of the poleval 2019 shared task 6: First dataset and open shared task for automatic cyberbullying detection in polish twitter. *Proceedings of the PolEval 2019 Workshop*, page 89.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena et al. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *CoRR*, abs/1910.10683.
- Piotr Rybak, Robert Mroczkowski, Janusz Tracz, and Ireneusz Gawlik. 2020. [KLEJ: Comprehensive benchmark for polish language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1191–1201, Online. Association for Computational Linguistics.
- Mike Schuster and Kaisuke Nakajima. 2012. Japanese and korean voice search. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5149–5152. IEEE.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Association for Computational Linguistics (ACL)*, pages 1715–1725.
- Antti Virtanen, Jenna Kanerva, Rami Ilo, Jouni Luoma, Juhani Luotolahti, Tapio Salakoski et al. 2019. [Multilingual is not enough: Bert for finnish](#).
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In the Proceedings of ICLR.
- Wei Wang, Bin Bi, Ming Yan, Chen Wu, Jiangnan Xia, Zuyi Bao et al. 2020. [Structbert: Incorporating language structures into pre-training for deep language understanding](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- B. L. Welch. 1947. [THE GENERALIZATION OF ‘STUDENT’S’ PROBLEM WHEN SEVERAL DIFFERENT POPULATION VARIANCES ARE INVOLVED](#). *Biometrika*, 34(1-2):28–35.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin et al. 2020. [CCNet: Extracting high quality monolingual datasets from web crawl data](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4003–4012, Marseille, France. European Language Resources Association.
- Alina Wróblewska. 2018. Extended and enhanced polish dependency bank in universal dependencies format. In *Proceedings of the Second Workshop on Universal Dependencies (UDW 2018)*, pages 173–182. Association for Computational Linguistics.
- Alina Wróblewska. 2020. [Towards the conversion of National Corpus of Polish to Universal Dependencies](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5308–5315, Marseille, France. European Language Resources Association.
- Alina Wróblewska and Katarzyna Krasnowska-Kieraś. 2017. Polish evaluation dataset for compositional distributional semantics models. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 784–792.
- Liang Xu, Xuanwei Zhang, Lu Li, Hai Hu, Chenjie Cao, Weitang Liu et al. 2020. [Clue: A chinese language understanding evaluation benchmark](#).
- Daniel Zeman, Joakim Nivre, Mitchell Abrams, Noëmi Aeppli, Željko Agić, Lars Ahrenberg et al. 2019. [Universal dependencies 2.5](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.