

# Wine is Not v i n.

## On the Compatibility of Tokenizations Across Languages

Antonis Maronikolakis\* and Philipp Dufter\* and Hinrich Schütze  
Center for Information and Language Processing (CIS), LMU Munich, Germany  
antmarakis, philipp@cis.lmu.de

### Abstract

The size of the vocabulary is a central design choice in large pretrained language models, with respect to both performance and memory requirements. Typically, subword tokenization algorithms such as byte pair encoding and WordPiece are used. In this work, we investigate the compatibility of tokenizations for multilingual static and contextualized embedding spaces and propose a measure that reflects the compatibility of tokenizations across languages. Our goal is to prevent incompatible tokenizations, e.g., "wine" (word-level) in English vs. "v i n" (character-level) in French, which make it hard to learn good multilingual semantic representations. We show that our compatibility measure allows the system designer to create vocabularies across languages that are compatible – a desideratum that so far has been neglected in multilingual models.

## 1 Introduction

Pretrained language models (Howard and Ruder, 2018; Peters et al., 2018; Devlin et al., 2019) have become the de-facto standard in natural language processing (NLP). Due to memory constraints and to overcome sparsity, subword tokenization models such as byte pair encoding (Sennrich et al., 2016) or WordPiece tokenization (Schuster and Nakajima, 2012) are most commonly used. Multilingual pretrained language models (PLMs) such as mBERT (Devlin et al., 2019) are mostly trained in an unsupervised fashion without any cross-lingual supervision. They rely on the fact that the underlying distributions across languages can be well aligned (Conneau et al., 2020b).

However, there is little work that investigates the compatibility of tokenizations across languages. We argue that this is an important research question for several reasons. **i)** Recent work suggests that the tokenization in multilingual models influences

\*Equal contribution.




Figure 1: Compatibility of static embedding vocabulary sizes measured by SVG (singular value gap) as a function of vocabulary size of English (x-axis) and Russian/Chinese (y-axis). Blue (red) indicates larger (smaller) compatibility. Black stars indicate vocabulary sizes where absolute compression rates of both languages are equal. The figure confirms our intuitions about vocabulary size compatibility: for English/Russian, similar vocabulary sizes are compatible, but for Chinese, we only find compatibility for larger English vocabulary sizes. This is because a Chinese logogram by itself is usually a meaning bearing unit, so that most subword (or sub-root) elements in English have no analog in Chinese.

the multilinguality of these models (Rust et al., 2021). This corresponds to a strong intuition: in case of heavily diverging tokenizations, e.g., word tokenization in English and character tokenization in French it will be hard to train common cross-lingual representations. **ii)** Given a limited vocabulary budget, e.g., from memory constraints, a researcher is naturally interested in how to distribute the budget across multiple languages. **iii)** Currently, a balance of the tokenization is heuristically achieved through up- and downsampling (e.g., Conneau et al., 2020a). With a better tokenization strategy, one might be able to achieve performance improvements.

Research in multilinguality has so far focused, among other directions, on analyzing and improving cross-lingual transfer (Artetxe et al., 2020a; K et al., 2020; Wu and Dredze, 2019; Conneau et al., 2020b) and modeling tasks (Conneau and Lam-

ple, 2019; Huang et al., 2019). Earlier work on tokenization focused on the development of algorithms (Schuster and Nakajima, 2012; Gage, 1994) while analysis of the effect of tokenization has been undertaken both in the general (Bostrom and Durrett, 2020) and in the multilingual domain (Wei et al., 2021).

We propose to investigate the compatibility of tokenizations across languages systematically. To this end, we train both static and contextualized embeddings using different vocabulary sizes and WordPiece tokenization (Schuster and Nakajima, 2012). Subsequently we investigate the similarity of the embedding spaces and their degree of multilinguality.

Figure 1 shows an example of our research question: here we trained static embedding spaces and investigated the similarity of the embedding spaces by computing the singular value gap (SVG) measure. The plots show that different vocabulary sizes yield very different similarities. In addition, language pairs behave quite differently. See the caption for more information.

In summary, our contributions are:<sup>1</sup>

1. We propose an experimental setting that allows the comparison of embeddings at different tokenization levels with each other.
2. We provide evidence that vocabulary size compatibility matters, i.e., if tokenizations across languages are incompatible it is hard to achieve multilingual representations.
3. We propose a measure that indicates the compatibility of tokenizations.

## 2 Related Work

In a position paper, Artetxe et al. (2020b) advocated for more rigor in cross-lingual research. The topic of tokenization was raised, especially the fact that current word and subword tokenization algorithms do not adequately capture morphological nuances and do not do justice to the great variety of languages in the world like those that use logographic scripts. These considerations motivate us to propose a method that makes better use of vocabulary space, a facet of multilingual models that is often overlooked.

<sup>1</sup>Code available at [https://github.com/antmarakis/vocab\\_size\\_compat](https://github.com/antmarakis/vocab_size_compat)

### 2.1 Multilingual Research

Investigation of the multilinguality of BERT has gathered momentum in recent years. Wu and Dredze (2019); Nooralahzadeh et al. (2020); Artetxe et al. (2020a); Conneau et al. (2020b) investigate cross-lingual transfer, with an extensive study of multilinguality at-scale presented in Conneau et al. (2020a). K et al. (2020) research how linguistic variation in language properties affects multilingual BERT models, with Dufter and Schütze (2020) further investigating the effect of model architecture and linguistic properties on BERT’s multilinguality. Conneau and Lample (2019) focus on variations of modeling tasks and their effect on downstream tasks such as XNLI (Conneau et al., 2018) and machine translation. A novel set of cross-lingual pre-training tasks was introduced in Huang et al. (2019). Anastasopoulos and Neubig (2020) show that performance in bilingual experiments is not optimal when English is used as the “hub”, but it depends on the language pair.

Work on the representation space of BERT has also been abundant. Singh et al. (2019) investigate the representation space of mBERT and its properties and Pires et al. (2019) initially researched the degree of multilinguality of mBERT.

In our work, we further investigate the zero-shot capabilities of bilingual models and present a method for comparing multilingual embeddings across tokenizations.

### 2.2 Tokenization Research

Another vein of research into model analysis has been the study of the effect of tokenization. Heinzlering and Strube (2018) compute BPE embeddings for 275 languages and perform a cross-lingual study on various tokenization methods. Further study into BPEs is conducted by Wei et al. (2021), where byte-level BPEs are compared against their character-level counterparts as it pertains to multilingual models while BPE vocabulary size is examined in the context of neural machine translation by Gowda and May (2020). Bostrom and Durrett (2020) show that byte-pair encoding is suboptimal when pretraining language models. In our work, we investigate another prevalent tokenization method, WordPiece, and show that vocabulary size plays an important role in achieving multilinguality in BERT.

Research into the selection of tokenization algorithms has produced alternative methods. For

example, Kudo (2018) present a subword regularization method where a model is trained with multiple tokenizations in order to reinforce robustness in the model; and Aguilar et al. (2020) bridge the gap between subword and character-level models, proposing a module that learns to approximate subword embeddings given characters. Asgari et al. (2020); Provilkov et al. (2020) further research subword tokenization methods. With our work we introduce a measure that can help NLP researchers pick vocabulary sizes with a more robust procedure.

### 3 Vocabulary Sizes

#### 3.1 Tokenization

Let  $\mathcal{C}$  be the set of unicode code points. We define a tokenizer as a function  $\tau : \mathcal{C}^t \rightarrow \mathcal{V}_\tau^s$  that maps a sequence of  $t$  unicode code points to a sequence of  $s$  vocabulary tokens.  $\mathcal{V}_\tau$  is the vocabulary of the tokenizer and  $n_\tau := |\mathcal{V}_\tau|$  is the vocabulary size of the tokenizer.

Usually the tokenizer is trained on a corpus  $\mathcal{U}$ . Popular tokenizers are byte-pair encoding (Sennrich et al., 2016), WordPiece (Schuster and Nakajima, 2012) and SentencePiece (Kudo and Richardson, 2018). An example tokenization using the WordPiece tokenizer of the PLM bert-base-multilingual-cased for the input  $c = \text{“Exceptional weather.”}$  is  $\tau(c) = [\text{“Ex”, “##ception”, “##al”, “weather”, “.”}]$ , where “##” is a continuation symbol.

Throughout this paper we use the WordPiece tokenizer. The main motivation is that important successful PLMs use this tokenization method. In future work we plan to extend this analysis to other tokenization algorithms as well.

#### 3.2 Compression Rate

Rust et al. (2021) use two measures to compare tokenizations across languages. However, they assume that each language has the notion of “words” that are separated by whitespace. For example they compute the proportion of words that are split into multiple tokens by the tokenizer.

We posit the following desiderata for a measure that compares tokenizations across languages: **i)** Assume that a text consists of a sequence of unicode code points. Do not require the existence of “words”. **ii)** Take into account vocabulary size and alphabet size. **iii)** Take into account the frequency of tokens in a corpus.

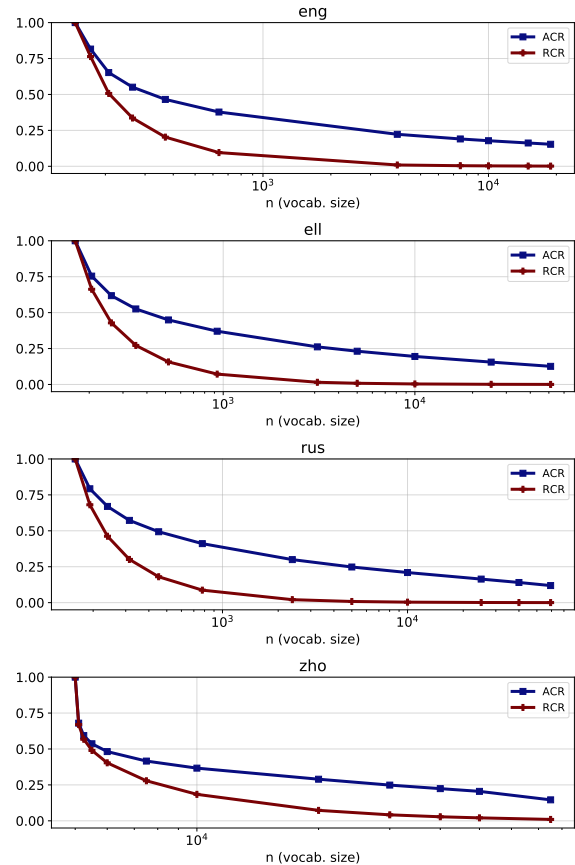


Figure 2: Compression rates for four languages and different vocabulary sizes. One can see that Chinese using logograms has a much steeper curve.

The first two desiderata should ensure that the measure is as inclusive as possible with regard to the variety of languages; the last desideratum ensures that the measure considers how “important” different tokens are in a corpus. Based on these desiderata, we propose two measures, an absolute and a relative compression rate.

**ACR  $r^{ABS}$ .** Let  $c \in \mathcal{C}^t$  be a corpus with  $t$  unicode code points and  $\tau_n(c)$  be the tokenized version of the corpus for a tokenizer with vocabulary size  $n$ . The absolute compression rate is then simply

$$r^{ABS}(n) = \frac{|\tau_n(c)|}{|\tau_{n_{\min}}(c)|},$$

where  $\tau_{n_{\min}}$  is the tokenizer with minimal vocabulary size  $n_{\min}$ . Note that the tokenization provided by  $\tau_{n_{\min}}$  can be different from character tokenization and highly depends on the behavior of the tokenizer. For example for the WordPiece implementation that we use<sup>2</sup> we obtain all unicode code

<sup>2</sup>We use <https://github.com/huggingface/tokenizers>

points twice in the vocabulary, once with a continuation symbol and once without.

**RCR**  $r^{REL}$ . One potential disadvantage of ACR is that it does not take into consideration how many different unicode code points are used. Thus we introduce the relative compression rate that considers in addition how “inflated” the vocabulary is.

$$r^{REL}(n) = \frac{|\tau_n(c)|}{|\tau_{n_{\min}}(c)|} \frac{n_{\min}}{n}.$$

One can also interpret the measure as a relative inverse type-token ratio

$$r^{REL}(n) = \frac{|\tau_n(c)|}{n} / \frac{|\tau_{n_{\min}}(c)|}{n_{\min}}.$$

The measures have convenient properties. For example they are monotonically decreasing and  $r^{ABS}(n_{\min}) = r^{REL}(n_{\min}) = 1.0$ . In the WordPiece implementation that we use the maximum tokenization length is limited by the number of whitespaces in the text. Thus  $r^{ABS}$  is always above a language specific threshold  $a$  and  $r^{ABS}(n_{\text{whitespace}}) = |\tau_{n_{\text{whitespace}}}(c)| / |\tau_{n_{\min}}(c)| =: a$ . For other tokenizers that treat whitespace as a normal unicode code point the whole text might be considered a single token, provided that the vocabulary size is sufficiently large, and thus  $r^{ABS}(n_{\max}) = 1 / |\tau_{n_{\min}}(c)| =: a$ . For the relative compression rate one can see that  $r^{REL}(\infty) = 0.0$ .

Figure 2 shows both compression rates for four languages as computed on the PBC.

### 3.3 Vocabulary Size Selection

In our study we want to compare different vocabulary sizes. However, it is not trivial to come up with meaningful vocabulary sizes. Uniformly chosen sizes are not interesting as increasing the vocabulary size from, e.g., 15,000 to 16,000 in English has only a marginal effect on the actual tokenization as it mostly affects rare words. Similarly, a vocabulary size of 1000 can be decent in English, but is most likely way too small for Chinese.

We thus try to model the function  $r^{ABS}$  for each language in order to be able to assess how important different vocabulary sizes are. To this end, we sample  $k$  vocabulary sizes uniformly from  $[n_{\min}, 100000]$  and obtain data pairs  $[(n_1, r^{ABS}(n_1)), \dots, (n_k, r^{ABS}(n_k))]$ . We observe that the compression rate follows an exponential pattern and thus we assume the following statistical model

$$y := r_{\theta}(x) = \alpha(x/n_{\min})^{\beta} + \gamma,$$

ENG	ELL	RUS	ZHO
129 <sub>1.0</sub>	143 <sub>1.0</sub>	137 <sub>1.0</sub>	6189 <sub>1.0</sub>
147 <sub>0.9</sub>	167 <sub>0.9</sub>	161 <sub>0.9</sub>	7054 <sub>0.9</sub>
172 <sub>0.8</sub>	201 <sub>0.8</sub>	195 <sub>0.8</sub>	8733 <sub>0.8</sub>
209 <sub>0.7</sub>	251 <sub>0.7</sub>	244 <sub>0.7</sub>	17083 <sub>0.7</sub>
267 <sub>0.6</sub>	330 <sub>0.6</sub>	323 <sub>0.6</sub>	33786 <sub>0.68</sub>
374 <sub>0.5</sub>	475 <sub>0.5</sub>	471 <sub>0.5</sub>	-
653 <sub>0.4</sub>	821 <sub>0.4</sub>	827 <sub>0.4</sub>	-
4139 <sub>0.3</sub>	2524 <sub>0.3</sub>	2632 <sub>0.3</sub>	-
18865 <sub>0.29</sub>	51077 <sub>0.24</sub>	59341 <sub>0.24</sub>	-

Table 1: Selected vocabulary sizes for each language. Subscript shows ACR.

where  $\theta = (\alpha, \beta, \gamma)$ . Based on the asymptotic behavior of our function we assume  $\gamma = a$  and  $\alpha = 1 - a$ . This gives us  $r_{\theta}(x) = (1 - a)(x/n_{\min})^{\beta} + a$  and the only parameter is  $\beta$ .

We fit  $\beta$  with a simple linear regression without intercept, i.e.,  $\log\left(\frac{(y-a)}{(1-a)}\right) = \beta \log\left(\frac{x}{n_{\min}}\right)$ . Overall we obtain an estimate  $\hat{\beta}$  and thus a model  $r_{\hat{\beta}}^{ABS}$ .

We can now sample interesting vocabulary sizes e.g., by sampling in areas where the rate of change  $(r_{\hat{\beta}}^{ABS})'$  is high. Alternatively, we go uniformly through possible compression rates: starting from a maximum compression rate we select values at 0.1-mark intervals (1.0, 0.9, ...) until we reach the minimum compression rate. Using  $(r_{\hat{\beta}}^{ABS})^{-1}$  we obtain vocabulary sizes from the chosen compression rates.

Table 1 shows selected vocabulary sizes for each language.

## 4 Experiments

### 4.1 Data

The Parallel Bible Corpus (PBC) by Mayer and Cysouw (2014) is a multi-parallel corpus spanning 1259 languages and up to 30k verses per translation. We chose this corpus for three reasons: **i)** For a clean experimental setup in our sentence retrieval task we require a multi-parallel corpus with additional paraphrases within the same language. This is possible in PBC as there are multiple translations for languages such as ENG, ELL or RUS.<sup>3</sup> **ii)** The overall corpus size is small and allows for efficient training of multilingual static and contextualized models. Dufter and Schütze (2020) provided evidence that experimenting on a small setup and verifying key results in a larger setup is feasible. **iii)** The corpus has a vast language coverage.

<sup>3</sup>Language codes follow the ISO 639-3 standard: <https://iso639-3.sil.org/>



We compute some general statistics (see Appendix Table 3) on PBC, calculating the maximum, average and minimum verse character length for each examined language, as well as the length covering 95% of all instances. This information was used in conjunction with computational considerations to set the maximum sequence length for our experiments.

#### 4.1.1 Language Selection

We perform our static embedding analysis on all 61 complete “newworld” editions in the PBC. These are very literal translations spanning the same set of verses across languages and thus ensure high comparability. For contextualized representations, experiments are more expensive. Thus we chose to focus on languages ENG, ELL, RUS and ZHO, which span different scripts, degrees of morphology and language families.

#### 4.1.2 Vocabulary Size Selection

Instead of hand-picking arbitrary sizes without guidance behind the selection process we follow our proposal from §3.3 and select sizes based on compression rates. Starting from a maximum compression rate (i.e., smallest vocabulary size) we choose sizes at 0.1-mark intervals (1.0, 0.9, ...) until we reach a minimum compression rate (i.e., maximum vocabulary size).

#### 4.1.3 Fake-English

We also experiment using Fake-English, denoting the English to Fake-English pair by *EngFake*. Fake-English is a concept examined in K et al. (2020), where English characters are converted to some separate, special characters such that the model processes English and Fake-English as two different languages. The “new” data that is created exhibits the same grammar and structure as English, but the script is different. This offers both a simple upper performance limit and a scrying glass into potential bugs or shortcomings of our approach. Note that to generate the Fake-English data, we used the same edition as the English Bible set, to ensure better comparability.

### 4.2 Static Embeddings

#### 4.2.1 Training

We evaluate vocabulary compatibility in static embedding spaces. To this end we tokenize comparable corpora using learned WordPiece vocabularies with different sizes. Subsequently we obtain

static embeddings using fastText (Bojanowski et al., 2017) with default parameters.

#### 4.2.2 Compatibility Evaluation

Comparing static embedding spaces at different granularities is challenging. Tasks like word translation or sentence retrieval disqualify, as obtaining word or sentence representations through mean pooling is too crude to get discriminative evaluation results. Training a neural network on top of the embeddings to evaluate multilinguality requires additional input, might obliterate the effect we want to investigate (i.e., the compatibility of the embedding spaces) and is not feasible when processing many language pairs. Therefore we employ a popular method to predict the cross-lingual performance using the spectral similarity of embedding spaces. Dubossarsky et al. (2020) introduced the measure *singular value gap* (SVG) and showed that it correlates well with cross-lingual downstream performance such as bilingual lexicon induction. Given two embedding spaces for two languages  $\mathbf{X}_l \in \mathbb{R}^{n_l \times d}$  and  $\mathbf{X}_e \in \mathbb{R}^{n_e \times d}$  it is computed as

$$SVG(\mathbf{X}_l, \mathbf{X}_e) := \sum_{i=1}^d \left( \log(\sigma_i^e) - \log(\sigma_i^l) \right)^2$$

with  $(\sigma_i^k)$  being the sorted singular values of  $\mathbf{X}_k$ . Instead of computing the sum across all dimensions we follow (Dubossarsky et al., 2020) and consider only the first 40 singular values. We found that experiments with the two alternative measures cond-hm and econd-hm that Dubossarsky et al. (2020) suggest yield similar results and thus focus only on SVG for the sake of clarity.

### 4.3 Contextualized Embeddings

#### 4.3.1 Pretraining

**Model Details.** For our experiments, we made use of BERT (Devlin et al., 2019). We performed two separate sets of experiments, one on PBC and another on the Wikipedia corpus. The Wikipedia experiments serve as a verification of the generalization of our findings to a larger scale.

A thorough and rigorous analysis on the effect of vocabulary sizes can be prohibitively costly on the larger scale. When comparing the vocabulary sizes of two languages with  $m$  and  $n$  different sizes each, we need  $m \times n$  different models. For example, to compare English ( $m = 9$ ) with Greek ( $n = 9$ ) we need to train  $9 \times 9 = 81$  models. Pretraining all

these models on Wikipedia with the default BERT parameters is prohibitively expensive.

Thus, we split our experiment into the small-scale Bible setup and the large-scale Wikipedia setup (§6). We also downsized the models for the Bible experiments. These models use a single attention head, which has been shown to have competitive performance compared to the multi-head approach with a substantial drop in required computational resources (Michel et al., 2019). We also downsized the hidden layer size, from 3072 to 256. We call this model BERT<sub>bible</sub>.

For the Wikipedia experiments, the default BERT hyperparameters were kept. We use BERT<sub>wiki</sub> to denote this model architecture.

**Bible Experiments.** For our work, we pre-trained the downsized BERT<sub>bible</sub> models on the PBC data. Then we evaluated on a bidirectional sentence retrieval task after finetuning following the SBERT example (Reimers and Gurevych, 2019).

We select the English (ENG), Greek (ELL), Russian (RUS) and Chinese (ZHO) versions of the Bible. Development and test sets were defined using common verses across the languages. The training set varies slightly for each language, since we included all verses not in the dev/test sets. Details on these splits can be found in Table 5.

For pretraining, we pair ENG with ELL/RUS/ZHO. We denote these pairs with *EngEll*, *EngRus* and *EngZho*. For each pair, we take the Cartesian product of the corresponding vocabulary sizes. The training sets for the two languages are then concatenated and two tokenizers are trained with the examined vocabulary sizes. Model pretraining then takes place for 75 epochs using BERT<sub>bible</sub>. Batch size is set to 128 and the maximum sequence length is also set to 128. We used ADAM (Kingma and Ba, 2015), with a learning rate of 2e-3, a weight decay of 0.01 and an epsilon value of 1e-6. Three runs were performed and the results averaged.

### 4.3.2 SBERT Finetuning

Each pretrained model is finetuned under the SBERT paradigm (Reimers and Gurevych, 2019). The model learns to identify whether two sentences are similar or not. Sentence representations are computed as usual and then pushed through a pooling layer. These representations are compared with cosine similarity, with the mean-squared-error loss used as the objective function.

The main motivation for this finetuning is to obtain meaningful sentence representations from different tokenization granularities. Overall, this allows us to analyze the multilinguality of the representations even though the tokens in both languages might be very different, such as words and characters. During finetuning, the model also learns how to separate dissimilar sentences and bring similar ones closer together in the representation space, which allows us to evaluate multilinguality using cross-lingual sentence retrieval.

The data used for this task comes from the PBC development sets. For each language we use two different editions, that is, two separate translations in the same language, i.e., paraphrases. Across the two editions, the verses with the same identifier are paired and labeled as similar, while for the dissimilar examples we pair each verse from one edition to a different, random verse in the other edition. We keep this pairing within each language, to ensure that our cross-lingual evaluation setup is zero-shot.

### 4.3.3 Compatibility Evaluation: Sentence Retrieval

To evaluate these pretrained and finetuned models, we use a cross-lingual sentence retrieval task. Data for this evaluation task comes from the development sets for English and the paired language. Given an English verse, we are tasked with retrieving the corresponding verse in the paired language. A mean-pooling layer takes as input the token representations of a sentence after a forward pass and outputs the final verse representation. Then, for each English sentence, we retrieve the 10 most similar sentences from the paired language, as calculated via cosine similarity. Our evaluation measure is precision@10. This method was repeated analogously in the other direction: for each sentence in the paired language, we retrieve the most similar sentences in English. Scores from both directions are averaged and results are analyzed in §5.2.2.

Note that the cross-lingual sentence retrieval is zero-shot in the sense that any multilinguality can only stem from the pretraining as no cross-lingual information was used during SBERT finetuning.

## 5 Results

### 5.1 Static Embedding Learning

We show results for static embeddings in Figure 3. For better visualization we interpolated the results




Figure 3: Figures showing the interpolated SVG of static embedding spaces for different language pairs computed on the PBC. Blue indicates high compatibility (i.e., low SVG) and red the contrary. Black stars indicate vocabulary sizes where  $r^{ABS}$  of both languages are equal. One can see different patterns: alphabetic languages tend to favor the diagonal. Logograms require much larger vocabulary sizes in English. Japanese uses a mixture of logograms and a syllabic writing system and thus exhibits a special pattern.

using `LinearTriInterpolator` in the Matplotlib library (Hunter, 2007). For alphabetic languages it is favorable to have comparable vocabulary sizes, which can be seen from the clear diagonal pattern of ELL, RUS and SPA. Chinese, being based on logograms, requires large vocabulary sizes in English. This makes intuitive sense: with small vocabulary sizes in English one arrives at a character tokenization and it is hard to imagine the equivalent of a Latin character or character n-gram in Chinese. Japanese exhibits a unique pattern, which can be explained by the fact that Japanese uses both syllables and logograms. Plots for all languages can be found in Appendix E.




Figure 4: Results showing the interpolated SVG of embedding spaces for different language pairs computed on Wikipedia.

## 5.2 Contextualized Embedding Learning

### 5.2.1 Fake-English Experiments

We begin our evaluation with the Fake-English experiments. As described in §4.1.3, we generate a modified “fake” English dataset from the English Bible corpus. These experiments act as our guide and debugging assistant. The behavior of these models can tell us a great deal about what to expect in the other experiments. Based on these results, hyperparameters were tuned both for pretraining and for finetuning. Results (on the test set) are shown in the rightmost plot of Figure 5. In the *EngFake* experiments, the diagonal is very prominent. Overall, performance in these experiments is quite high. On the diagonal, the precision@10 score is the highest in the bottom left corner, while it diminishes slightly towards the top right (detailed results can be found in Table 7 of the appendix). The fact that BERT is able to produce compatible representations between English and Fake-English at different tokenization granularities is in line with findings in (K et al., 2020).

### 5.2.2 Bible Experiments

As previously described in §4.3, we conducted experiments on PBC for the *EngEll*, *EngRus* and *EngZho* language pairs. Test results are shown in Figure 5 (with detailed tables in Appendix C).

*EngEll* and *EngRus* both show stronger performance along the diagonal. Performance is low with small sizes (both for English and the paired

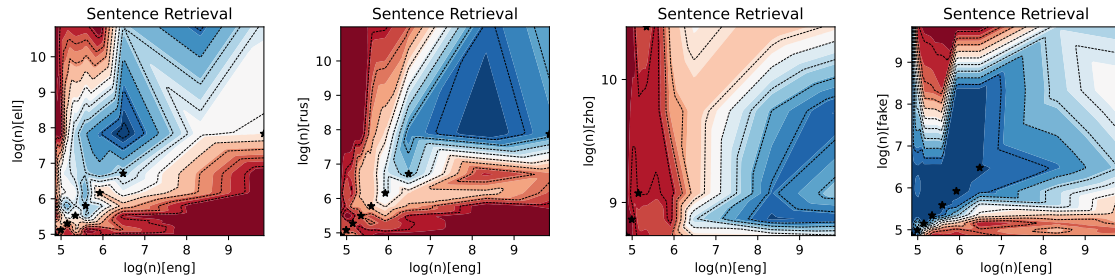


Figure 5: Figures showing the sentence retrieval test results from our Bible experiments. Blue indicates a high sentence retrieval score and red a low score. From left to right: *EngEll*, *EngRus*, *EngZho* and *EngFake*.

language). *EngEll* shows more consistently high precision across the diagonal, from small to large sizes. On the other hand, *EngRus* performance, while still better along the diagonal, is higher in the top right of the figure where both English and Russian vocabulary sizes are larger.

Despite these similar trends, performance in the *EngRus* experiment is higher than in *EngEll* (as indicated by the deep blue in the top right of the figure). Further, in *EngEll*, performance seems to maximize towards the top-middle section of Figure 5. For *EngZho*, instead of a diagonal we see a left-to-right pattern. The larger the English vocabulary size, the better performance is. With the smaller Chinese vocabulary sizes (bottom right of the figure), we see more consistently high results as well. Overall the results show similar trends to findings from the static embeddings in §5.1.

These results indicate language-specific patterns exist across tokenization sizes, something that can aid us in creating more compatible vocabularies.

## 6 Wikipedia Experiments

### 6.1 Data and Task

To verify whether our findings generalize to a large-scale setting, we also evaluate on XNLI (Conneau et al., 2018). The models we have previously pre-trained on the Bible corpus do not have the capacity to tackle this larger task, therefore we opted to instead pretrain new models on Wikipedia data. To this end, English, Greek, Russian and Chinese Wikipedia dumps were used. Due to a lack of computational resources, we kept only 7.4GB of English data, 3.3GB of Chinese and 6.3GB of Russian. For the Greek experiments, all articles were kept resulting in around 0.8GB of data in total. For all other languages, articles were randomly sampled. Details on data size can be found in Table 6 of the appendix.

Statistics for XNLI are shown in Table 4 of the appendix. Large differences were observed between the two text fields (premise and hypothesis), as well as between Chinese and the other languages (attributed to the difference in scripts).

### 6.2 Training

To accommodate the larger dataset, BERT<sub>wiki</sub> was used. The number of epochs was set to 1 and the maximum sequence length was reduced to 128 due to computational restrictions. Analogously to the small-scale Bible experiments, for pretraining, English data was concatenated with the paired language data, with separate tokenizers trained on the two corpora.

Because evaluating for all possible combinations of size pairs is prohibitively resource-consuming, we settled on three size pairs for each language. These pairs were selected according to the Bible evaluation we ran beforehand. Namely, for each language, we picked the best, the one at the 75th percentile (*perc\_75*) and the worst pair sizes to compare according to their sentence retrieval precision@10 score.

After the pretraining of these models on Wikipedia, we finetuned them on the English XNLI training set and then performed cross-lingual zero-shot evaluation on the other languages. As before, Greek, Russian and Chinese were evaluated.

### 6.3 Evaluation and Results

Figure 4 shows the same experiments as described in §5.1, this time conducted on the Wikipedia dataset. The key trends are observable here as well.

In Table 2 we show the accuracy of the three examined pairs (as defined in §6.2: best, *perc\_75*, worst) for each language in the XNLI task. Note that a majority baseline has an accuracy of 1/3. The pairs are sorted according to their performance in the Bible experiments (precision@10 on sentence



<i>EngEll</i>	Acc.
14098 <sub>0.4</sub> - 19789 <sub>0.3</sub>	48.8
7414 <sub>0.9</sub> - 12539 <sub>0.4</sub>	46.4
22158 <sub>0.3</sub> - 6177 <sub>0.9</sub>	44.6
<i>EngRus</i>	Acc.
22158 <sub>0.3</sub> - 15892 <sub>0.3</sub>	59.1
9752 <sub>0.6</sub> - 15892 <sub>0.3</sub>	58.3
6956 <sub>1.0</sub> - 15892 <sub>0.3</sub>	56.4
<i>EngZho</i>	Acc.
22158 <sub>0.3</sub> - 43005 <sub>0.9</sub>	45.1
14098 <sub>0.4</sub> - 39393 <sub>1.0</sub>	44.7
6956 <sub>1.0</sub> - 81426 <sub>0.7</sub>	44.8

Table 2: Zero-shot XNLI accuracy for English (source) and three target languages. We give the vocabulary sizes and, as a subscript, Absolute Compression Rate (ACR). In each block of three results (one each for ELL, RUS, ZHO), the pair of vocabulary sizes with the best performance on sentence retrieval is at the top, followed by *perc\_75* and the worst performing pair. Wikipedia results are a good predictor of XNLI performance for Greek and Russian. See text for more details.

retrieval, largest at the top). For all language pairs, performance on the English XNLI test set was similar, slightly fluctuating ( $\pm 0.75$ ) around 74%. These scores are hence omitted from Table 2 since they do not add much to the discussion.

The best combinations (from the Bible experiments) for *EngEll* and *EngRus* have the best XNLI performance as well, followed by *perc\_75* and finally the worst-performing combination. Thus, our findings for *EngEll* and *EngRus* are corroborated: the best performances are found when the absolute compression rate is similar.

For *EngZho*, we see little variation between the three chosen size combinations although the best-performing combination is still at the top. Maybe this is due to the fact that Chinese behaves differently from the two “alphabetic” languages and therefore a more extensive hyperparameter optimization would be required. We leave this question for future research.

#### 6.4 Correlation Analysis

One major objective is to be able to compare the compatibility of tokenizations across languages. To this end we compute  $\log(r_e^{ABS})/\log(r_f^{ABS})$  for two languages  $e, f$  and analogously for  $r^{REL}$ . We compute the correlation of this compatibility measure using the original mBERT tokenizer with the downstream performance of mBERT on XNLI as

reported by (Hu et al., 2020). More details can be found in the supplementary.

The Pearson correlation is .40 and .34, respectively. This indicates that our measures capture the compatibility of tokenizations and correlate with zero-shot transfer downstream performance. The strength of the correlation is similar to what Rust et al. (2021) find.

## 7 Conclusion

We investigated tokenization compatibility across languages, both for static and contextualized embeddings. To this end, we proposed compression rates and a method to select meaningful vocabulary sizes in an automated manner for any language. We introduced a compatibility measure and showed that it correlates with downstream performance.

We gave evidence for two key findings that hold for both static and contextualized embeddings. **i)** Tokenization compatibility can have a significant impact on multilingual performance: performance is generally higher when vocabulary sizes are compatible. **ii)** Tokenization compatibility varies significantly across language pairs: pairs of alphabetic languages show a stronger performance on the diagonal (i.e., for comparable vocabulary sizes) while an alphabetic language like English requires large vocabulary sizes when paired with a logographic language like Chinese.

Our study has clear limitations: we only experimented with the WordPiece tokenizer and with a small number of language pairs and there are only a few larger-scale experiments on XNLI. We will continue research in this direction and hope that this paper will spark interest by other researchers in the compatibility of tokenizations across languages and its effect on multilinguality.

## Acknowledgements

This work was supported by the European Research Council (# 740516). The second author was supported by the Bavarian research institute for digital transformation (bidt) through their fellowship program. We thank Ehsaneddin Asgari and the anonymous reviewers for valuable comments and fruitful discussions.

## References

Gustavo Aguilar, Bryan McCann, Tong Niu, Nazneen Fatema Rajani, Nitish Shirish Keskar, and

- Thamar Solorio. 2020. [Char2subword: Extending the subword embedding space from pre-trained models using robust character compositionality](#). *CoRR*, abs/2010.12730.
- Antonios Anastasopoulos and Graham Neubig. 2020. [Should all cross-lingual embeddings speak English?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8658–8679, Online. Association for Computational Linguistics.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020a. [On the cross-lingual transferability of monolingual representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4623–4637, Online. Association for Computational Linguistics.
- Mikel Artetxe, Sebastian Ruder, Dani Yogatama, Gorka Labaka, and Eneko Agirre. 2020b. [A call for more rigor in unsupervised cross-lingual learning](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7375–7388, Online. Association for Computational Linguistics.
- Ehsaneddin Asgari, Masoud Jalili Sabet, Philipp Dufter, Christoph Ringlstetter, and Hinrich Schütze. 2020. [Subword sampling for low resource word alignment](#). *CoRR*, abs/2012.11657.
- 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.
- Kaj Bostrom and Greg Durrett. 2020. [Byte pair encoding is suboptimal for language model pretraining](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4617–4624, Online. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020a. [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 *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel R. Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2475–2485. Association for Computational Linguistics.
- Alexis Conneau, Shijie Wu, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2020b. [Emerging cross-lingual structure in pretrained language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6022–6034, Online. Association for Computational Linguistics.
- 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, Minneapolis, Minnesota. Association for Computational Linguistics.
- Haim Dubossarsky, Ivan Vulić, Roi Reichart, and Anna Korhonen. 2020. [The secret is in the spectra: Predicting cross-lingual task performance with spectral similarity measures](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2377–2390, Online. Association for Computational Linguistics.
- Philipp Dufter and Hinrich Schütze. 2020. [Identifying elements essential for BERT’s multilinguality](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4423–4437, Online. Association for Computational Linguistics.
- Philip Gage. 1994. A new algorithm for data compression. *The C Users Journal*, 12(2):23–38.
- Thamme Gowda and Jonathan May. 2020. [Finding the optimal vocabulary size for neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3955–3964, Online. Association for Computational Linguistics.
- Benjamin Heinzerling and Michael Strube. 2018. [BPEmb: Tokenization-free pre-trained subword embeddings in 275 languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421. PMLR.

- Haoyang Huang, Yaobo Liang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, and Ming Zhou. 2019. [Unicoder: A universal language encoder by pre-training with multiple cross-lingual tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2485–2494, Hong Kong, China. Association for Computational Linguistics.
- J. D. Hunter. 2007. [Matplotlib: A 2d graphics environment](#). *Computing in Science & Engineering*, 9(3):90–95.
- Karthikeyan K, Zihan Wang, Stephen Mayhew, and Dan Roth. 2020. [Cross-lingual ability of multilingual BERT: an empirical study](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Taku Kudo. 2018. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia. 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.
- Thomas Mayer and Michael Cysouw. 2014. [Creating a massively parallel Bible corpus](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 3158–3163, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Paul Michel, Omer Levy, and Graham Neubig. 2019. [Are sixteen heads really better than one?](#) In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Farhad Nooralahzadeh, Giannis Bekoulis, Johannes Bjerva, and Isabelle Augenstein. 2020. [Zero-shot cross-lingual transfer with meta learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4547–4562, Online. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Ivan Provilkov, Dmitrii Emelianenko, and Elena Voita. 2020. [BPE-dropout: Simple and effective subword regularization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1882–1892, Online. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Phillip Rust, Jonas Pfeiffer, Ivan Vulić, Sebastian Ruder, and Iryna Gurevych. 2021. [How good is your tokenizer? on the monolingual performance of multilingual language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3118–3135, 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 2012, Kyoto, Japan, March 25-30, 2012*, pages 5149–5152. IEEE.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Jasdeep Singh, Bryan McCann, Richard Socher, and Caiming Xiong. 2019. [BERT is not an interlingua and the bias of tokenization](#). In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 47–55, Hong Kong, China. Association for Computational Linguistics.
- Junqiu Wei, Qun Liu, Yinpeng Guo, and Xin Jiang. 2021. [Training multilingual pre-trained language model with byte-level subwords](#). *CoRR*, abs/2101.09469.

Shijie Wu and Mark Dredze. 2019. **Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT.** In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 833–844, Hong Kong, China. Association for Computational Linguistics.



## A Data Statistics

Here we present details for the PBC and Wikipedia data.

PBC	Max.	Avg.	Min.	perc_95
ENG	474	138	9	249
ELL	523	150	3	272
RUS	398	122	3	221
ZHO	264	62	8	113

Table 3: Character length statistics for the PBC data.

XNLI	Max.		Avg.		Min.		perc_95	
ENG	1815	393	113	56	5	1	252	101
ELL	350	229	120	59	13	9	225	106
RUS	309	195	111	55	11	6	210	100
ZHO	112	87	33	16	4	3	62	29

Table 4: Character length statistics for the XNLI data. We report the lengths of the premise and the hypothesis) separately.

Language	Train	Dev	Test
ENG	23,633	5,000	2,500
ELL	23,673	5,000	2,500
RUS	23,673	5,000	2,500
ZHO	23,657	5,000	2,500

Table 5: Number of verses for each language and split combination in the PBC.

Language	Size
ENG	7436MB
ELL	809MB
RUS	6291MB
ZHO	3252MB

Table 6: Sizes for the languages in our Wikipedia data.

## B Computational Details

For each individual PBC experiment, pretraining took around 2 hours and finetuning/evaluation 25 minutes. For each individual Wikipedia experiment, pretraining took around 145 hours and finetuning/evaluation 5 hours. A multi-GPU server was used, with GeForce GTX 1080Ti devices.

## C Bible Experiments Detailed Results

Here we present the detailed tables for the PBC results. Three runs were made with results averaged. Standard deviations were small, so they were omitted for readability.

	129	147	172	209	267	374	653	4139	18865
129	1.00	0.99	0.99	0.99	0.99	0.99	1.00	1.00	1.00
147	0.96	0.99	0.68	0.97	0.98	0.93	0.97	0.61	0.11
172	0.28	0.53	1.00	0.99	0.99	0.99	1.00	0.22	0.06
209	0.11	0.26	0.71	0.99	1.00	1.00	1.00	0.21	0.05
267	0.15	0.20	0.44	0.94	1.00	0.99	1.00	0.16	0.05
374	0.13	0.17	0.37	0.72	0.94	0.98	1.00	0.99	0.03
653	0.18	0.22	0.10	0.47	0.57	0.94	0.97	0.99	0.03
4139	0.26	0.35	0.22	0.39	0.66	0.66	0.89	0.67	0.65
18865	0.28	0.21	0.32	0.36	0.37	0.51	0.69	0.49	0.68

Table 7: Similarity Matrix for *EngFake*. Rows denote English and columns Fake-English.

	143	167	201	251	330	475	821	2524	51077
129	0.00	0.03	0.09	0.01	0.01	0.06	0.04	0.02	0.00
147	0.05	0.09	0.11	0.09	0.16	0.12	0.14	0.00	0.03
172	0.04	0.02	0.07	0.14	0.20	0.13	0.14	0.12	0.04
209	0.05	0.10	0.10	0.12	0.15	0.13	0.16	0.17	0.03
267	0.03	0.01	0.09	0.18	0.17	0.19	0.18	0.20	0.03
374	0.03	0.02	0.08	0.11	0.16	0.09	0.19	0.22	0.00
653	0.00	0.05	0.04	0.02	0.09	0.14	0.17	0.26	0.18
4139	0.00	0.00	0.00	0.01	0.00	0.07	0.07	0.13	0.23
18865	0.01	0.00	0.00	0.00	0.00	0.01	0.00	0.12	0.14

Table 8: Similarity Matrix for *EngEll*. Rows denote English and columns Greek.

	137	161	195	244	323	471	827	2632	59341
129	0.09	0.09	0.07	0.06	0.01	0.03	0.02	0.02	0.01
147	0.12	0.08	0.07	0.01	0.06	0.10	0.06	0.00	0.00
172	0.14	0.08	0.07	0.03	0.13	0.04	0.07	0.00	0.00
209	0.06	0.03	0.06	0.10	0.13	0.16	0.17	0.00	0.01
267	0.03	0.05	0.11	0.08	0.17	0.16	0.16	0.13	0.04
374	0.01	0.06	0.08	0.15	0.20	0.26	0.26	0.19	0.05
653	0.00	0.05	0.05	0.04	0.14	0.18	0.33	0.37	0.09
4139	0.00	0.00	0.00	0.00	0.00	0.17	0.05	0.44	0.41
18865	0.00	0.00	0.00	0.00	0.04	0.09	0.11	0.37	0.33

Table 9: Similarity Matrix for *EngRuS*. Rows denote English and columns Russian.

	6189	7054	8733	17083	33786
129	0.00	0.01	0.01	0.00	0.01
147	0.02	0.02	0.05	0.03	0.01
172	0.05	0.06	0.04	0.03	0.07
209	0.04	0.06	0.04	0.03	0.07
267	0.06	0.07	0.05	0.03	0.03
374	0.07	0.03	0.03	0.11	0.13
653	0.16	0.24	0.13	0.13	0.18
4139	0.17	0.36	0.33	0.28	0.13
18865	0.19	0.34	0.26	0.35	0.17

Table 10: Similarity Matrix for *EngZho*. Rows denote English and columns Chinese.

## **D Correlation Analysis**

We show detailed results for the correlation analysis in Table 11.

## **E SVG Plots**

In the next pages, we show the complete SVG plots for all 61 languages.

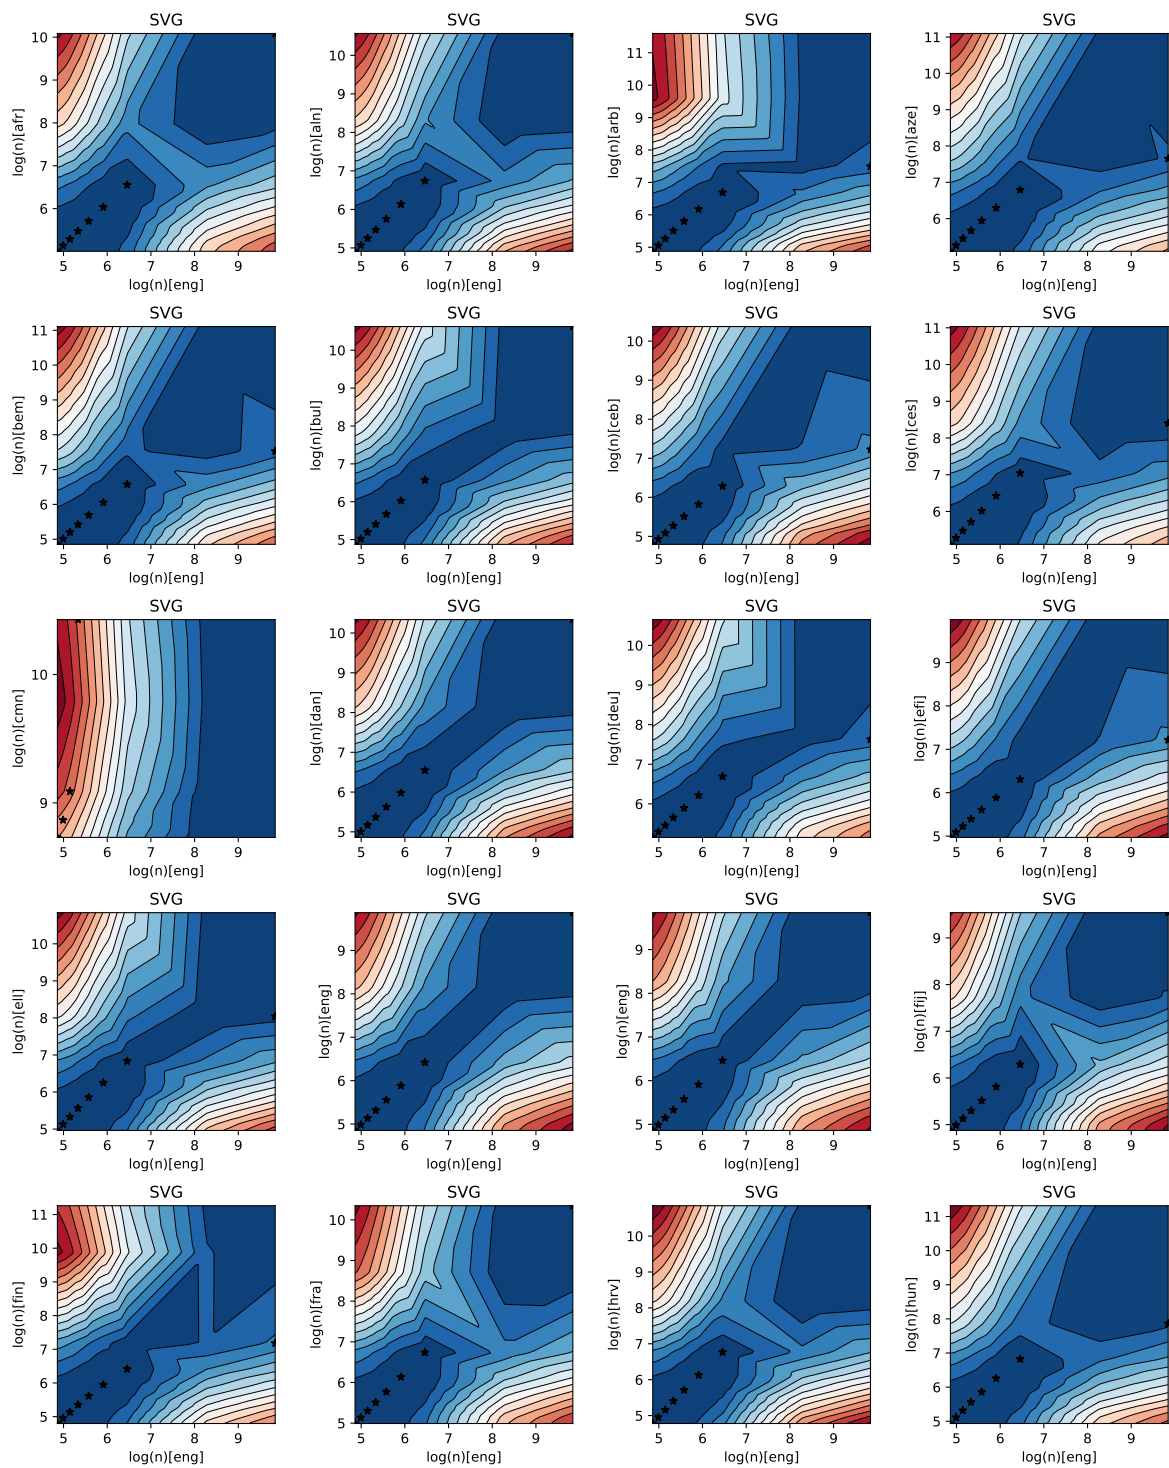


Figure 6: Figures showing the interpolated SVG for different language pairs as computed on the PBC (1/4).




Figure 7: Figures showing the interpolated SVG for different language pairs as computed on the PBC (2/4).



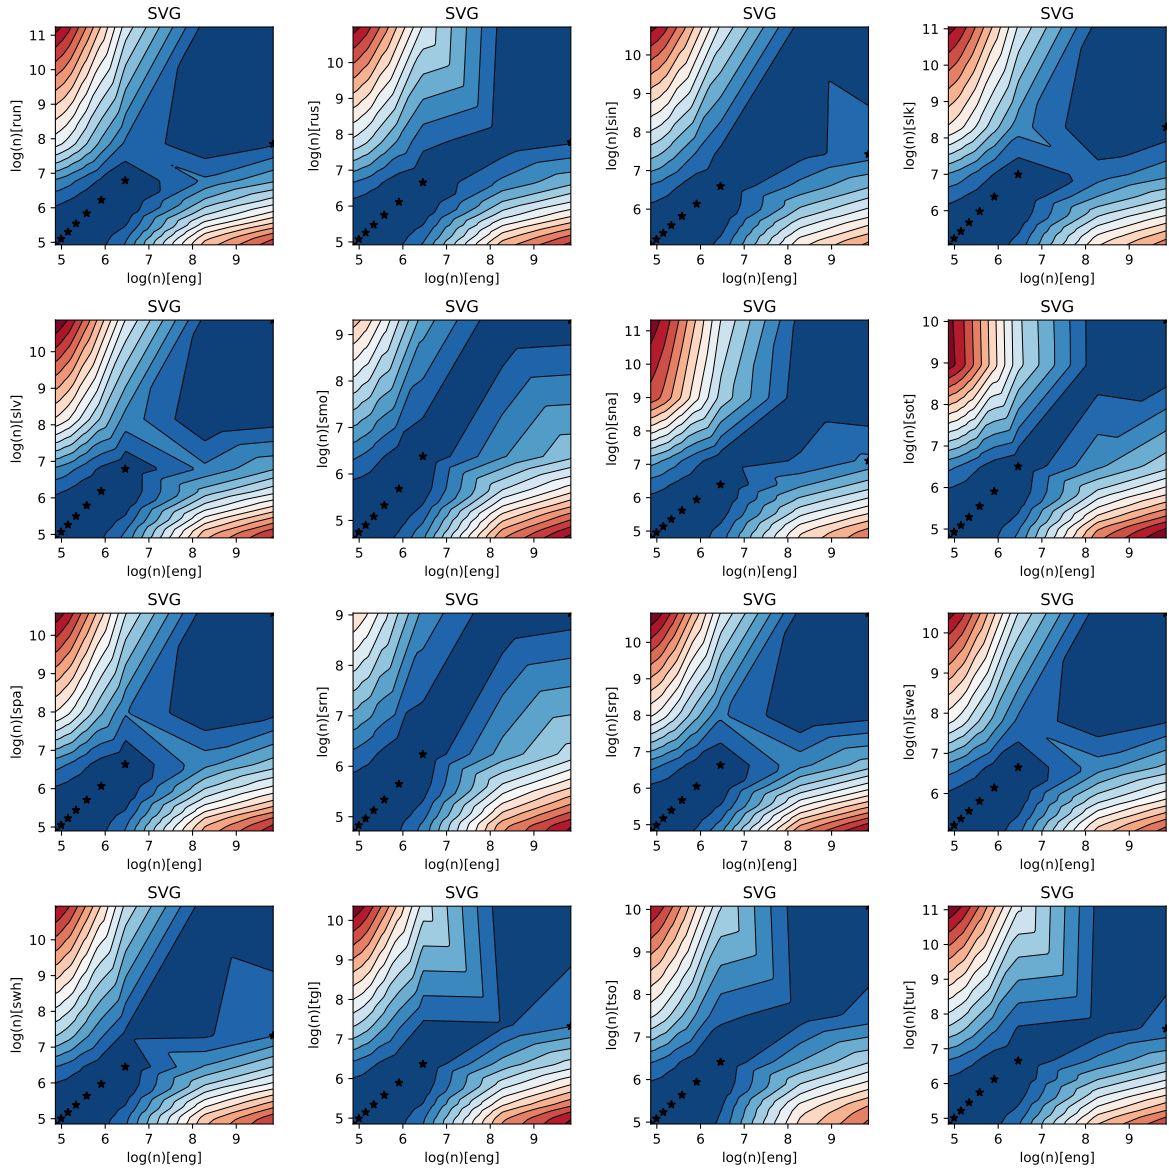


Figure 8: Figures showing the interpolated SVG for different language pairs as computed on the PBC (3/4).

Language	XNLI Acc.	$r^{ABS}$	$r^{REL}$	$\log(r_x^{ABS})/\log(r_{eng}^{ABS})$	$\log(r_x^{REL})/\log(r_{eng}^{REL})$
ara	64.9	0.46	0.06	0.64	0.68
bul	68.9	0.39	0.03	0.77	0.85
deu	71.1	0.31	0.01	0.95	1.06
ell	66.4	0.57	0.06	0.46	0.65
eng	81.4	0.30	0.01	1.00	1.00
esp	74.3	0.32	0.02	0.95	0.93
fra	73.8	0.33	0.02	0.91	0.95
hin	60.0	0.53	0.08	0.53	0.61
rus	69.0	0.37	0.02	0.82	0.90
swa	50.4	0.39	0.03	0.78	0.84
tha	55.8	0.68	0.12	0.31	0.49
tur	61.6	0.37	0.03	0.81	0.81
urd	58.0	0.48	0.07	0.60	0.63
vie	69.5	0.45	0.08	0.65	0.60
zho	69.3	0.96	0.69	0.04	0.09
Pearson Correlation				0.40	0.34

Table 11: Compression rates computed using the mBERT tokenizer. XNLI results by [Hu et al. \(2020\)](#).




Figure 9: Figures showing the interpolated SVG for different language pairs as computed on the PBC (4/4).