

Allocating Large Vocabulary Capacity for Cross-lingual Language Model Pre-training

Bo Zheng^{†*}, Li Dong[‡], Shaohan Huang[‡],
Saksham Singhal[‡], Wanxiang Che[†], Ting Liu[†], Xia Song[‡], Furu Wei[‡]

[†]Harbin Institute of Technology

[‡]Microsoft Corporation

{bzheng, car, tliu}@ir.hit.edu.cn

{lidong1, shaohanh, saksingh, xiaso, fuwei}@microsoft.com

Abstract

Compared to monolingual models, cross-lingual models usually require a more expressive vocabulary to represent all languages adequately. We find that many languages are under-represented in recent cross-lingual language models due to the limited vocabulary capacity. To this end, we propose an algorithm VOCAP to determine the desired vocabulary capacity of each language. However, increasing the vocabulary size significantly slows down the pre-training speed. In order to address the issues, we propose k -NN-based target sampling to accelerate the expensive softmax. Our experiments show that the multilingual vocabulary learned with VOCAP benefits cross-lingual language model pre-training. Moreover, k -NN-based target sampling mitigates the side-effects of increasing the vocabulary size while achieving comparable performance and faster pre-training speed. The code and the pretrained multilingual vocabularies are available at <https://github.com/bozheng-hit/VoCapXLM>.

1 Introduction

Pretrained cross-lingual language models (Conneau and Lample, 2019; Conneau et al., 2020; Chi et al., 2021b; Xue et al., 2020) have recently shown great success in improving cross-lingual transferability. These models encode texts from different languages into universal representations with a shared multilingual vocabulary and a shared Transformer encoder (Vaswani et al., 2017). By pre-training cross-lingual language models on the large-scale multilingual corpus, the models achieve state-of-the-art performance on various downstream tasks, e.g., cross-lingual question answering and cross-lingual sentence classification.

Although the Transformer architecture used in most pretrained monolingual and cross-lingual language models are almost identical, the vocabularies

are quite different. The vocabulary sizes in existing pretrained monolingual language models typically range from 30K to 60K subword units (Devlin et al., 2019; Liu et al., 2019; Dong et al., 2019; Bao et al., 2020). Meanwhile, state-of-the-art pretrained cross-lingual language models use the shared multilingual vocabulary of 250K subword units to represent more than 100 languages (Conneau et al., 2020; Chi et al., 2021b; Xue et al., 2020). Although some subword units are shared across languages, no more than 2.5K language-specific subword units on average are allocated for each language, which is still relatively small. Besides, the multilingual vocabulary is trained on the combined multilingual corpus with subword segmentation algorithms like BPE (Sennrich et al., 2015) and unigram language model (Kudo, 2018). During vocabulary construction, these algorithms tend to select more subword units shared across languages with common scripts like Latin and Cyrillic (Chung et al., 2020b), but have a lower chance to select language-specific subword units. It is hard to determine how much vocabulary capacity a particular language requires and whether the shared multilingual vocabulary has allocated enough vocabulary capacity to represent the language.

In this paper, we propose VOCAP, an algorithm to allocate large vocabulary for cross-lingual language model by separately evaluating the required vocabulary capacity of each language. First, we use the *average log probability* (ALP) to evaluate the ability of a vocabulary to represent a particular language. We find that ALP is highly correlated to the downstream task performance, and we use it as an indicator to allocate language-specific vocabulary capacity. In addition, the language-specific pre-training corpus size should also be considered since the pretrained model can only learn limited knowledge from low-resource languages where the pre-training data is scarce. Therefore, allocating too much vocabulary capacity for low-resource lan-

*Contribution during internship at Microsoft Research.

guages is inefficient. VOCAP leverages both ALP and pre-training corpus size to evaluate the required vocabulary capacity of each language. We finally allocate a multilingual vocabulary with 500K subword units with VOCAP and show it can significantly improve the model performance.

However, increasing the vocabulary size has two practical drawbacks: slow pre-training speed and heavy model size. To address the pre-training speed issue, we propose k -NN-based target sampling, an approximate algorithm to improve the computing efficiency in the expensive softmax caused by the large vocabulary. We pre-train the model with a small subset of the entire vocabulary constructed with k nearest neighbors of the target words in current mini-batch data, evaluated with the inner product of subword embeddings. As for the model size, we halve the embedding dimension and draw a different conclusion from [Conneau et al. \(2020\)](#) that increasing vocabulary from 250K to 500K with a fixed capacity model can also improve the performance.

Our contributions are summarized as follows:

- We propose VOCAP, an algorithm to allocate appropriate vocabulary capacity for each language in the shared multilingual vocabulary of cross-lingual language models.
- We propose k -NN-based target sampling, a softmax approximation algorithm to improve the computing efficiency during cross-lingual language model pre-training.
- We evaluate our methods on the XTREME benchmark ([Hu et al., 2020](#)), including three different tasks on seven datasets. Experiments show that VOCAP consistently outperforms previous vocabulary construction methods. Meanwhile, our k -NN-based target sampling enables effective acceleration while achieving comparable performance.

2 VOCAP: Language-Specific Vocabulary Capacity Allocation

We attribute the main factors that affect the performance of a particular language in a cross-lingual language model to language-specific pre-training corpus size and vocabulary capacity. While previous work adjusts pre-training corpus size with an exponentially smoothed sampling distribution ([Conneau and Lample, 2019](#); [Conneau et al.,](#)

[2020](#)), few existing works have explored the effect of the language-specific vocabulary capacity in pretrained cross-lingual language models.

In this section, we first investigate the correlation between the language-specific vocabulary capacity and downstream task performance through experiments. Then we introduce our proposed multilingual vocabulary allocation algorithm VOCAP.

2.1 Investigating Language-Specific Vocabulary Capacity

We start by introducing *average log probability* (ALP) to quantify the language-specific vocabulary capacity in the shared multilingual vocabulary for a specific language.¹ Given a monolingual corpus composed of sentences $\mathcal{D}_i = \{s_1, \dots, s_{|\mathcal{D}_i|}\}$ from the i -th language and tokenized with vocabulary V , the average log probability is defined as follows:

$$\text{ALP}(\mathcal{D}_i, V) = \frac{1}{|\mathcal{D}_i|} \sum_{j=1}^{|\mathcal{D}_i|} \sum_{k=1}^{|s_j|} \log p_{uni}(s_j^k) \quad (1)$$

where s_j^k is the k -th subword of the sentence s_j , and $p_{uni}(\cdot)$ is the unigram distribution counted on the monolingual corpus \mathcal{D}_i . It is difficult to count the language-specific subword units in multilingual vocabularies since the raw text contains a lot of code-switched data. By contrast, ALP is a more convenient indicator of language-specific vocabulary capacity and it is penalized by the subword units with low-frequency.

To investigate the impact of language-specific vocabulary capacity, we first learn monolingual vocabularies in different sizes to obtain vocabularies with different ALP, i.e., language-specific vocabulary capacity. Then we conduct pre-training with these monolingual vocabularies on their corresponding monolingual corpora. Finally, we evaluate these monolingual models on downstream tasks and study the correlation between language-specific vocabulary capacity and downstream task performance.

2.1.1 Setup

To alleviate the bias from the languages' characteristics, we first select four languages with different pre-training corpus sizes from different language families, which are Hindi (hi), Persian (fa), Italian (it), Russian (ru). We first learn thirty monolingual

¹For brevity and consistency, we refer to the parameterized tokenizer also as vocabulary, e.g., SentencePiece model.




Figure 1: ALP of different monolingual vocabularies with different vocabulary sizes.




Figure 2: F1 score on POS task with different vocabularies versus their ALP on the monolingual corpus.




Figure 3: F1 score on NER task with different vocabularies versus their ALP on the monolingual corpus.




Figure 4: Comparison of vocabulary capacity of different-resourced languages. Shorter bars indicate larger vocabulary capacity.

vocabularies for each language on the corresponding monolingual corpus, with vocabulary size ranging from 1K to 30K. Then we pretrain monolingual language models with the corresponding monolingual vocabularies. We evaluate these pretrained models on two downstream tasks: NER (Pan et al., 2017) and POS (Zeman et al., 2019) from the XTREME benchmark since there is annotated task data for a large number of languages. The vocabularies are learned on the reconstructed CommonCrawl corpus (Chi et al., 2021b; Conneau et al., 2020) using SentencePiece (Kudo and Richardson, 2018) with the unigram language model (Kudo, 2018). The unigram distributions are also counted on the CommonCrawl corpus. The Wikipedia corpus is used for all pre-training experiments in this paper since it is easier to run experiments due to its smaller size. More details about the pre-training data can be found in the appendix.

2.1.2 Observations

Increasing vocabulary size affects ALP of different languages in varying degrees. In Fig-

ure 1, we show the correlation between vocabulary size and ALP of four different languages. We observe the ALP varies across different languages, mainly because ALP correlates with the lexicon granularity of the language, i.e., the average number of tokens per sentence. Besides, when the vocabulary size is larger than 10,000, the gains of increasing monolingual vocabulary size in hi and fa are less than it and ru. We attribute it to that hi and fa does not have extensive compoundings. Another observation is that for each language, every time we increase the vocabulary size by 1K, the increment in ALP is monotonically decreasing.

ALP correlates positively with downstream task performance. In Figure 2 and Figure 3, we

illustrate downstream task performance of models pretrained with monolingual vocabularies on corresponding monolingual corpora. We observe that ALP correlates positively with downstream task performance, making language-specific ALP a valid indicator to allocate multilingual vocabulary. Another natural option to allocate multilingual vo-

Algorithm 1 Allocating Multilingual Vocabulary with VOCAP

Input: size of target multilingual vocabulary T ; monolingual vocabularies of N languages $\{V_{t_i}^i\}_{i=1}^N$; monolingual corpus of N languages $\{D_i\}_{i=1}^N$

Output: multilingual vocabulary V

```
1: for  $i \leftarrow 1$  to  $N$  do
2:   for  $j \leftarrow 1$  to 50 do
3:      $a_{i,j \times 1000} \leftarrow \text{ALP}(D_i, V_{j \times 1000}^i)$ 
4:    $t_i \leftarrow 0$ 
5:    $a_{i,0} \leftarrow -\infty$ 
6: do
7:    $j \leftarrow 0$ 
8:    $\delta \leftarrow 0$ 
9:   for  $i \leftarrow 1$  to  $N$  do
10:    if  $\delta < a_{i,t_i+1000} - a_{i,t_i}$  then
11:       $\delta \leftarrow a_{i,t_i+1000} - a_{i,t_i}$ 
12:       $j \leftarrow i$ 
13:    $t_j \leftarrow t_j + 1000$ 
14:    $V \leftarrow |\bigcup_{i=1}^N V_{t_i}^i|$ 
15: while  $|V| < T$ 
16: if  $|V| > T$  then
17:   Clip the size of  $V$  to  $T$ 
```

cabulary is directly using monolingual vocabulary size to indicate language-specific vocabulary capacity. We compare ALP against vocabulary size and observe that ALP correlates better than vocabulary size with the downstream task performance. Besides, ALP reflects the language-specific characteristics, while vocabulary size does not. The detailed comparison is shown in the appendix.

2.2 Allocating Multilingual Vocabulary with VOCAP

Based on the observations in Section 2.1.2, we first give the implementation of our proposed vocabulary allocation algorithm VOCAP. Then we compare the multilingual vocabulary learned with VOCAP and directly learned with SentencePiece on the multilingual corpus.

2.2.1 VOCAP Implementation

We formulate the vocabulary construction of VOCAP as the problem of finding the optimal way to allocate language-specific vocabulary size to each language, such that the overall ALP of all languages is maximized. In addition to language-specific vocabulary capacity measured with ALP from Equation (1), the language-specific pre-training corpus size also affects the downstream task performance. Considering the two factors, the

procedure of VOCAP can be formulated as follows:

$$\operatorname{argmax}_{t_1, \dots, t_N} \sum_{i=1}^N q_i^\beta \text{ALP}(D_i, V_{t_i}^i) \quad \text{s.t.} \quad \left| \bigcup_{i=1}^N V_{t_i}^i \right| = T \quad (2)$$

where $t_i \in \{x \times 1000 \mid x \leq 50, x \in \mathbb{N}^+\}$ is the number of subword units allocated to the i -th language,² β is a rescaling factor, $V_{t_i}^i$ is the vocabulary of the i -th language with t_i subword units, T is the size of the target multilingual vocabulary, and q_i is the probability of sampling training instances from i -th language during pre-training (Conneau and Lample, 2019; Conneau et al., 2020):

$$q_i = \frac{f_i^\alpha}{\sum_{j=1}^N f_j^\alpha} \quad \text{with} \quad f_i = \frac{n_i}{\sum_{k=1}^N n_k} \quad (3)$$

where n_i is the number of instances in the i -th language, α is a rescaling factor used to alleviate the bias towards high-resource languages. Since the increment in ALP when increasing the vocabulary size by a certain number is monotonically decreasing, Equation (2) can be solved with the greedy algorithm in Algorithm 1.

2.2.2 Intrinsic Analysis

We compare the multilingual vocabulary learned with VOCAP and directly learned with SentencePiece on the multilingual corpus. The multilingual corpus to learn vocabularies in this paper is the concatenation of sentences sampled randomly from the monolingual corpora. Sentences from the i -th language is sampled with probability q_i from Equation (3) and use $\alpha = 0.7$. We filter languages with corpus size larger than 0.1 GB, resulting in 86 languages.

We evaluate the multilingual vocabularies with their ALP on each language’s monolingual corpus, and show results of different-resourced languages in Figure 4. We refer to languages with less than 1GB and more than 10GB pre-training corpus in the reconstructed CommonCrawl as low-resource and high-resource languages, respectively, otherwise mid-resource languages. When directly learning vocabulary on the multilingual corpus using SentencePiece, the vocabulary with 500K subword units (JOINT_{500K}) only has a negligible improvement compared to the vocabulary with 250K subword units (JOINT_{250K}). Meanwhile, our method

²Since the cost of learning monolingual vocabularies with arbitrary sizes and getting the corresponding ALP is unaffordable, we learn monolingual vocabularies with vocabulary size range from 1K to 50K at intervals of 1K.

(VOCAP_{500K}) consistently outperforms JOINT_{500K} in different-resourced languages, especially in mid and low-resource languages. The statistics of the allocated vocabulary size for each language in VOCAP_{500K} are shown in the appendix.

3 Accelerate Large-Vocabulary Language Model Pre-Training

Although extending the multilingual vocabulary benefits cross-lingual language models, pre-training with such large vocabularies brings two practical issues: slow pre-training speed and heavy model size. To tackle the issues, we first introduce our k -NN-based target sampling in Section 3.1, which is a softmax approximation algorithm to improve computing efficiency. Then we describe how we reallocate the model parameters to keep the model size fixed in Section 3.2.

3.1 k -NN-Based Target Sampling

To reduce the expensive computation cost of the softmax function, we propose k -NN-based target sampling to approximate the expensive softmax. The original masked language modeling objective minimizes the cross-entropy loss for every masked subword w_i on the extensive multilingual vocabulary V . The proposed k -NN-based target sampling instead uses a smaller vocabulary subset V' . The approximation of the masked language modeling loss for the masked subword w_i is defined as follows:

$$\mathcal{L}(w_i) = -\log \frac{\exp(h^T v_{w_i} + b_{w_i})}{\sum_{w_j \in V'} \exp(h^T v_{w_j} + b_{w_j})} \quad (4)$$

where h is the corresponding output vector of the penultimate network layer, i.e., the output vector of the Transformer encoder, v_{w_i} is the embedding of the subword unit w_i , and b_{w_i} is a bias term. We formulate the construction of the vocabulary subset V' as follows:

$$V' = \bigcup_{w_i \in \mathcal{W}} \mathcal{I}_k(w_i) \quad (5)$$

$$\mathcal{I}_k(w_i) = \text{top-k}(\{v_{w_i}^T v_{w_j} \mid w_j \in V\}) \quad (6)$$

where \mathcal{W} denotes the set of target masked subword units in the current mini-batch, and $\mathcal{I}_k(w_i)$ denotes the k most similar subwords measured with the inner product of the subword embedding v_{w_i} and v_{w_j} .

However, retrieving $\mathcal{I}_k(w_i)$ at every training step for every subword unit $w_i \in \mathcal{W}$ requires as much

Algorithm 2 Pre-training with k -NN-based target sampling

Input: multilingual corpus \mathcal{D}_m ; size k of k -NN-based target sampling; multilingual vocabulary V ; learning rate τ

Output: model parameters θ

- 1: **while** not converged **do**
 - 2: Sample n mini-batches $\{\mathcal{X}^{(t)}, \mathcal{W}^{(t)}\}_{t=1}^n \sim \mathcal{D}_m$ \triangleright
 $\mathcal{X}^{(t)}$ is a mini-batch of monolingual text, and $\mathcal{W}^{(t)}$ is the set of masked subwords.
 - 3: Update $\mathcal{I}_k(w_i)$ for every $w_i \in V$
 - 4: **for** $t \leftarrow 1$ **to** n **do** \triangleright Train the model for n steps.
 - 5: $V' \leftarrow \bigcup_{w_i \in \mathcal{W}^{(t)}} \mathcal{I}_k(w_i)$
 - 6: $g \leftarrow \sum_{w_i \in \mathcal{W}^{(t)}} \nabla_{\theta} \mathcal{L}(w_i)$
 - 7: $\theta \leftarrow \theta - \tau g$
-

computation cost as softmax, which is unaffordable. As an alternative, we compute $\mathcal{I}_k(w_i)$ for every subword $w_i \in V$ according to the current subword embeddings every n training steps and replace the previous version of $\mathcal{I}_k(w_i)$ with the new one. We determine the value of n such that $|V| \ll n \times |\mathcal{W}|$. We illustrate the pre-training procedure with k -NN-based target sampling in Algorithm 2.

From a practical point of view under the cross-lingual setting, the previous sampling-based softmax approximation methods either sample subwords from recent mini-batches or samples subwords from unigram distribution, the task becomes simpler since a considerable part of the subword samples is from different languages. Meanwhile, our k -NN-based target sampling uses subwords with similar representations like synonyms, which enforces the model focus on discriminating the ground-truth subword from a set of noise samples that are not easy to distinguish. When using an approximate algorithm, the key point is to remain the difficult part of the original masked language modeling objective as much as possible.

3.2 Reducing the Embedding Dimension

In order to keep the number of model parameters fixed while increasing the vocabulary size, we follow (Lan et al., 2020) and (Chung et al., 2020a) to reduce both the input and output embedding dimension and linearly project the embeddings to the hidden dimension of the Transformer blocks. More precisely, we halve the embedding dimension when the vocabulary size is doubled. This rebalancing strategy only slightly degrades the model performance but improves pre-training speed and decreases the model size.

Conneau et al. (2020) also studied the relation between the size of the shared multilingual vocabulary and downstream task performance with multi-

Model	# Params	Speed	Pair Sentence		Structure Prediction		Question Answering			
			XNLI	PAWS-X	POS	NER	XQuAD	MLQA	TyDiQA	
			Acc.	Acc.	F1	F1	F1/EM	F1/EM	F1/EM	Avg.
XLM-R _{250K}	265M	1.00x	68.7	82.6	72.1	60.6	63.4/47.4	57.2/39.6	45.2/29.6	60.7
JOINT _{250K}	265M	1.00x	69.2	83.3	72.4	59.7	63.9/47.9	58.9/40.7	45.4/29.6	61.1
JOINT _{500K}	448M	0.72x	69.4	82.2	72.1	60.5	64.7/48.0	58.2/40.3	48.0/32.6	61.4
VOCAP _{250K}	265M	1.00x	69.3	82.0	71.4	60.0	66.2/50.3	60.1/42.6	45.6/30.6	61.5
VOCAP _{500K}	448M	0.72x	70.5	83.0	72.9	62.7	66.8/ 50.6	60.9/ 42.9	50.0/34.5	63.1
+ k -NN	448M	1.18x	70.8	82.6	72.5	61.8	67.1 /49.8	61.4 /42.5	56.3 / 39.3	63.7
+ half emb	265M	0.94x	70.3	83.0	72.0	61.7	65.8/49.0	61.0/42.3	49.3/33.0	62.5
+ k -NN & half emb	265M	1.35x	69.8	83.4	72.1	60.1	66.6/49.5	60.8/42.7	50.2/33.9	62.5

Table 1: Evaluation results on the XTREME benchmark. “XLM-R_{250K}” denotes using the XLM-R (Conneau et al., 2020) vocabulary with 250K subword units. “ k -NN” and “half emb” denote our k -NN-based target sampling method and using half embedding dimension, respectively.

lingual models of the fixed number of parameters. They keep the overall number of parameters constant by adjusting the width (i.e., hidden size) of the Transformer. Notice that we only reduce the embedding dimension while keeping the Transformer blocks untouched.

4 Experiment

4.1 Setup

Fine-Tuning Datasets To validate the effectiveness of our methods, we conduct experiments on three types of cross-lingual understanding tasks from XTREME benchmark (Hu et al., 2020), including two classification datasets: XNLI (Conneau et al., 2018), PAWS-X (Yang et al., 2019), three span extraction datasets: XQuAD (Artetxe et al., 2020), MLQA (Lewis et al., 2020), TyDiQA-GoldP (Clark et al., 2020), and two sequence labeling datasets: NER (Pan et al., 2017), POS (Zeman et al., 2019). The statistics of the datasets are shown in the appendix.

Implementation Details We adapt the Transformer architecture from the base model setting in Conneau et al. (2020), i.e., 12 layers and 768 hidden dimension size. We use masked language modeling objective to train our models for 1 million updates on eight 32GB Nvidia V100 GPUs with a batch size of 256. We update the top- k indices for every word in the multilingual vocabulary every 1,000 training steps and use $k = 50$ in k -NN-based target sampling. The learning rate is scheduled with a polynomial decay with 10K warmup steps, where the peak learning rate is set as 0.0001. We adapt other hyper-parameters in pre-training from Chi et al. (2021b). All fine-tuning results are averaged over five random seeds. The fine-tuning pipeline

is based on the code base of (Zheng et al., 2021). The fine-tuning implementation details are shown in the appendix.

4.2 Results

Table 1 shows XTREME fine-tuning results with models pretrained using different vocabularies and acceleration strategies. Compared to vocabulary directly learned on multilingual corpus with SentencePiece, i.e., XLM-R_{250K} and JOINT_{250K}, our VOCAP_{250K} improves on question answering datasets but degrades on PAWS-X, POS and NER. Then increasing the vocabulary from VOCAP_{250K} to VOCAP_{500K} mitigates the gap and bring improvements on six datasets except for PAWS-X, which only includes seven high-resource languages. However, increasing the size of vocabulary directly learned with Sentencepiece from JOINT_{250K} to JOINT_{500K} does not improve the performance as our VOCAP method does, showing the importance of selecting language-specific subword units and leveraging how much vocabulary capacity each language requires.

Since increasing vocabulary size brings the issues of model size and pre-training speed, we study the proposed method to accelerate pre-training: k -NN-based target sampling (k -NN) and using half embedding dimension (half emb). Our k -NN method improves pre-training speed with a 500K vocabulary so that the speed is 1.18 times that vanilla pre-training with a 250K vocabulary. Meanwhile, pre-training with our k -NN method does not significantly degrade the performance, it even brings improvement on XNLI, MLQA, and TyDiQA. Then we halve the embedding dimension of the models with 500K vocabulary and results in a similar number of parameters to models with 250K

Method	XNLI	POS	MLQA	Speed
VOCAP _{500K}	69.2	72.9	59.9/41.7	1.00x
+ k -NN	69.3	72.1	59.6/40.3	1.64x
+ target sampling	68.8	71.3	57.6/38.8	1.56x
+ NCE	56.0	61.8	41.1/26.2	1.40x
+ NEG	56.5	62.9	40.1/25.6	1.40x

Table 2: Comparison between different sampling-based softmax approximation approaches with vocabulary VOCAP_{500K}. Models are pretrained for 0.5M steps.

Method	XNLI	POS	MLQA	Speed
VOCAP _{500K}	69.2	71.8	59.9/41.7	1.00x
+ k -NN ($k=5$)	68.5	71.3	58.6/40.0	1.76x
+ k -NN ($k=10$)	69.3	71.4	58.9/39.6	1.74x
+ k -NN ($k=25$)	69.2	71.7	59.8/40.9	1.69x
+ k -NN ($k=50$)	69.3	72.1	59.6/40.3	1.64x
+ k -NN ($k=100$)	69.5	72.1	60.0/41.3	1.57x

Table 3: Comparison between different k values in k -NN-based sampling method. Models are pretrained for 0.5M steps.

vocabulary. The overall performance degrades by 0.6-points but still consistently improves over models with 250K vocabularies while the speed is comparable. Combining the two methods above, we achieve a 1.35-times speed-up and more than 1 point improvement with a similar model size compared to models with 250K vocabularies.

4.3 Analysis and Discussion

We conduct a thorough analysis to understand the impact of our proposed methods on cross-lingual language models. To reduce the computation load, we only pre-train the cross-lingual language models for 500K steps for some of our settings.

k -NN-based target sampling outperforms previous sampling-based approaches. To verify the effectiveness of our proposed k -NN-based sampling method, we compare it against previous sampling-based approaches used to approximate softmax, which are target sampling (Jean et al., 2015), noise contrastive estimation (Mnih and Teh (2012), NCE) and negative sampling (Mikolov et al. (2013), NEG). The results are shown in Table 2. To make a fair comparison, since our k -NN-based sampling method using $k = 50$ samples vocabulary subset with less than 50,000 subword units per batch on average, we here sample 50,000 negative subword units per batch for target sampling, NCE, and NEG. Among the four methods, NCE and NEG are significantly worse than k -NN and

Method	XNLI	POS	NER	MLQA
$\beta=0$	66.9	71.8	61.5	58.6/41.0
$\beta=0.3$	69.0	71.7	61.6	59.2/40.1
$\beta=0.7$	69.2	71.8	61.5	59.9/41.7
$\beta=1.0$	69.5	71.8	60.9	58.4/40.3

Table 4: Impact of adjusting high-resource versus low-resource vocabulary capacity trade-off with β . $\beta = 0$ indicates the vocabulary is allocated without considering pre-training corpus size. Models are pretrained for 0.5M steps.

target sampling. We attribute it that NCE and NEG need more training steps to converge (Mnih and Teh, 2012). Besides, the original NCE typically sample different negative samples for every target word, while we here use 50,000 negative samples for all target word in current mini-batch, which is more efficient on GPUs.

Effect of the value of k in k -NN-based target sampling.

We illustrate the downstream task performance when using different values of k in our k -NN-based target sampling in Table 3. While a smaller k indicates faster pre-training speed, we observe even with a small value like 5, the result does not significantly degrade compared to using the original softmax. We attribute this to that by retrieving subword samples that are most similar to the target subword, the model can focus on the difficult part of the original masked language modeling objective. More precisely, the model focus on discriminating the ground-truth subword from a set of noise samples that are not easy to distinguish. Considering the overall performance, the pre-training speed, and running memory to store k -NN indices, we use $k = 50$ in all our experiments.

Language-specific pre-training corpus should also be considered when allocating vocabulary capacity.

The pre-training corpus size varies across different languages. It is inefficient to allocate a large vocabulary capacity for low-resource languages with rare pre-training data since the pre-trained model can only learn limited knowledge from these languages. Here we study the value of rescaling factor β from Equation (2) in multilingual vocabulary construction in Table 4. The rescaling factor β controls the number of selected language-specific subword units. Increasing the value of β improves the performance of XNLI, where most languages are high-resource languages. However, it degrades the performance of NER, where more

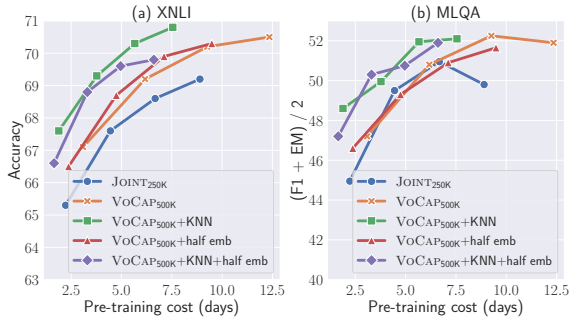


Figure 5: Performance on XNLI and MLQA versus the cross-lingual language models’ pre-training cost.

low-resources languages exist. When considering overall performance, we decide to use $\beta = 0.7$ in our experiments.

The proposed acceleration strategies significantly improve the downstream task performance under the same pre-training cost. Increasing the vocabulary size slows the pre-training speed, even though there is almost no difference in fine-tuning speed. We study the relationship between the downstream task performance and the pre-training cost under different model settings in Figure 5. We observe VO_{CAP}_{500K}+ k -NN achieves the best performance. Models trained with 500K vocabulary consistently outperform 250K vocabulary on XNLI. Besides, we observe the performance on MLQA with the model trained using 250K vocabulary degrades as the training continues while models trained using 500K vocabulary does not, indicating the sufficient vocabulary capacity is essential for question answering task.

VO_{CAP} gains more improvement on mid and low-resource languages than high-resource languages. In Figure 4 in Section 2, we show that the vocabulary learned with VO_{CAP} benefits the vocabulary capacity of low-resource languages more than high-resource languages, indicating the improvements should mainly come from low-resource languages. To verify this, we compare VO_{CAP} against SentencePiece baseline on the performance of different-resourced languages on XNLI and NER in Figure 6. We observe that the vocabulary learned with VO_{CAP} significantly outperforms the vocabularies directly learned with SentencePiece on mid and low-resource languages. This observation is also consistent with the ALP results in Figure 4.




Figure 6: Impact of VO_{CAP} on the performance of different-resourced languages on XNLI and NER.

5 Related Work

Pretrained Cross-Lingual Language Models

Recent work pre-trains Transformer models (Vaswani et al., 2017) on the large-scale multilingual corpus to obtain pretrained cross-lingual language models (Conneau and Lample, 2019; Conneau et al., 2020; Chi et al., 2020, 2021a,b,c,d; Chung et al., 2020a; Xue et al., 2020; Ma et al., 2020, 2021). These models are capable of encoding texts from different languages into universal representations and significantly improves cross-lingual transferability.

Multilingual Vocabulary Construction

Cross-lingual language models need large vocabularies to ensure all languages are adequately represented. Recent research work on constructing multilingual vocabulary for cross-lingual language models can be categorized into two groups. mBERT (Devlin et al., 2019), XLM (Conneau and Lample, 2019), and XLM-R (Conneau et al., 2020) learn vocabularies on a combined multilingual corpus with WordPiece (Wu et al., 2016), BPE (Sennrich et al., 2015), and unigram language model (Kudo, 2018) from SentencePiece (Kudo and Richardson, 2018), respectively. Chung et al. (2020b) propose to balance the trade-off between optimizing for cross-lingual subword sharing and the need for robust representation of individual languages. They first group languages into clusters and learn vocabularies individually on each cluster, then combine all cluster-vocabularies to form a single unified multilingual vocabulary. Compared to Chung et al. (2020b), our advantage is that we separately quantify the vocabulary capacity each language needs with average log probability and balance the construction procedure with pre-training corpus size.

Softmax Approximation Approximating the softmax was a core problem in training NLP tasks with a large vocabulary, e.g., neural machine translation, language modeling. With the rise of subword representations (Sennrich et al., 2015; Wu et al., 2016; Kudo, 2018), the vocabulary size significantly decreases, and the problem has been less studied recently. Nevertheless, the need for training cross-lingual language models with a large multilingual vocabulary has drawn our attention again to the softmax approximation approaches. The existing softmax approximation approaches can be grouped into softmax-based and sampling-based approaches. Softmax-based approaches include hierarchical softmax (Morin and Bengio, 2005), differentiated softmax (Chen et al., 2016), and CNN-softmax (Kim et al., 2016). However, these approaches improve the softmax efficiency by changing its architecture, which is unsuitable for either training on GPUs or multilingual settings. Sampling-based approaches instead optimize some other easy-to-compute loss function to approximate the original softmax, including target sampling (Jean et al., 2015), noise contrastive estimation (Mnih and Teh, 2012), negative sampling (Mikolov et al., 2013). Our k -NN-based target sampling is also a sampling-based approach.

6 Conclusion

In this paper, we study pre-training cross-lingual language models with large vocabulary capacity. First, we propose VOCAP to construct large multilingual vocabulary in cross-lingual language models. We conduct a quantitative analysis to show that average log probability is a valid indicator of vocabulary capacity for a particular language, which also correlates with downstream task performance on the language. VOCAP uses the language-specific average log probability and pre-training corpus size to allocate appropriate vocabulary capacity for each language in the multilingual vocabulary. Moreover, we propose k -NN-based target sampling to accelerate pre-training with the allocated large multilingual vocabulary by approximating the expensive softmax. We also show that reducing the embedding dimension is an effective way to keep the improvement brought by the large vocabulary without increasing the number of model parameters. The experiments demonstrate the effectiveness of the proposed vocabulary construction method as well as the acceleration methods.

Acknowledgments

We would like to acknowledge Zewen Chi, Shuming Ma for the helpful discussions. This work was supported by the National Key R&D Program of China via grant 2020AAA0106501 and the National Natural Science Foundation of China (NSFC) via grant 61976072 and 61772153. Wanxiang Che is the corresponding author.

References

- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. [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.
- Hangbo Bao, Li Dong, Furu Wei, Wenhui Wang, Nan Yang, Xiaodong Liu, Yu Wang, Jianfeng Gao, Songhao Piao, Ming Zhou, and Hsiao-Wuen Hon. 2020. [UniLMv2: Pseudo-masked language models for unified language model pre-training](#). In *Proceedings of the 37th International Conference on Machine Learning*, pages 7006–7016.
- Wenlin Chen, David Grangier, and Michael Auli. 2016. [Strategies for training large vocabulary neural language models](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Zewen Chi, Li Dong, Shuming Ma, Shaohan Huang, Xian-Ling Mao, Heyan Huang, and Furu Wei. 2021a. [mT6: Multilingual pretrained text-to-text transformer with translation pairs](#). *arXiv preprint arXiv:2104.08692*.
- Zewen Chi, Li Dong, Furu Wei, Wenhui Wang, Xian-Ling Mao, and Heyan Huang. 2020. [Cross-lingual natural language generation via pre-training](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 7570–7577. AAAI Press.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2021b. [InfoXLM: An information-theoretic framework for cross-lingual language model pre-training](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3576–3588, Online. Association for Computational Linguistics.
- Zewen Chi, Li Dong, Bo Zheng, Shaohan Huang, Xian-Ling Mao, Heyan Huang, and Furu Wei. 2021c.

- Improving pretrained cross-lingual language models via self-labeled word alignment. 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 3418–3430, Online. Association for Computational Linguistics.
- Zewen Chi, Shaohan Huang, Li Dong, Shuming Ma, Saksham Singhal, Payal Bajaj, Xia Song, and Furu Wei. 2021d. XLM-E: Cross-lingual language model pre-training via ELECTRA. *ArXiv*, abs/2106.16138.
- Hyung Won Chung, Thibault Févry, Henry Tsai, Melvin Johnson, and Sebastian Ruder. 2020a. Re-thinking embedding coupling in pre-trained language models. *CoRR*, abs/2010.12821.
- Hyung Won Chung, Dan Garrette, Kiat Chuan Tan, and Jason Riesa. 2020b. Improving multilingual models with language-clustered vocabularies. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 4536–4546. Association for Computational Linguistics.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. TyDi QA: A benchmark for information-seeking question answering in typologically diverse languages. *Transactions of the Association for Computational Linguistics*, 8:454–470.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 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 *Advances in Neural Information Processing Systems*, pages 7057–7067. Curran Associates, Inc.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel 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*, pages 2475–2485, Brussels, Belgium. 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.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. In *Advances in Neural Information Processing Systems*, pages 13063–13075. Curran Associates, Inc.
- 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 generalization. *arXiv preprint arXiv:2003.11080*.
- Sébastien Jean, KyungHyun Cho, Roland Memisevic, and Yoshua Bengio. 2015. On using very large target vocabulary for neural machine translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 1–10. The Association for Computer Linguistics.
- Yoon Kim, Yacine Jernite, David A. Sontag, and Alexander M. Rush. 2016. Character-aware neural language models. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 2741–2749. AAAI Press.
- 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, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 66–75. 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, EMNLP 2018: System Demonstrations, Brussels, Belgium, October 31 - November 4, 2018*, pages 66–71. Association for Computational Linguistics.
- 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.
- Patrick Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2020. MLQA: Evaluating cross-lingual extractive question answering. In

- Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7315–7330, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Shuming Ma, Li Dong, Shaohan Huang, Dongdong Zhang, Alexandre Muzio, Saksham Singhal, Hany Hassan Awadalla, Xia Song, and Furu Wei. 2021. DeltaLM: Encoder-decoder pre-training for language generation and translation by augmenting pretrained multilingual encoders. *ArXiv*, abs/2106.13736.
- Shuming Ma, Jian Yang, H. Huang, Zewen Chi, Li Dong, Dongdong Zhang, Hany Hassan Awadalla, Alexandre Muzio, Akiko Eriguchi, Saksham Singhal, Xia Song, Arul Menezes, and Furu Wei. 2020. XLM-T: Scaling up multilingual machine translation with pretrained cross-lingual transformer encoders. *ArXiv*, abs/2012.15547.
- 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.
- Andriy Mnih and Yee Whye Teh. 2012. A fast and simple algorithm for training neural probabilistic language models. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*. icml.cc / Omnipress.
- Frederic Morin and Yoshua Bengio. 2005. Hierarchical probabilistic neural network language model. In *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics, AISTATS 2005, Bridgetown, Barbados, January 6-8, 2005*. Society for Artificial Intelligence and Statistics.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. Cross-lingual name tagging and linking for 282 languages. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1958, Vancouver, Canada. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008. Curran Associates, Inc.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, abs/1609.08144.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2020. mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934*.
- Yinfei Yang, Yuan Zhang, Chris Tar, and Jason Baldridge. 2019. PAWS-X: A cross-lingual adversarial dataset for paraphrase identification. 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 3687–3692, Hong Kong, China. Association for Computational Linguistics.
- Daniel Zeman, Joakim Nivre, Mitchell Abrams, and 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.
- Bo Zheng, Li Dong, Shaohan Huang, Wenhui Wang, Zewen Chi, Saksham Singhal, Wanxiang Che, Ting Liu, Xia Song, and Furu Wei. 2021. Consistency regularization for cross-lingual fine-tuning. 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 3403–3417, Online. Association for Computational Linguistics.

A Correlation between Language-Specific Vocabulary Capacity and Task Performance

We compare the Pearson correlation coefficients between ALP and downstream task performance with the coefficients between vocabulary size and downstream task performance in Table 5. The results show that ALP correlates better than vocabulary size with downstream task performance.

B Statistics of XTREME Datasets

Language	Task	$\rho(\text{ALP}, \text{F1})$	$\rho(V , \text{F1})$
hi	POS	0.922	0.787
	NER	0.879	0.890
fa	POS	0.905	0.700
	NER	0.912	0.872
it	POS	0.665	0.422
	NER	0.899	0.900
ru	POS	0.423	0.327
	NER	0.872	0.833

Table 5: Pearson correlation coefficients between ALP and downstream task performance and between vocabulary size and downstream task performance.

Task	Dataset	Train	Lang
Classification	XNLI	392K	15
	PAWS-X	49.4K	7
Structured Prediction	POS	21K	33
	NER	20K	40
Question Answering	XQuAD	87K	11
	MLQA	87K	7
	TyDiQA	3.7K	9

Table 6: Statistics for the datasets in the XTREME benchmark. we report the number of training examples (|Train|), and the number of languages (|Lang|).

C Fine-tuning Settings

Implementation Details For the POS dataset, we use the average-pooling strategy on subwords to obtain word representation since part-of-speech is related to different parts of words, depending on the language. We tune the hyper-parameter and select the model with the best average results over all the languages’ development set. There are two datasets without development set in multi-languages. For XQuAD, we tune the hyper-parameters with the development set of MLQA since they share the same training set and have a higher degree of overlap in languages. For TyDiQA-GoldP, we use the English test set as the development set.

Hyper-Parameters For XNLI, PAWS-X, POS, and NER, we fine-tune 10 epochs. For XQuAD and MLQA, we fine-tune 4 epochs. For TyDiQA-GoldP, we fine-tune 6 or 8 epochs and select the best number of epochs with the English test set as the development set. For learning rate, we select

Code	Size (GB)	Code	Size (GB)	Code	Size (GB)
af	0.2	hu	9.5	pl	28.6
am	0.4	hy	0.7	ps	0.4
ar	16.1	id	17.2	pt	39.4
as	0.1	is	0.5	ro	11.0
az	0.8	it	47.2	ru	253.3
ba	0.2	ja	86.8	sa	0.2
be	0.5	ka	1.0	sd	0.2
bg	7.0	kk	0.6	si	1.3
bn	5.5	km	0.2	sk	13.6
ca	3.0	kn	0.3	sl	6.2
cs	14.9	ko	40.0	sq	3.0
cy	0.4	ky	0.5	sr	7.2
da	6.9	la	0.3	sv	60.4
de	99.0	lo	0.2	sw	0.3
el	13.1	lt	2.3	ta	7.9
en	731.6	lv	1.3	te	2.3
eo	0.5	mk	0.6	tg	0.7
es	85.6	ml	1.3	th	33.0
et	1.4	mn	0.4	tl	1.2
eu	1.0	mr	0.5	tr	56.4
fa	19.0	ms	0.7	tt	0.6
fi	5.9	mt	0.2	ug	0.2
fr	89.9	my	0.4	uk	13.4
ga	0.2	ne	0.6	ur	3.0
gl	1.5	nl	25.9	uz	0.1
gu	0.3	nn	0.4	vi	74.5
he	4.4	no	5.5	yi	0.3
hi	5.0	or	0.3	zh	96.8
hr	1.4	pa	0.8		

Table 7: The statistics of the reconstructed CommonCrawl corpus for learning vocabularies.

in [7e-6, 1e-5] for XNLI and PAWS-X, [1e-5, 2e-5] for POS and NER, [2e-5, 3e-5] for XQuAD, MLQA and TyDiQA-GoldP.

D Pre-Training Data

We use the reconstruct CommonCrawl corpus in Chi et al. (2021b) to learn vocabularies in our paper. Because tokenizing the pre-training data is time-consuming, we instead conduct our pre-training on Wikipedia since it has a smaller size. We only consider the languages that are shared by the reconstructed CommonCrawl corpus and Wikipedia. The statistics of the Wikipedia corpus and the reconstructed CommonCrawl corpus are listed in Table 8 and Table 7.

Code	Size (GB)	Code	Size (GB)	Code	Size (GB)
af	0.12	hu	0.8	pl	1.55
am	0.01	hy	0.6	ps	0.04
ar	1.29	id	0.52	pt	1.5
as	0.04	is	0.05	ro	0.42
az	0.24	it	2.69	ru	5.63
ba	0.13	ja	2.65	sa	0.04
be	0.31	ka	0.37	sd	0.02
bg	0.62	kk	0.29	si	0.09
bn	0.41	km	0.12	sk	0.21
ca	1.1	kn	0.25	sl	0.21
cs	0.8	ko	0.56	sq	0.1
cy	0.06	ky	0.1	sr	0.74
da	0.33	la	0.05	sv	1.7
de	5.43	lo	0.01	sw	0.03
el	0.73	lt	0.19	ta	0.46
en	12.58	lv	0.12	te	0.44
eo	0.25	mk	0.34	tg	0.04
es	3.38	ml	0.28	th	0.52
et	0.23	mn	0.05	tl	0.04
eu	0.24	mr	0.1	tr	0.43
fa	0.66	ms	0.2	tt	0.09
fi	0.68	mt	0.01	ug	0.03
fr	4.0	my	0.15	uk	2.43
ga	0.03	ne	0.06	ur	0.13
gl	0.27	nl	1.38	uz	0.06
gu	0.09	nn	0.13	vi	0.76
he	1.11	no	0.54	yi	0.02
hi	0.38	or	0.04	zh	1.08
hr	0.28	pa	0.1		

Table 8: The statistics of the Wikipedia corpus used for pre-training.

Code	Size (K)	Code	Size (K)	Code	Size (K)
af	2	hu	12	pl	20
am	3	hy	5	ps	3
ar	15	id	13	pt	20
as	2	is	3	ro	13
az	5	it	22	ru	34
ba	2	ja	23	sa	1
be	3	ka	4	sd	2
bg	9	kk	4	si	3
bn	6	km	4	sk	11
ca	8	kn	2	sl	8
cs	14	ko	17	sq	7
cy	3	ky	3	sr	10
da	9	la	3	sv	18
de	24	lo	2	sw	3
el	17	lt	7	ta	6
en	23	lv	6	te	4
eo	4	mk	4	tg	5
es	26	ml	3	th	14
et	5	mn	3	tl	4
eu	4	mr	3	tr	18
fa	9	ms	4	tt	3
fi	9	mt	3	ug	3
fr	25	my	2	uk	12
ga	2	ne	3	ur	5
gl	5	nl	14	uz	2
gu	2	nn	3	vi	12
he	6	no	7	yi	2
hi	6	or	2	zh	30
hr	6	pa	3		

Table 9: The statistics of the allocated vocabulary size for each language.