

# A Closer Look at Few-Shot Crosslingual Transfer: The Choice of Shots Matters

Mengjie Zhao<sup>1\*</sup> Yi Zhu<sup>2\*</sup> Ehsan Shareghi<sup>3,2</sup> Ivan Vulić<sup>2</sup>

Roi Reichart<sup>4</sup> Anna Korhonen<sup>2</sup> Hinrich Schütze<sup>1</sup>

<sup>1</sup>CIS, LMU Munich <sup>2</sup>LTL, University of Cambridge

<sup>3</sup>Department of Data Science & AI, Monash University

<sup>4</sup>Faculty of Industrial Engineering and Management, Technion, IIT  
mzhao@cis.lmu.de, {yz568, iv250, alk23}@cam.ac.uk,  
ehsan.shareghi@monash.edu, roiri@technion.ac.il

## Abstract

Few-shot crosslingual transfer has been shown to outperform its zero-shot counterpart with pretrained encoders like multilingual BERT. Despite its growing popularity, little to no attention has been paid to standardizing and analyzing the design of few-shot experiments. In this work, we highlight a fundamental risk posed by this shortcoming, illustrating that the model exhibits a high degree of sensitivity to the selection of few shots. We conduct a large-scale experimental study on 40 sets of sampled few shots for six diverse NLP tasks across up to 40 languages. We provide an analysis of success and failure cases of few-shot transfer, which highlights the role of lexical features. Additionally, we show that a straightforward full model finetuning approach is quite effective for few-shot transfer, outperforming several state-of-the-art few-shot approaches. As a step towards standardizing few-shot crosslingual experimental designs, we make our sampled few shots publicly available.<sup>1</sup>

## 1 Introduction

Multilingual pretrained encoders like multilingual BERT (mBERT; Devlin et al. (2019)) and XLM-R (Conneau et al., 2020) are the top performers in crosslingual tasks such as natural language inference (Conneau et al., 2018), document classification (Schwenk and Li, 2018; Artetxe and Schwenk, 2019), and argument mining (Toledo-Ronen et al., 2020). They enable transfer learning through language-agnostic representations in crosslingual setups (Hu et al., 2020).

A widely explored transfer scenario is *zero-shot crosslingual transfer* (Pires et al., 2019; Conneau and Lample, 2019; Artetxe and Schwenk, 2019),

where a pretrained encoder is finetuned on abundant task data in the source language (e.g., English) and then directly evaluated on target-language test data, achieving surprisingly good performance (Wu and Dredze, 2019; Hu et al., 2020). However, there is evidence that zero-shot performance reported in the literature has large variance and is often not reproducible (Keung et al., 2020a; Rios et al., 2020); the results in languages distant from English fall far short of those similar to English (Hu et al., 2020; Liang et al., 2020).

Lauscher et al. (2020) stress the importance of *few-shot crosslingual transfer* instead, where the encoder is first finetuned on a source language and then further finetuned with a small amount (10–100) of examples (**few shots**) of the target language. The few shots substantially improve model performance of the target language with negligible annotation costs (Garrette and Baldridge, 2013; Hedderich et al., 2020).

In this work, however, we demonstrate that the gains from few-shot transfer exhibit a *high degree of sensitivity to the selection of few shots*. For example, different choices for the few shots can yield a performance variance of over 10% accuracy in a standard document classification task. Motivated by this, we propose to fix the few shots for fair comparisons between different crosslingual transfer methods, and provide a benchmark resembling the standard “*N*-way *K*-shot” few-shot learning configuration (Fei-Fei et al., 2006; Koch et al., 2015). We also evaluate and compare several state-of-the-art (SotA) few-shot finetuning techniques, in order to understand their performance and susceptibility to the variance related to few shots.

We also demonstrate that the effectiveness of few-shot crosslingual transfer depends on the type of downstream task. For syntactic tasks such as named-entity recognition, the few shots can improve results by up to  $\approx 20 F_1$  points. For chal-

\* Equal contribution.

<sup>1</sup>Code and resources are available at <https://github.com/fsxlt>

lenging tasks like adversarial paraphrase identification, the few shots do not help and even sometimes lead to worse performance than zero-shot transfer. To understand these phenomena, we conduct additional in-depth analyses, and find that the models tend to utilize shallow lexical hints (Geirhos et al., 2020) in the target language, rather than leveraging abstract crosslingual semantic features learned from the source language.

Our **contributions**: **1)** We show that few-shot crosslingual transfer is prone to large variations in task performance; this property hinders unbiased assessments of the effectiveness of different few-shot methods. **2)** To remedy this issue, we publish fixed and standardized few shots to support fair comparisons and reproducibility. **3)** We empirically verify that few-shot crosslingual transfer has different performance impact on structurally different tasks; we provide in-depth analyses concerning the source of performance gains. **4)** We analyze several SotA few-shot learning methods, and show that they underperform simple full model finetuning. We hope that our work will shed new light on the potential and current difficulties of few-shot learning in crosslingual setups.

## 2 Background and Related Work

**Zero-/Few-Shot Crosslingual Transfer.** Multilingual pretrained encoders show strong zero-shot crosslingual transfer (**ZS-XLT**) ability in various NLP tasks (Pires et al., 2019; Hsu et al., 2019; Artetxe and Schwenk, 2019). In order to guide and measure the progress, standardized benchmarks like XTREME (Hu et al., 2020) and XGLUE (Liang et al., 2020) have been developed.

Recently, Lauscher et al. (2020) and Hedderich et al. (2020) extended the focus on few-shot crosslingual transfer (**FS-XLT**): They assume the availability of a handful of labeled examples in a target language,<sup>2</sup> which are used to further finetune a source-trained model. The extra few shots bring large performance gains at low annotation cost. In this work, we systematically analyze this recent FS-XLT scenario.

FS-XLT resembles the intermediate-task transfer (STILT) approach (Phang et al., 2018; Pruksachatkun et al., 2020). In STILT, a pretrained encoder is finetuned on a resource-rich intermedi-

<sup>2</sup>According to Garrette and Baldrige (2013), it is possible to collect  $\approx 100$  POS-annotated sentences in two hours even for low-resource languages such as Malagasy.

ate task, and then finetuned on a (resource-lean) target task. Likewise, FS-XLT focuses on transferring knowledge and general linguistic intelligence (Yogatama et al., 2019), although such transfer is between *languages* in the same task instead of between different tasks.

**Few-shot learning** was first explored in computer vision (Miller et al., 2000; Fei-Fei et al., 2006; Koch et al., 2015); the aim there is to learn new concepts with only few images. Methods like prototypical networks (Snell et al., 2017) and model-agnostic meta-learning (MAML; Finn et al. (2017)) have also been applied to many monolingual (typically English) NLP tasks such as relation classification (Han et al., 2018; Gao et al., 2019), named-entity recognition (Hou et al., 2020a), word sense disambiguation (Holla et al., 2020), and text classification (Yu et al., 2018; Yin, 2020; Yin et al., 2020; Bansal et al., 2020; Gupta et al., 2020). However, recent few-shot learning methods in computer vision consisting of two simple finetuning stages, first on base-class images and then on new-class few shots, have been shown to outperform MAML and achieve SotA scores (Wang et al., 2020; Chen et al., 2020; Tian et al., 2020; Dhillion et al., 2020). Inspired by this work, we compare various few-shot finetuning methods from computer vision in the context of FS-XLT.

**Task Performance Variance.** Deep neural networks’ performance on NLP tasks is bound to exhibit large variance. Reimers and Gurevych (2017) and Dror et al. (2019) stress the importance of reporting *score distributions* instead of a single score for fair(er) comparisons. Dodge et al. (2020), Mosbach et al. (2021), and Zhang et al. (2021) show that finetuning pretrained encoders with different random seeds yields performance with large variance. In this work, we examine a specific source of variance: We show that the choice of the few shots in crosslingual transfer learning also introduces large variance in performance; consequently, we offer standardized few shots for more controlled and fair comparisons.

## 3 Method

Following Lauscher et al. (2020) and Hedderich et al. (2020), our FS-XLT method comprises two stages. First, we conduct **source-training**: The pretrained mBERT is finetuned with abundant annotated data in the source language. Similar to Hu et al. (2020), Liang et al. (2020) and due to

Name	Metric	Task	$ \mathcal{T} $	TS	# of lang.
XNLI	Acc.	Natural language inference	3	No	15
PAWSX	Acc.	Paraphrase identification	2	No	7
MLDoc	Acc.	News article classification	4	Yes	8
MARC	Acc.	Amazon reviews	5	Yes	6
POS	F1	Part-of-speech tagging	17	Yes	29
NER	F1	Named-entity recognition	7	Yes	40

Table 1: Evaluation datasets.  $|\mathcal{T}|$ : Number of classes (classification tasks) and label set size (POS and NER). TS: availability of a training split in the target language.

the abundant labeled data for many NLP tasks, we choose English as the source in our experiments. Directly evaluating the source-trained model after this stage corresponds to the widely studied ZS-XLT scenario. The second stage is **target-adapting**: The source-trained model from previous stage is adapted to a target language using few shots. We discuss details of sampling the few shots in §4. The development set of the target language is used for model selection in this stage.

## 4 Experimental Setup

We consider three types of tasks requiring varying degrees of semantic and syntactic knowledge transfer: Sequence classification (**CLS**), named-entity recognition (**NER**), and part-of-speech tagging (**POS**) in up to 40 typologically diverse languages (cf., Appendix §B).

### 4.1 Datasets and Selection of Few Shots

For the CLS tasks, we sample few shots from four multilingual datasets: News article classification (MLDoc; Schwenk and Li (2018)); Amazon review classification (MARC; Keung et al. (2020b)); natural language inference (XNLI; Conneau et al. (2018); Williams et al. (2018)); and crosslingual paraphrase adversaries from word scrambling (PAWSX; Zhang et al. (2019); Yang et al. (2019)). We use treebanks in Universal Dependencies (Nivre et al., 2020) for POS, and WikiANN dataset (Pan et al., 2017; Rahimi et al., 2019) for NER. Table 1 reports key information about the datasets.

We adopt the conventional few-shot sampling strategy (Fei-Fei et al., 2006; Koch et al., 2015; Snell et al., 2017), and conduct “ $N$ -way  $K$ -shot” sampling from the datasets;  $N$  is the number of classes and  $K$  refers to the number of shots per class. A group of  $N$ -way  $K$ -shot data is referred to as a **bucket**. We set  $N$  equal to the number of labels  $|\mathcal{T}|$ . Following Wang et al. (2020), we sample 40 buckets for each target (i.e., non-English)

language of a task to get a reliable estimation of model performance.

**CLS Tasks.** For MLDoc and MARC, each language has a train/dev/test split. We sample the buckets without replacement from the training set of each target language, so that buckets are disjoint from each other. Target languages in XNLI and PAWSX only have dev/test splits. We sample the buckets from the dev set; the remaining data serves as a single new dev set for model selection during target-adapting. For all tasks, we use  $K \in \{1, 2, 4, 8\}$ .

**POS and NER.** For the two structured prediction tasks, “ $N$ -way  $K$ -shot” is not well-defined because each sentence contains one or more labeled tokens. We use a similar sampling principle as with CLS, where  $N$  is the size of the label set for each language and task, but  $K$  is set to the minimum number of occurrences for each label. In particular, we utilize the *Minimum-Including Algorithm* (Hou et al., 2020b,a) to satisfy the following criteria when sampling a bucket: 1) each label appears at least  $K$  times, and 2) at least one label will appear less than  $K$  times if any sentence is removed from the bucket. Appendix §C gives sampling details. In contrast to sampling for CLS, we do not enforce samples from different buckets to be disjoint due to the small amount of data in some low-resource languages. We only use  $K \in \{1, 2, 4\}$  and exclude  $K = 8$ , as 8-shot buckets already have lots of labeled tokens, and thus (arguably) might not be considered few-shot.

### 4.2 Training Setup

We use the pretrained cased mBERT model (Devlin et al., 2019), and rely on the PyTorch-based (Paszke et al., 2019) HuggingFace Transformers repository (Wolf et al., 2019) in all experiments.

For *source-training*, we finetune the pretrained encoder for 10 epochs with batch size 32. For *target-adapting to every* target language, the few-shot data is a sampled bucket in this language, and we finetune on the bucket for 50 epochs with early-stopping of 10 epochs. The batch size is set to the number of shots in the bucket. Each target-adapting experiment is repeated 40 times using the 40 buckets. We use the Adam optimizer (Kingma and Ba, 2015) with default parameters in both stages with learning rates searched over  $\{1e-5, 3e-5, 5e-5, 7e-5\}$ . For CLS tasks, we use mBERT’s [CLS] token as the final represen-




Figure 1: Histograms of dev set accuracies. Top: 40 runs with different random seeds. Bottom: 40 runs with different 1-shot buckets. Left: DE MARC. Right: ES MLDoc. The variance due to buckets is larger.

tation. For NER and POS, following Devlin et al. (2019), we use a linear classifier layer on top of the representation of each tokenized word, which is its last wordpiece (He and Choi, 2020).

We set the maximum sequence length to 128 after wordpiece tokenization (Wu et al., 2016), in all experiments. Further implementation details are shown in our Reproducibility Checklist in Appendix §A.

## 5 Results and Discussion

### 5.1 Source-Training Results

The ZS-XLT performance from English (EN) to target languages of the four CLS tasks are shown in the  $K = 0$  column in Table 2. For NER and POS, the results are shown in Figure 2.

For XTREME tasks (XNLI, PAWSX, NER, POS), our implementation delivers results comparable to Hu et al. (2020). For MLDoc, our results are comparable to (Dong and de Melo, 2019; Wu and Dredze, 2019; Eisenschlos et al., 2019). It is worth noting that reproducing the exact results is challenging, as suggested by Keung et al. (2020a). For MARC, our zero-shot results are worse than Keung et al. (2020b)’s who use the dev set of each target language for model selection while we use EN dev, following the common true ZS-XLT setup.

### 5.2 Target-Adapting Results

**Variance of Few-Shot Transfer.** We hypothesize that FS-XLT suffers from large variance (Dodge et al., 2020) due to the large model complexity and small amount of data in a bucket. To test this empirically, we first conduct two experiments on MLDoc and MARC. First, for a *fixed random seed*, we repeat 1-shot target-adapting 40 times using different 1-shot buckets in German (DE) and Spanish (ES). Second, for a fixed 1-shot bucket, we repeat the same experiment 40 times using random seeds

in  $\{0 \dots 39\}$ . Figure 1 presents the dev set performance distribution of the 40 runs with 40 random seeds (top) and 40 1-shot buckets (bottom).

With exactly the same training data, using different random seeds yields a 1–2 accuracy difference of FS-XLT (Figure 1 top). A similar phenomenon has been observed in finetuning monolingual encoders (Dodge et al., 2020) and multilingual encoders with ZS-XLT (Keung et al., 2020a; Wu and Dredze, 2020b; Xia et al., 2020); we show this observation also holds for FS-XLT. The key takeaway is that varying the buckets is a more severe problem. It causes much larger variance (Figure 1 bottom): The maximum accuracy difference is  $\approx 6$  for DE MARC and  $\approx 10$  for ES MLDoc. This can be due to the fact that difficulty of individual examples varies in a dataset (Swayamdipta et al., 2020), resulting in different amounts of information encoded in buckets.

This large variance could be an issue when comparing different few-shot learning algorithms. The bucket choice is a strong confounding factor that may obscure the strength of a promising few-shot technique. Therefore, for fair comparison, *it is necessary to work with a fixed set of few shots*. We propose to fix the sampled buckets for unbiased comparison of different FS-XLT methods. We publish the sampled buckets from the six multilingual datasets as a fixed and standardized few-shot evaluation benchmark.

In what follows, each FS-XLT experiment is repeated 40 times using 40 different buckets with the same fixed random seed; we report mean and standard deviation. As noted, the variance due to random seeds is smaller (cf., Figure 1) and has been well studied before (Reimers and Gurevych, 2017; Dodge et al., 2020). In this work, we thus focus our attention and limited computing resources on understanding the impact of buckets, the *newly detected* source of variance. However, we encourage practitioners to report results with both factors considered in the future.

**Different Numbers of Shots.** A comparison concerning the number of shots ( $K$ ), based on the few-shot results in Table 2 and Figure 2, reveals that the buckets largely improve model performance on a majority of tasks (MLDoc, MARC, POS, NER) over zero-shot results. This is in line with prior work (Lauscher et al., 2020; Hedderich et al., 2020) and follows the success of work on using bootstrapped data (Chaudhary et al., 2019; Sherborne

	K=0	K=1	K=2	K=4	K=8		
MLDoc	EN	96.88	-	-	-	-	
	DE	88.30	90.36 ± 1.48	90.77 ± 0.87	91.85 ± 0.83	91.98 ± 0.82	
	FR	83.05	88.94 ± 2.46	89.71 ± 1.68	90.80 ± 0.88	91.01 ± 0.94	
	ES	81.90	83.99 ± 2.35	85.65 ± 1.60	86.30 ± 1.85	88.46 ± 1.90	
	IT	74.13	74.97 ± 2.04	75.29 ± 1.57	76.43 ± 1.41	78.12 ± 1.25	
	RU	72.33	77.40 ± 4.27	80.57 ± 1.37	81.33 ± 1.33	81.91 ± 1.21	
	ZH	84.38	87.18 ± 1.45	87.31 ± 1.53	88.33 ± 1.11	88.72 ± 1.05	
	JA	74.58	76.23 ± 1.59	76.71 ± 2.12	78.60 ± 2.43	81.17 ± 1.72	
	MARC	EN	64.52	-	-	-	-
DE		49.62	51.50 ± 1.58	52.76 ± 0.87	52.78 ± 1.00	53.32 ± 0.59	
FR		47.30	49.32 ± 1.34	49.70 ± 1.43	50.64 ± 0.94	51.23 ± 0.76	
ES		48.44	49.72 ± 1.24	49.96 ± 1.12	50.45 ± 1.22	51.25 ± 0.93	
ZH		40.40	43.19 ± 1.76	44.45 ± 1.36	45.40 ± 1.26	46.40 ± 0.93	
JA		38.84	41.95 ± 2.09	43.63 ± 1.30	43.98 ± 0.89	44.44 ± 0.69	
XNLI		EN	82.67	-	-	-	-
		DE	70.32	70.58 ± 0.36	70.60 ± 0.34	70.61 ± 0.39	70.70 ± 0.50
		FR	73.57	73.41 ± 0.48	73.74 ± 0.46	73.57 ± 0.49	73.77 ± 0.44
	ES	73.71	73.84 ± 0.40	73.87 ± 0.44	73.74 ± 0.48	73.87 ± 0.46	
	RU	68.70	68.81 ± 0.52	68.76 ± 0.54	68.87 ± 0.55	68.81 ± 0.77	
	ZH	69.32	69.73 ± 0.94	69.75 ± 0.94	70.56 ± 0.76	70.62 ± 0.86	
	AR	64.97	64.75 ± 0.36	64.82 ± 0.23	64.82 ± 0.23	64.94 ± 0.37	
	BG	67.58	68.15 ± 0.69	68.19 ± 0.75	68.55 ± 0.67	68.32 ± 0.70	
	EL	65.67	65.64 ± 0.40	65.73 ± 0.36	65.80 ± 0.41	66.00 ± 0.53	
	HI	56.57	56.94 ± 0.82	57.07 ± 0.82	57.21 ± 1.14	57.82 ± 1.18	
	SW	48.08	50.33 ± 1.08	50.28 ± 1.24	51.08 ± 0.62	51.01 ± 0.79	
	TH	46.17	49.43 ± 2.60	50.08 ± 2.42	51.32 ± 2.07	52.16 ± 2.43	
	TR	60.40	61.02 ± 0.68	61.20 ± 0.61	61.35 ± 0.49	61.31 ± 0.56	
	UR	57.05	57.56 ± 0.85	57.83 ± 0.91	58.20 ± 0.93	58.67 ± 1.03	
	VI	69.82	70.04 ± 0.59	70.14 ± 0.75	70.23 ± 0.63	70.41 ± 0.70	
	PAWSX	EN	93.90	-	-	-	-
DE		83.80	84.14 ± 0.40	84.08 ± 0.42	84.04 ± 0.47	84.23 ± 0.66	
FR		86.90	87.07 ± 0.27	87.06 ± 0.37	87.03 ± 0.31	86.94 ± 0.41	
ES		88.25	87.90 ± 0.54	87.80 ± 0.56	87.84 ± 0.53	87.85 ± 0.75	
ZH		77.75	77.71 ± 0.37	77.63 ± 0.47	77.68 ± 0.51	77.82 ± 0.64	
JA		73.30	73.78 ± 0.75	73.71 ± 1.04	73.48 ± 0.69	73.79 ± 1.28	
KO		72.05	73.75 ± 1.30	73.11 ± 1.05	73.79 ± 0.92	73.31 ± 0.61	

Table 2: Zero-shot (column  $K = 0$ ) and few-shot (columns  $K > 0$ ) results (Acc. in %) on the test set for CLS tasks. Green [red]: few-shot transfer outperforms [underperforms] zero-shot transfer.

et al., 2020).

In general, we observe that: **1)** 1-shot buckets bring the largest relative performance improvement over ZS-XLT; **2)** the gains follow the increase of  $K$ , but with diminishing returns; **3)** the performance variance across the 40 buckets decreases as  $K$  increases. These observations are more pronounced for POS and NER; e.g., 1-shot EN to Urdu (UR) POS transfer shows gains of  $\approx 22 F_1$  points (52.40 with zero-shot, 74.95 with 1-shot).

For individual runs, we observe that models in FS-XLT tend to overfit the buckets quickly at small  $K$  values. For example, in around 32% of NER 1-shot buckets, the model achieves the best dev score right after the first epoch; continuing the training only degrades performance. Similar observations hold for semantic tasks like MARC, where in 10 out of 40 DE 1-shot buckets, the dev set performance peaks at epoch 1 (cf. learning curve in Appendix §D Figure 6). This suggests the necessity of running the target-adapting experiments on multiple buckets if reliable conclusions are to be drawn.

**Different Downstream Tasks.** The models for different tasks present various levels of sensitiv-

ity to FS-XLT. Among the CLS tasks that require semantic reasoning, FS-XLT benefits MLDoc the most. This is not surprising given the fact that keyword matching can largely solve MLDoc (Artetxe et al., 2020a,b): A few examples related to target language keywords are expected to significantly improve performance. FS-XLT also yields prominent gains on the Amazon review classification dataset MARC. Similar to MLDoc, we hypothesize that just matching a few important opinion and sentiment words (Liu, 2012) in the target language brings large gains already. We provide further qualitative analyses in §5.4.

XNLI and PAWSX behave differently from MLDoc and MARC. XNLI requires higher level semantic reasoning on pairs of sentences. FS-XLT performance improves modestly (XNLI) or even decreases (PAWSX-ES) compared to ZS-XLT, even with large  $K$ . PAWSX requires a model to distinguish adversarially designed non-paraphrase sentence pairs with large lexical overlap like “Flights from New York to Florida” and “Flights from Florida to New York” (Zhang et al., 2019). This poses a challenge for FS-XLT, given the small amount of target language information in the buckets. Therefore, *when buckets are small (e.g.,  $K = 1$ ) and for challenging semantic tasks like PAWSX, the buckets do not substantially help*. Annotating more shots in the target language is an intuitive solution. Designing task-specific pretraining/finetuning objectives could also be promising (Klein and Nabi, 2020; Ram et al., 2021).

Unlike CLS tasks, POS and NER benefit from FS-XLT substantially. We speculate that there are two reasons: **1)** Both tasks often require little to no high-level semantic understanding or reasoning; **2)** due to i.i.d. sampling, train/dev/test splits are likely to have overlapping vocabulary, and the labels in the buckets can easily propagate to dev and test. We delve deeper into these conjectures in §5.4.

**Different Languages.** For languages that are more distant from EN, e.g., with different scripts, small lexical overlap, or fewer common typological features (Pires et al., 2019; Wu and Dredze, 2020a), FS-XLT introduces crucial lexical and structural information to guide the update of embedding and transformer layers in mBERT.

We present several findings based on the NER and POS results for a typologically diverse language sample. Figure 2 shows that for languages with non-Latin scripts (different from EN), despite

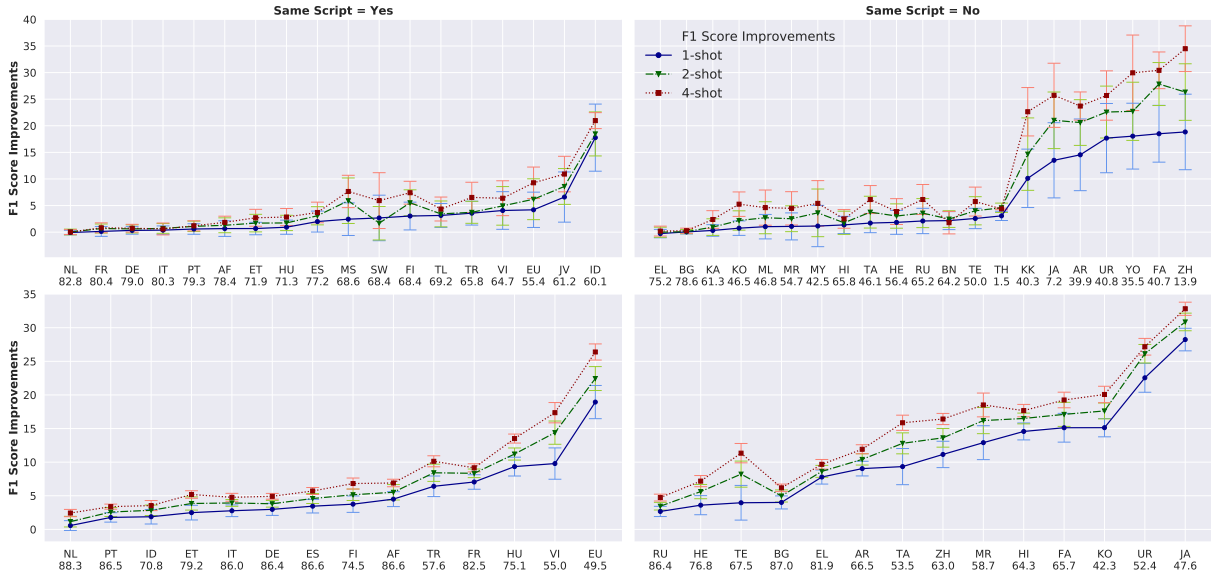


Figure 2: Improvement in  $F_1$  (mean and standard deviation) of FS-XLT over ZS-XLT (numbers shown on x-axis beneath each language) for NER (top) and POS (bottom) for three different bucket sizes. See Appendix §D (Tables 12 and 13) for absolute numerical values.

Task	Factor	S	P
NER	lexical overlap	-0.34	-0.35
	# of common linguistic features	-0.37	-0.10
POS	lexical overlap	-0.63	-0.50
	# of common linguistic features	-0.57	-0.54

Table 3: Correlations between FS-XLT  $F_1$  score gains and the two factors (lexical overlap and the number of common linguistic features with EN) when considered independently for POS and NER: S/R denotes Spearman’s/Pearson’s  $\rho$ . See Footnotes 3, 4 for information on the two factors.

their small to non-existent lexical overlap<sup>3</sup> and diverging typological features (see Appendix §D Tables 9 and 14), the performance boosts are generally larger than those in the same-script target languages: 6.2 vs. 3.0 average gain in NER and 11.4 vs. 5.4 in POS for  $K = 1$ . This clearly manifests the large information discrepancy between target-language buckets and source-language data. EN data is less relevant to these languages, so they obtain very limited gain from source-training, reflected by their low ZS-XLT scores. With a small amount of target-language knowledge in the buckets, the performance is improved dramatically, highlighting the effectiveness of FS-XLT.

Table 3 shows that, besides script form, lexical overlap and the number of linguistic features com-

<sup>3</sup>We define lexical overlap as  $\frac{|V_L \cap V_{EN}|}{|V_{EN}|}$  where  $V$  denotes vocabulary.  $|V|_L$  is computed with the 40 buckets of a target language  $L$ .

mon with EN<sup>4</sup> also contribute directly to FS-XLT performance difference among languages: There is a moderate *negative* correlation between  $F_1$  score gains vs. the two factors when considered independently for both syntactic tasks: The fewer overlaps/features a target language shares with EN, the larger the gain FS-XLT achieves.

This again stresses the importance of buckets – they contain target-language-specific knowledge about a task that cannot be obtained by ZS-XLT, which solely relies on language similarity. Interestingly, Pearson’s  $\rho$  indicates that common linguistic features are much less linearly correlated with FS-XLT gains in NER than in POS.

### 5.3 Importance of Source-Training

Table 4 reports the performance *drop* when directly carrying out target-adapting, without any prior source-training of mBERT. We show the scores for MLDoc and PAWSX as a simple and a challenging CLS task, respectively. For NER and POS, we select two high- (Russian (RU), ES), mid- (Vietnamese (VI), Turkish (TR)), and low-resource languages (Tamil (TA), Marathi (MR)) each.<sup>5</sup>

The results clearly indicate that omitting the

<sup>4</sup>Following Pires et al. (2019), we use six WALS features: 81A (Order of Subject, Object and Verb), 85A (Order of Adposition and Noun), 86A (Order of Genitive and Noun), 87A (Order of Adjective and Noun), 88A (Order of Demonstrative and Noun), and 89A (Order of Numeral and Noun).

<sup>5</sup>The categorization based on resource availability is according to WikiSize (Wu and Dredze, 2020a).

	MLDoc		PAWSX			POS		NER	
	K=1	K=8	K=1	K=8		K=1	K=4	K=1	K=4
DE	-37.73	-7.67	-31.11	-30.82	RU	-15.89	-3.20	-48.19	-35.77
FR	-38.14	-13.21	-33.02	-32.34	ES	-9.51	-0.93	-63.98	-41.53
ES	-33.69	-14.38	-33.76	-33.97	VI	-7.82	-0.36	-54.41	-41.45
IT	-33.63	-12.62	-	-	TR	-15.05	-8.08	-54.35	-34.52
RU	-30.66	-11.08	-	-	TA	-13.72	-4.40	-34.70	-24.81
ZH	-37.31	-12.57	-23.74	-23.65	MR	-11.34	-3.63	-40.10	-25.68
JA	-29.82	-14.32	-20.97	-20.82	-	-	-	-	-
KO	-	-	-19.83	-19.68	-	-	-	-	-

Table 4: Performance drop when conducting target-adapting without source-training.




Figure 3: Normalized (with softmax) Jaccard index (%) of a bucket (row) and the improved predictions achieved with 10 buckets (column).

source-training stage yields large performance drops. Even larger variance is also observed in this scenario (cf. Appendix §D Table 11). Therefore, the model indeed learns, when trained on the source language, some transferable crosslingual features that are beneficial to target languages, both for semantic and syntactic tasks.

#### 5.4 Importance of Lexical Features

We now investigate the sources of gains brought by FS-XLT over ZS-XLT.

For syntactic tasks, we take Persian (FA) POS as an example. Figure 3 visualizes the lexical overlap, measured by the Jaccard index, of 10 1-shot buckets (rows) and the improved word-label predictions introduced by target-adapting on each of the buckets (columns). In more detail, for column  $c$ , we collect the set (denoted as  $C_c$ ) of all test set words whose label is incorrectly predicted by the zero-shot model, but correctly predicted by the model trained on the  $c$ -th bucket. For row  $i$ , we denote with  $B_i$  the set of words occurring in bucket  $i$ . The figure shows in cell  $(i, k)$  the Jaccard index of  $B_i$  and  $C_k$ . The bright color (i.e., higher lexical overlap) on the diagonal reflects that the improvements

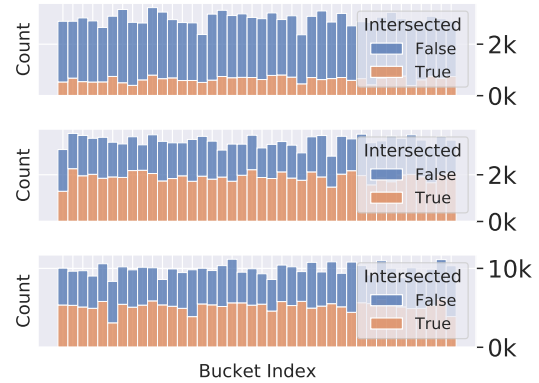


Figure 4: Improvement of word-label predictions introduced by a bucket (x-axis) in FA (top), UR (mid), and HI (bottom), in relation to the words' presence in the bucket (True or False).

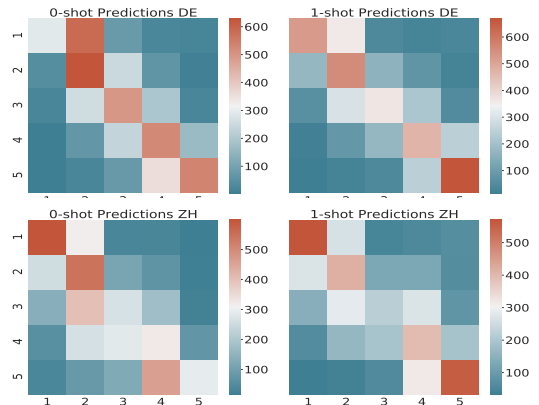


Figure 5: MARC (5 classes) test set prediction confusion matrices. Top: DE. Bottom: ZH. Left: zero-shot models. Right: 1-shot models. Colorbar numbers represent the number of instances in that cell.

introduced by a bucket are mainly<sup>6</sup> those word-label predictions that are lexically more similar to the bucket than to other buckets.

We also investigate the question: How many word-label predictions that are improved after FS-XLT occur in the bucket, i.e., in the training data? Figure 4 plots this for the 40 1-shot buckets in FA, UR, and Hindi (HI). We see that many test words do occur in the bucket (shown in orange), in line with recent findings (Lewis et al., 2021; Elangovan et al., 2021). These analyses shed light on why the buckets benefit NER/POS – which heavily rely on lexical information – more than higher level semantic tasks.

For the CLS task MARC, which requires un-

<sup>6</sup>Note that the sampled buckets for POS are not completely disjoint (cf. sampling strategy in §4).

token	[SEP]	.	nicht	!	Die	sehr
$\Delta\text{Attn}$	+4.13	+2.91	+1.84	-1.75	-0.92	-0.81

Table 5: Tokens with the highest attention change from [CLS], comparing zero-shot with a 1-shot DE bucket.

derstanding product reviews, Figure 5 visualizes the confusion matrices of test set predictions for DE and Chinese (ZH) zero- and 1-shot models; axis ticks are review scores in  $\{1, 2, 3, 4, 5\}$ . The squares on the diagonals in the two left heatmaps show that parameter initialization on EN is a good basis for well-performing ZS-XLT: This is particularly true for DE, which is linguistically closer to EN. Two extreme review scores – 1 (for DE) and 5 (for ZH) – have the largest confusions. The two right heatmaps show that improvements brought by the 1-shot buckets are mainly achieved by correctly predicting more cases of the two extreme review scores:  $2 \rightarrow 1$  (DE) and  $4 \rightarrow 5$  (ZH). But the more challenging cases (reviews with scores 2, 3, 4), which require non-trivial reasoning, are not significantly improved, or even become worse.

We inspect examples that are incorrectly predicted by the few-shot model (predicting 1), but are correctly predicted by the zero-shot model (predicting 2). Specifically, we compute the difference of where [CLS] attends to, before and after adapting the model on a 1-shot DE bucket. We extract and average attentions computed by the 12 heads from the topmost transformer layer.

Table 5 shows that “nicht” (“not”) draws high attention change from [CLS]. “Nicht” (i.e., negation) by itself is not a reliable indicator of sentiment, so giving the lowest score to reviews solely because they contain “nicht” is not a good strategy. The following review is classified as 1 by the 1-shot model, but 2 is the gold label (as the review is not entirely negative):

“Die Uhr ging nicht einmal eine Minute ... **Optisch allerdings sehr schön.**” (“The clock didn’t even work one minute ... **Visually, however, very nice.**”)

Pretrained multilingual encoders are shown to learn and store “language-agnostic” features (Pires et al., 2019; Zhao et al., 2020); §5.3 shows that source-training mBERT on EN substantially benefits other languages, even for difficult semantic tasks like PAWSX. Conditioning on such language-agnostic features, we expect that the buckets should lead to good understanding and reasoning capabilities for a target language. However, plain few-shot finetuning still relies heavily on unintended shallow

lexical cues and shortcuts (Niven and Kao, 2019; Geirhos et al., 2020) that generalize poorly. Other open research questions for future work arise: How do we overcome this excessive reliance on lexical features? How can we leverage language-agnostic features with *few shots*? Our standardized buckets, baseline results, and analyses are the initial step towards researching and answering these questions.

## 5.5 Target-Adapting Methods

SotA few-shot learning methods (Chen et al., 2019; Wang et al., 2020; Tian et al., 2020; Dhillon et al., 2020) from computer vision consist of two stages: 1) training on base-class images, and 2) few-shot finetuning using new-class images. Source-training and target-adapting stages of FS-XLT, albeit among languages, follow an approach very similar to these methods. Therefore, we test their effectiveness for crosslingual transfer. These methods are built upon cosine similarity that imparts inductive bias about distance and is more effective than a fully-connected classifier layer (FC) with small  $K$  (Wang et al., 2020). Following (Chen et al., 2019; Wang et al., 2020; Tian et al., 2020), we freeze the embedding and transformer layers of mBERT, and explore four variants of the target-adapting stage using MARC.

**COS+Pooler.** We randomly initialize a trainable weight matrix  $\mathbf{W} \in \mathbb{R}^{h \times c}$  where  $h$  is the hidden dimension size and  $c$  is the number of classes. Rewriting  $\mathbf{W}$  as  $[\mathbf{w}_1, \dots, \mathbf{w}_i, \dots, \mathbf{w}_c]$ , we compute the logits of an input sentence representation  $\mathbf{x} \in \mathbb{R}^h$  (from mBERT) belonging to class  $i$  as

$$\alpha \cdot \frac{\mathbf{x}^\top \mathbf{w}_i}{\|\mathbf{x}\|_2 \cdot \|\mathbf{w}_i\|_2},$$

where  $\alpha$  is a scaling hyperparameter, set to 10 in all experiments. During training,  $\mathbf{W}$  and mBERT’s pooler layer containing a linear layer and a  $\tanh$  non-linearity are updated.

**FC+Pooler.** During training, we update the linear classifier layer and mBERT’s pooler layer.

**FC only.** During training, we only update the linear classifier layer. This variant largely reduces model complexity and exhibit lower variance when  $K$  is small.

**FC(reset)+Pooler.** Similar to FC+Pooler, but the source-trained linear classifier layer is randomly re-initialized before training.

Table 6 shows the performance of these methods along with full model finetuning (without freezing). *FC+Pooler* performs the best among the



	Full-Model Finetuning			FC only		FC + Pooler		COS + Pooler		FC (reset) + Pooler	
	K=0	K=1	K=8	K=1	K=8	K=1	K=8	K=1	K=8	K=1	K=8
DE	49.62	51.50 ± 1.58	53.32 ± 0.59	50.82 ± 1.17	52.58 ± 0.63	<b>51.18 ± 1.13</b>	<b>53.17 ± 0.58</b>	37.98 ± 5.53	45.85 ± 2.14	38.52 ± 6.64	49.46 ± 2.21
FR	47.30	49.32 ± 1.34	51.23 ± 0.76	48.19 ± 0.78	49.05 ± 0.93	<b>48.60 ± 1.02</b>	<b>49.97 ± 0.77</b>	39.93 ± 3.50	44.41 ± 1.95	40.12 ± 5.04	47.77 ± 2.00
ES	48.44	49.72 ± 1.24	51.25 ± 0.93	49.03 ± 0.73	49.69 ± 0.57	<b>49.28 ± 0.85</b>	<b>50.21 ± 0.63</b>	40.01 ± 4.33	45.35 ± 2.37	40.89 ± 4.96	47.73 ± 2.33
ZH	40.40	43.19 ± 1.76	46.40 ± 0.93	41.90 ± 1.15	43.34 ± 0.88	<b>42.30 ± 1.37</b>	<b>44.42 ± 0.65</b>	33.10 ± 5.48	38.31 ± 1.87	31.83 ± 7.00	42.07 ± 2.19
JA	38.84	41.95 ± 2.09	44.44 ± 0.69	40.76 ± 1.76	43.14 ± 0.76	<b>41.40 ± 1.74</b>	<b>43.81 ± 0.56</b>	34.36 ± 4.19	38.95 ± 1.80	32.80 ± 5.17	41.18 ± 1.68

Table 6: Accuracy (%) on MARC when varying classifier head configurations. Full-Model Finetuning updates all parameters during training; the other four methods only update a subset as described in §5.5. The best results (excluding Full-Model Finetuning) are in bold.

four for both  $K = 1$  and  $K = 8$  in all languages. However, it underperforms the full model finetuning, especially when  $K = 8$ . *FC only* is sub-optimal; yet the decrease in comparison to *FC+Pooler* is small, highlighting that EN-trained mBERT is a strong feature extractor. *COS+Pooler* and *FC(reset)+Pooler* perform considerably worse than the other two methods and zero-shot transfer – presumably because their new parameters need to be trained from scratch with few shots.

We leave further exploration of other possibilities of exploiting crosslingual features through collapse-preventing regularization (Aghajanyan et al., 2021) or contrastive learning (Gunel et al., 2021) to future work. Integrating prompting (Brown et al., 2020; Schick and Schütze, 2020; Gao et al., 2020; Liu et al., 2021) – a strong performing few-shot learning methodology for NLP – into the crosslingual transfer learning pipeline is also a promising direction.

## 6 Conclusion and Future Work

We have presented an extensive study of *few-shot crosslingual transfer*. The focus of the study has been on an empirically detected performance variance in few-shot scenarios: The models exhibit a high level of sensitivity to the choice of few shots. We analyzed and discussed the major causes of this variance across six diverse tasks for up to 40 languages. Our results show that large language models tend to overfit to few shots quickly and mostly rely on shallow lexical features present in the few shots, though they have been trained with abundant data in English. Moreover, we have empirically validated that state-of-the-art few-shot learning methods in computer vision do not outperform a conceptually simple alternative: Full model finetuning.

Our study calls for more rigor and accurate reporting of the results of few-shot crosslingual transfer experiments. They should include score distributions over standardized and fixed few shots. To

aid this goal, we have created and provided such fixed few shots as a standardized benchmark for six multilingual datasets.

Few-shot learning is promising for crosslingual transfer, because it mirrors how people acquire new languages, and that the few-shot data annotation is feasible. In future work, we will investigate more sophisticated techniques and extend the work to more NLP tasks.

## Acknowledgments

This work was funded by the European Research Council: ERC NonSequeToR (#740516) and ERC LEXICAL (#648909). We thank the anonymous reviewers and Fei Mi for their helpful suggestions.

## References

- Armen Aghajanyan, Akshat Shrivastava, Anchit Gupta, Naman Goyal, Luke Zettlemoyer, and Sonal Gupta. 2021. [Better fine-tuning by reducing representational collapse](#). In *International Conference on Learning Representations*.
- 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.
- Mikel Artetxe and Holger Schwenk. 2019. [Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond](#). *Transactions of the Association for Computational Linguistics*, 7:597–610.
- Trapit Bansal, Rishikesh Jha, and Andrew McCallum. 2020. [Learning to few-shot learn across diverse](#)

- natural language classification tasks. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5108–5123, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Aditi Chaudhary, Jiateng Xie, Zaid Sheikh, Graham Neubig, and Jaime Carbonell. 2019. [A little annotation does a lot of good: A study in bootstrapping low-resource named entity recognizers](#). 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 5164–5174, Hong Kong, China. Association for Computational Linguistics.
- Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. 2019. [A closer look at few-shot classification](#). In *International Conference on Learning Representations*.
- Yinbo Chen, Xiaolong Wang, Zhuang Liu, Huijuan Xu, and Trevor Darrell. 2020. [A new meta-baseline for few-shot learning](#). *arXiv preprint arXiv:2003.04390*.
- 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 7059–7069.
- 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*. 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.
- Guneet Singh Dhillon, Pratik Chaudhari, Avinash Ravichandran, and Stefano Soatto. 2020. [A baseline for few-shot image classification](#). In *International Conference on Learning Representations*.
- Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah Smith. 2020. [Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping](#).
- Xin Dong and Gerard de Melo. 2019. [A robust self-learning framework for cross-lingual text classification](#). 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 6306–6310, Hong Kong, China. Association for Computational Linguistics.
- Rotem Dror, Segev Shlomov, and Roi Reichart. 2019. [Deep dominance - how to properly compare deep neural models](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2773–2785, Florence, Italy. Association for Computational Linguistics.
- Julian Eisenschlos, Sebastian Ruder, Piotr Czapla, Marcin Kadras, Sylvain Gugger, and Jeremy Howard. 2019. [MultiFiT: Efficient multi-lingual language model fine-tuning](#). 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 5702–5707, Hong Kong, China. Association for Computational Linguistics.
- Aparna Elangovan, Jiayuan He, and Karin Verspoor. 2021. [Memorization vs. generalization : Quantifying data leakage in NLP performance evaluation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1325–1335, Online. Association for Computational Linguistics.
- L. Fei-Fei, R. Fergus, and P. Perona. 2006. [One-shot learning of object categories](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(4):594–611.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. [Model-agnostic meta-learning for fast adaptation of deep networks](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Re-*

- search, pages 1126–1135, International Convention Centre, Sydney, Australia. PMLR.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2020. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*.
- Tianyu Gao, Xu Han, Hao Zhu, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2019. **FewRel 2.0: Towards more challenging few-shot relation classification**. 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 6251–6256, Hong Kong, China. Association for Computational Linguistics.
- Dan Garrette and Jason Baldridge. 2013. **Learning a part-of-speech tagger from two hours of annotation**. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 138–147, Atlanta, Georgia. Association for Computational Linguistics.
- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. 2020. **Shortcut learning in deep neural networks**. *Nature Machine Intelligence*, 2(11):665–673.
- Beliz Gunel, Jingfei Du, Alexis Conneau, and Veselin Stoyanov. 2021. **Supervised contrastive learning for pre-trained language model fine-tuning**. In *International Conference on Learning Representations*.
- Aakriti Gupta, Kapil Thadani, and Neil O’Hare. 2020. **Effective few-shot classification with transfer learning**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1061–1066, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Xu Han, Hao Zhu, Pengfei Yu, Ziyun Wang, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2018. **FewRel: A large-scale supervised few-shot relation classification dataset with state-of-the-art evaluation**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4803–4809, Brussels, Belgium. Association for Computational Linguistics.
- Han He and Jinho D. Choi. 2020. **Establishing Strong Baselines for the New Decade: Sequence Tagging, Syntactic and Semantic Parsing with BERT**. In *Proceedings of the 33rd International Florida Artificial Intelligence Research Society Conference, FLAIRS’20*. Best Paper Candidate.
- Michael A. Hedderich, David Adelani, Dawei Zhu, Jesujoba Alabi, Udia Markus, and Dietrich Klakow. 2020. **Transfer learning and distant supervision for multilingual transformer models: A study on African languages**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2580–2591, Online. Association for Computational Linguistics.
- Nithin Holla, Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2020. **Learning to learn to disambiguate: Meta-learning for few-shot word sense disambiguation**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4517–4533, Online. Association for Computational Linguistics.
- Yutai Hou, Wanxiang Che, Yongkui Lai, Zhihan Zhou, Yijia Liu, Han Liu, and Ting Liu. 2020a. **Few-shot slot tagging with collapsed dependency transfer and label-enhanced task-adaptive projection network**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1381–1393. Association for Computational Linguistics.
- Yutai Hou, Jiafeng Mao, Yongkui Lai, Cheng Chen, Wanxiang Che, Zhigang Chen, and Ting Liu. 2020b. **Fewjoint: A few-shot learning benchmark for joint language understanding**. *CoRR*, abs/2009.08138.
- Tsung-Yuan Hsu, Chi-Liang Liu, and Hung-yi Lee. 2019. **Zero-shot reading comprehension by cross-lingual transfer learning with multi-lingual language representation model**. 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 5933–5940, Hong Kong, China. 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*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421, Virtual. PMLR.
- Phillip Keung, Yichao Lu, Julian Salazar, and Vikas Bhardwaj. 2020a. **Don’t use English dev: On the zero-shot cross-lingual evaluation of contextual embeddings**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 549–554, Online. Association for Computational Linguistics.
- Phillip Keung, Yichao Lu, György Szarvas, and Noah A. Smith. 2020b. **The multilingual Amazon reviews corpus**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4563–4568, Online. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. **Adam: A method for stochastic optimization**. In *ICLR (Poster)*.
- Tassilo Klein and Moin Nabi. 2020. **Contrastive self-supervised learning for commonsense reasoning**. In

- Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7517–7523, Online. Association for Computational Linguistics.
- Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. 2015. Siamese neural networks for one-shot image recognition. In *ICML 2015 Deep Learning Workshop*.
- Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. 2020. [From zero to hero: On the limitations of zero-shot language transfer with multilingual transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4483–4499, Online. Association for Computational Linguistics.
- Patrick Lewis, Pontus Stenetorp, and Sebastian Riedel. 2021. [Question and answer test-train overlap in open-domain question answering datasets](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1000–1008, Online. Association for Computational Linguistics.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. [XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, Online. Association for Computational Linguistics.
- Bing Liu. 2012. Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies*, 5(1):1–167.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021. Gpt understands, too. *arXiv preprint arXiv:2103.10385*.
- Erik G Miller, Nicholas E Matsakis, and Paul A Viola. 2000. Learning from one example through shared densities on transforms. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No. PR00662)*, volume 1, pages 464–471. IEEE.
- Marius Mosbach, Maksym Andriushchenko, and Dietrich Klakow. 2021. [On the stability of fine-tuning {bert}: Misconceptions, explanations, and strong baselines](#). In *International Conference on Learning Representations*.
- Timothy Niven and Hung-Yu Kao. 2019. [Probing neural network comprehension of natural language arguments](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4658–4664, Florence, Italy. Association for Computational Linguistics.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D. Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. [Universal Dependencies v2: An evergrowing multilingual treebank collection](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4034–4043, Marseille, France. European Language Resources Association.
- 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.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems*, volume 32, pages 8026–8037. Curran Associates, Inc.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Jason Phang, Thibault Févry, and Samuel R Bowman. 2018. Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks. *arXiv preprint arXiv:1811.01088*.
- 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.
- Yada Pruksachatkun, Jason Phang, Haokun Liu, Phu Mon Htut, Xiaoyi Zhang, Richard Yuanzhe Pang, Clara Vania, Katharina Kann, and Samuel R. Bowman. 2020. [Intermediate-task transfer learning with pretrained language models: When and why does it work?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5231–5247, Online. Association for Computational Linguistics.
- Afshin Rahimi, Yuan Li, and Trevor Cohn. 2019. [Massively multilingual transfer for NER](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 151–164, Florence, Italy. Association for Computational Linguistics.

- Ori Ram, Yuval Kirstain, Jonathan Berant, Amir Globerson, and Omer Levy. 2021. Few-shot question answering by pretraining span selection. *arXiv preprint arXiv:2101.00438*.
- Nils Reimers and Iryna Gurevych. 2017. Reporting score distributions makes a difference: Performance study of LSTM-networks for sequence tagging. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 338–348, Copenhagen, Denmark. Association for Computational Linguistics.
- Annette Rios, Mathias Müller, and Rico Sennrich. 2020. Subword segmentation and a single bridge language affect zero-shot neural machine translation. In *Proceedings of the Fifth Conference on Machine Translation*, pages 526–535, Online. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2020. It’s not just size that matters: Small language models are also few-shot learners. *arXiv preprint arXiv:2009.07118*.
- Holger Schwenk and Xian Li. 2018. A corpus for multilingual document classification in eight languages. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Paris, France. European Language Resources Association (ELRA).
- Tom Sherborne, Yumo Xu, and Mirella Lapata. 2020. Bootstrapping a crosslingual semantic parser. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 499–517, Online. Association for Computational Linguistics.
- Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, volume 30, pages 4077–4087. Curran Associates, Inc.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 2020. Dataset cartography: Mapping and diagnosing datasets with training dynamics. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9275–9293, Online. Association for Computational Linguistics.
- Yonglong Tian, Yue Wang, Dilip Krishnan, Joshua B. Tenenbaum, and Phillip Isola. 2020. Rethinking few-shot image classification: A good embedding is all you need? In *Computer Vision – ECCV 2020*, pages 266–282, Cham. Springer International Publishing.
- Orith Toledo-Ronen, Matan Orbach, Yonatan Bilu, Artem Spector, and Noam Slonim. 2020. Multilingual argument mining: Datasets and analysis. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 303–317, Online. Association for Computational Linguistics.
- Xin Wang, Thomas Huang, Joseph Gonzalez, Trevor Darrell, and Fisher Yu. 2020. Frustratingly simple few-shot object detection. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9919–9928, Virtual. PMLR.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. 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 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- 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.
- Shijie Wu and Mark Dredze. 2020a. Are all languages created equal in multilingual bert? In *Proceedings of the 5th Workshop on Representation Learning for NLP, ReplANLP@ACL 2020, Online, July 9, 2020*, pages 120–130. Association for Computational Linguistics.
- Shijie Wu and Mark Dredze. 2020b. Do explicit alignments robustly improve multilingual encoders? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4471–4482, Online. Association for Computational Linguistics.
- 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, Łukasz 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.
- Patrick Xia, Shijie Wu, and Benjamin Van Durme. 2020. Which \*BERT? A survey organizing contextualized encoders. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7516–7533, Online. Association for Computational Linguistics.

- 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.
- Wenpeng Yin. 2020. Meta-learning for few-shot natural language processing: A survey.
- Wenpeng Yin, Nazneen Fatema Rajani, Dragomir Radev, Richard Socher, and Caiming Xiong. 2020. Universal natural language processing with limited annotations: Try few-shot textual entailment as a start. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8229–8239, Online. Association for Computational Linguistics.
- Dani Yogatama, Cyprien de Masson d’Autume, Jerome Connor, Tomas Kocisky, Mike Chrzanowski, Lingpeng Kong, Angeliki Lazaridou, Wang Ling, Lei Yu, Chris Dyer, and Phil Blunsom. 2019. Learning and evaluating general linguistic intelligence.
- Mo Yu, Xiaoxiao Guo, Jinfeng Yi, Shiyu Chang, Saloni Potdar, Yu Cheng, Gerald Tesauro, Haoyu Wang, and Bowen Zhou. 2018. Diverse few-shot text classification with multiple metrics. 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 1206–1215, New Orleans, Louisiana. Association for Computational Linguistics.
- Tianyi Zhang, Felix Wu, Arzoo Katiyar, Kilian Q Weinberger, and Yoav Artzi. 2021. Revisiting few-sample {bert} fine-tuning. In *International Conference on Learning Representations*.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019. PAWS: Paraphrase adversaries from word scrambling. 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 1298–1308, Minneapolis, Minnesota. Association for Computational Linguistics.
- Wei Zhao, Steffen Eger, Johannes Bjerva, and Isabelle Augenstein. 2020. Inducing language-agnostic multilingual representations. *arXiv preprint arXiv:2008.09112*.

## A Reproducibility Checklist

### A.1 mBERT Architecture and Number of Parameters

We use the “bert-base-multilingual-cased” model<sup>7</sup>. It contains 12 Transformer blocks with 768 hidden dimensions. Each block has 12 self attention heads. The model is pretrained on the concatenation of the Wikipedia dump of 104 languages.

There are about 179 million parameters in mBERT. For all the tasks, we use a linear output layer. Denoting the output dimension of a task as  $m$ , e.g.,  $m = 2$  for PAWSX. Then we have in total 179 million +  $768 \times m + m$  parameters for the task.

### A.2 Computing Infrastructure

All experiments are conducted on GeForce GTX 1080Ti. In the source-training stage, we use 4 GPUs with per-GPU batch size 32. In the target-adapting stage, we use a single GPU and the batch size is equal to the number of examples in a bucket.

### A.3 Evaluation Metrics and Validation Performance

We follow the standard evaluation metrics used in XTREME (Hu et al., 2020) and they are shown in Table 1; evaluation functions in scikit-learn (Pedregosa et al., 2011) and seqeval (<https://github.com/chakki-works/seqeval>) are used. Link to code: code/utils/eval\_meters.py.

The validation performance of the English-trained models are shown in the first row of Table 7; the optimal learning rate for each task is shown in the second row.

MLDoc	MARC	XNLI	PAWSX	POS	NER
98.1	65.1	83.5	94.5	95.6	84.3
1e-5	1e-5	3e-5	1e-5	1e-5	1e-5

Table 7: Source-training validation performance (%) and the optimal learning rate.

For all the FS-XLT experiments, we enclosed the validation scores in <https://github.com/fsxlt/running-logs>.

### A.4 Hyperparameter Search

For both source-training and target-adapting, the only hyperparameter we search is learning rate (from  $\{1e-5, 3e-5, 5e-5, 7e-5\}$ ) to reduce

<sup>7</sup><https://github.com/google-research/bert/blob/master/multilingual.md>

---

### Algorithm 1: Minimum-including

---

**Require:** # of shot  $K$ , language data  $\mathcal{D}$ , label set  $\mathcal{L}_{\mathcal{D}}$   
1: Initialize a bucket  $\mathcal{S} = \{\}$ ,  $\text{Count}_{\ell_j} = 0$  ( $\forall \ell_j \in \mathcal{L}_{\mathcal{D}}$ )  
2: **for**  $\ell$  in  $\mathcal{L}_{\mathcal{D}}$  **do**  
    **while**  $\text{Count}_{\ell} < K$  **do**  
        From  $\mathcal{D}$ , randomly sample a  
         $(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$  pair that  $\mathbf{y}^{(i)}$  includes  $\ell$   
        Add  $(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$  to  $\mathcal{S}$   
        Update all  $\text{Count}_{\ell_j}$  ( $\forall \ell_j \in \mathcal{L}_{\mathcal{D}}$ )  
3: **for each**  $(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$  in  $\mathcal{S}$  **do**  
    Remove  $(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$  from  $\mathcal{S}$   
    Update all  $\text{Count}_{\ell_j}$  ( $\forall \ell_j \in \mathcal{L}_{\mathcal{D}}$ )  
    **if any**  $\text{Count}_{\ell_j} < K$  **then**  
        Put  $(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$  back to  $\mathcal{S}$   
        Update all  $\text{Count}_{\ell_j}$  ( $\forall \ell_j \in \mathcal{L}_{\mathcal{D}}$ )  
4: Return  $\mathcal{S}$ 

---

the sensitivity of our results to hyperparameter selection.

### A.5 Datasets and Preprocessing

For tasks (XNLI, PAWSX, POS, NER) covered in XTREME (Hu et al., 2020), we utilize the provided preprocessed datasets. Our MLDoc dataset is obtained from <https://github.com/facebookresearch/MLDoc>. We retrieve MARC from [docs.opendata.aws/amazon-reviews-ml/readme.html](https://docs.opendata.aws/amazon-reviews-ml/readme.html). Table 8 shows example entries of the datasets. It is worth noting that MARC is a single sentence review classification task, however, we put the “review title” and “product category” in the “Text B” field, following Keung et al. (2020b).

We utilize the tokenizer in the HuggingFace Transformers package (Wolf et al., 2019) to preprocess all the texts. In all experiments, we use 128 maximum sequence length and truncate from the end of a sentence if its length exceeds the limit.

## B Languages

We work on 40 languages in total. They are shown in Table 9, together with their ISO 639-1 codes, writing script, and language features from WALS (<https://wals.info/>) used in our experiments.

## C Minimum-Including Algorithm

We utilize the *Minimum-including Algorithm* from Hou et al. (2020a,b) for sampling the buckets of POS and NER which have several labels in a sentence. Denoting as  $\mathbf{x}$  a sentence that consists of an array of words  $(x_1, \dots, x_n)$ , and the array  $\mathbf{y}$  that consists of a series of labels  $(y_1, \dots, y_n)$ . We sample the buckets by using Algorithm 1. Note that we

MARC	Text A	Très mignons et de bonne qualité. La figurine est assez imposante mais conforme à la taille indiquée dans le descriptif.
	Text B	Jolis détails . home
XNLI	Text A	Ich musste anfangen Seminare zu belegen .
	Text B	Ich brauchte keine Vorbereitung .
PAWSX	Text A	Lo entrenó John Velázquez y en sus carreras más importantes lo montó el jinete Dale Romans.
	Text B	Lo entrenó John Velázquez, y el jinete Dale Romans lo montó en las carreras más importantes.
POS	Text A	(Lo,PRON), (sanno,VERB), (oramai,ADV), (quasi,ADV), (tutti,PRON), (che,SCONJ), (un,DET), (respiro,NOUN), (affannoso,ADJ) ...
NER	Text A	(Sempat,O), (pindah,O), (ke,O), (HJK,B-ORG), (dan,O), (1899,B-ORG), (Hoffenheim,I-ORG), (yang,O), (meminjamkannya,O), (ke,O) ...

Table 8: Example entries of the datasets. We convert the raw text to the mBERT format “Text A” and “Text B” (Devlin et al., 2019). For POS and NER, we list (word, tag) pairs in the sentence. Following Schwenk and Li (2018), we provide document indices of MLDoc for retrieving the documents from RCV1 and RCV2.

Language	Writing Script	81A Order of Subject, Object and Verb	85A Order of adposition and noun	86A Order of genitive and noun	87A Order of adjective and noun	88A Order of demonstrative and noun	89A Order of numeral and noun
English (EN)	Latin	SVO	Prepositions	No dominant order	Adjective-noun	Demonstrative-noun	Numeral-noun
Afrikaans (AF)	Latin	-	-	-	-	-	-
Arabic (AR)	Arabic	VSO	Prepositions	Noun-genetive	Noun-adjective	Demonstrative-noun	Numeral-noun
Bulgarian (BG)	Cyrillic	SVO	Prepositions	No dominant order	Adjective-noun	Demonstrative-noun	Numeral-noun
Bengali (BN)	Brahmic	SOV	-	-	-	-	-
German (DE)	Latin	No dominant order	Prepositions	Noun-genetive	Adjective-noun	Demonstrative-noun	Numeral-noun
Greek (EL)	Greek	No dominant order	Prepositions	Noun-genetive	Adjective-noun	Demonstrative-noun	Numeral-noun
Spanish (ES)	Latin	SVO	Prepositions	Noun-genetive	Noun-adjective	Demonstrative-noun	Numeral-noun
Estonian (ET)	Latin	SVO	Postpositions	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun
Basque (EU)	Latin	SOV	Postpositions	Genetive-noun	Noun-adjective	Noun-demonstrative	Numeral-noun
Persian (FA)	Perso-Arabic	SOV	Prepositions	Noun-genetive	Noun-adjective	Demonstrative-noun	Numeral-noun
Finnish (FI)	Latin	SVO	Postpositions	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun
French (FR)	Latin	SVO	Prepositions	Noun-genetive	Noun-adjective	Demonstrative-noun	Numeral-noun
Hebrew (HE)	Hebrew	SVO	Prepositions	Noun-genetive	Noun-adjective	Noun-demonstrative	Numeral-noun
Hindi (HI)	Devanagari	SOV	Postpositions	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun
Hungarian (HU)	Latin	No dominant order	Postpositions	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun
Indonesian (ID)	Latin	SVO	Prepositions	Noun-genetive	Noun-adjective	Noun-demonstrative	Numeral-noun
Italian (IT)	Latin	SVO	Prepositions	Noun-genetive	Noun-adjective	Demonstrative-noun	Numeral-noun
Japanese (JA)	Ideograms	SOV	Postpositions	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun
Javanese (JV)	Latin	-	-	-	-	-	-
Georgian (KA)	Georgian	SOV	Postpositions	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun
Kazakh (KK)	Cyrillic	-	-	-	-	-	-
Korean (KO)	Hangul	SOV	Postpositions	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun
Malayalam (ML)	Brahmic	SOV	-	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun
Marathi (MR)	Devanagari	SOV	Postpositions	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun
Malay (MS)	Latin	-	-	-	-	-	-
Burmese (MY)	Brahmic	SOV	Postpositions	Genetive-noun	Noun-adjective	Demonstrative-noun	Noun-numeral
Dutch (NL)	Latin	No dominant order	Prepositions	Noun-genetive	Adjective-noun	Demonstrative-noun	Numeral-noun
Portuguese (PT)	Latin	SVO	Prepositions	Noun-genetive	Noun-adjective	Demonstrative-noun	-
Russian (RU)	Cyrillic	SVO	Prepositions	Noun-genetive	Adjective-noun	Demonstrative-noun	Numeral-noun
Swahili (SW)	Latin	SVO	Prepositions	Noun-genetive	Noun-adjective	Noun-demonstrative	Noun-numeral
Tamil (TA)	Brahmic	SOV	Postpositions	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun
Telugu (TE)	Brahmic	SOV	Postpositions	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun
Thai (TH)	Brahmic	SVO	Prepositions	Noun-genetive	Noun-adjective	Noun-demonstrative	Noun-numeral
Tagalog (TL)	Latin	VSO	-	Noun-genetive	No dominant order	Mixed	Numeral-noun
Turkish (TR)	Latin	SOV	Postpositions	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun
Urdu (UR)	Perso-Arabic	SOV	Postpositions	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun
Vietnamese (VI)	Latin	SVO	Prepositions	Noun-genetive	Noun-adjective	Noun-demonstrative	Numeral-noun
Yoruba (YO)	Latin	SVO	Prepositions	Noun-genetive	Noun-adjective	Noun-demonstrative	Noun-numeral
Chinese (ZH)	Chinese ideograms	SVO	No dominant order	Genetive-noun	Adjective-noun	Demonstrative-noun	Numeral-noun

Table 9: All languages for the experiments along with their ISO 639-1 codes, writing script, and linguistic features. “-” denotes lacking feature information from WALS.

sample with replacement for POS and NER.

## D Additional Results

### D.1 Learning Curve

Figure 6 visualizes the averaged learning curve of 10 out of 40 German 1-shot MARC buckets for which the best dev performance is obtained at epoch 1.

### D.2 Numerical Values

The numerical values of the POS and NER FS-XLT results are shown in Table 13 and Table 12. The absolute performances of few-shot transfer without English source-training are shown in Table 11. The lexical overlap of target languages with EN for NER and POS is shown in Table 14.

292	584	78	27	19	526	361	43	31	40
45	630	250	64	11	176	554	162	80	28
24	259	497	196	24	65	298	369	218	50
4	69	237	525	165	22	87	176	471	244
6	25	75	357	537	16	27	42	245	670
599	316	36	33	16	570	262	45	56	67
255	543	112	70	20	269	416	126	125	65
136	401	266	174	23	143	284	219	270	84
60	262	283	322	73	63	163	190	395	189
38	83	127	462	290	32	39	59	314	555

Table 10: Numerical value of the confusion matrices in Figure 5. For 1-shot confusion matrices (right), we average results of 5 buckets and then round to integers.



	MLDoc		PAWSX			POS		NER	
	K=1	K=8	K=1	K=8		K=1	K=4	K=1	K=4
DE	52.63 ± 8.98	84.31 ± 3.60	53.03 ± 1.67	53.41 ± 1.47	RU	73.18 ± 4.42	86.65 ± 1.32	19.11 ± 6.94	35.57 ± 6.23
FR	50.80 ± 8.50	77.80 ± 4.44	54.05 ± 1.33	54.60 ± 0.97	ES	80.54 ± 4.17	90.26 ± 0.99	15.21 ± 5.98	39.37 ± 5.33
ES	50.30 ± 8.30	74.08 ± 6.48	54.14 ± 1.53	53.88 ± 1.72	VI	56.97 ± 5.16	72.00 ± 1.99	14.36 ± 4.28	29.63 ± 5.55
IT	41.34 ± 6.82	65.50 ± 4.21	-	-	TR	48.96 ± 3.15	59.65 ± 1.83	15.02 ± 5.58	37.81 ± 5.63
RU	46.74 ± 9.48	70.83 ± 5.63	-	-	TA	49.12 ± 4.67	64.96 ± 2.16	13.11 ± 4.55	27.42 ± 4.82
ZH	49.87 ± 10.44	76.15 ± 5.10	53.97 ± 1.79	54.17 ± 1.38	MR	60.26 ± 5.72	73.58 ± 2.39	15.68 ± 7.09	33.50 ± 6.02
JA	46.41 ± 6.59	66.85 ± 6.54	52.81 ± 0.96	52.97 ± 1.15	-	-	-	-	-
KO	-	-	53.92 ± 0.78	53.63 ± 0.99	-	-	-	-	-

Table 11: Target-adapting results without source-training. Numbers are mean and standard deviation of 40 runs.

	K=0	K=1	K=2	K=4
EN	95.39	-	-	-
AF	86.60	91.10 ± 1.11	92.12 ± 1.15	93.50 ± 0.56
AR	66.55	75.64 ± 1.09	77.01 ± 0.84	78.52 ± 0.67
BG	87.02	91.01 ± 0.97	91.97 ± 0.90	93.18 ± 0.56
DE	86.38	89.38 ± 0.90	90.21 ± 0.50	91.32 ± 0.43
EL	81.89	89.69 ± 1.05	90.53 ± 0.89	91.58 ± 0.72
ES	86.64	90.05 ± 1.01	91.19 ± 0.74	92.31 ± 0.52
ET	79.17	81.69 ± 1.09	83.05 ± 0.98	84.39 ± 0.56
EU	49.51	68.44 ± 2.47	71.94 ± 1.78	75.89 ± 1.20
FA	65.73	80.82 ± 2.14	82.81 ± 1.79	84.95 ± 1.16
FI	74.49	78.25 ± 1.22	79.65 ± 0.85	81.32 ± 0.82
FR	82.54	89.55 ± 1.08	90.84 ± 0.64	91.66 ± 0.60
HE	76.79	80.40 ± 1.42	82.42 ± 1.06	83.98 ± 0.83
HI	64.29	78.87 ± 1.26	80.80 ± 0.80	81.97 ± 0.92
HU	75.10	84.44 ± 1.40	86.31 ± 0.90	88.61 ± 0.67
ID	70.80	72.68 ± 1.08	73.64 ± 0.78	74.34 ± 0.75
IT	85.97	88.77 ± 0.87	89.93 ± 0.50	90.77 ± 0.59
JA	47.60	75.84 ± 1.68	78.46 ± 1.31	80.42 ± 0.98
KO	42.29	57.43 ± 1.36	59.92 ± 1.18	62.37 ± 1.22
MR	58.70	71.60 ± 2.52	74.89 ± 1.95	77.21 ± 1.77
NL	88.35	88.97 ± 0.73	89.55 ± 0.79	90.83 ± 0.54
PT	86.45	88.18 ± 0.70	88.98 ± 0.66	89.78 ± 0.38
RU	86.36	89.07 ± 0.76	89.85 ± 0.57	91.13 ± 0.51
TA	53.51	62.84 ± 2.69	66.30 ± 1.56	69.36 ± 1.13
TE	67.48	71.46 ± 2.58	75.72 ± 1.94	78.84 ± 1.44
TR	57.58	64.01 ± 1.53	66.02 ± 1.28	67.73 ± 0.82
UR	52.40	74.95 ± 2.15	78.53 ± 1.38	79.57 ± 1.24
VI	54.96	64.79 ± 2.33	69.39 ± 1.73	72.36 ± 1.51
ZH	63.01	74.15 ± 1.96	76.62 ± 1.39	79.42 ± 0.83

Table 12: Zero- (column K=0) and few- (columns K>0) shot cross-lingual transfer results (%) on POS test set.

	K=0	K=1	K=2	K=4
EN	83.65	-	-	-
AF	78.36	79.07 ± 1.47	79.69 ± 1.40	80.24 ± 1.16
AR	39.91	54.44 ± 6.74	60.51 ± 4.30	63.61 ± 2.65
BG	78.59	78.65 ± 0.38	78.70 ± 0.39	78.87 ± 0.48
BN	64.17	66.37 ± 1.69	66.66 ± 1.57	65.98 ± 2.11
DE	79.00	79.33 ± 0.71	79.61 ± 0.76	79.74 ± 0.73
EL	75.20	74.93 ± 0.79	75.18 ± 0.95	75.40 ± 0.93
ES	77.16	79.19 ± 1.97	80.28 ± 1.71	80.90 ± 1.94
ET	71.88	72.58 ± 1.17	73.60 ± 1.65	74.60 ± 1.59
EU	55.35	59.60 ± 3.32	61.59 ± 3.84	64.68 ± 2.96
FA	40.73	59.20 ± 5.34	68.55 ± 4.04	71.13 ± 3.45
FI	68.43	71.43 ± 2.61	73.92 ± 2.44	75.81 ± 2.15
FR	80.38	80.54 ± 0.93	81.08 ± 0.85	81.22 ± 0.93
HE	56.36	58.24 ± 2.25	59.43 ± 2.29	60.27 ± 2.43
HI	65.84	67.16 ± 1.61	67.56 ± 2.18	68.29 ± 1.76
HU	71.28	72.23 ± 1.33	73.03 ± 1.44	74.14 ± 1.61
ID	60.10	77.87 ± 6.31	78.57 ± 4.14	81.07 ± 1.50
IT	80.30	80.68 ± 0.79	81.00 ± 0.92	80.90 ± 1.12
JA	7.16	20.71 ± 7.07	28.23 ± 5.32	32.93 ± 6.03
JV	61.18	67.80 ± 4.72	69.79 ± 3.37	72.12 ± 3.34
KA	61.26	61.62 ± 1.09	62.25 ± 1.56	63.68 ± 1.66
KK	40.29	50.42 ± 5.49	54.97 ± 6.81	62.94 ± 4.55
KO	46.50	47.25 ± 1.36	48.69 ± 1.82	51.76 ± 2.30
ML	46.77	47.83 ± 2.30	49.51 ± 3.01	51.41 ± 3.31
MR	54.70	55.78 ± 2.54	57.22 ± 2.43	59.18 ± 3.13
MS	68.61	71.04 ± 3.07	74.51 ± 4.28	76.25 ± 3.04
MY	42.45	43.55 ± 3.88	46.03 ± 4.48	47.81 ± 4.28
NL	82.77	82.73 ± 0.43	82.83 ± 0.54	82.82 ± 0.46
PT	79.28	79.89 ± 0.99	80.39 ± 0.98	80.49 ± 0.95
RU	65.20	67.30 ± 2.38	68.78 ± 2.73	71.34 ± 2.82
SW	68.36	71.07 ± 4.28	70.08 ± 3.15	74.33 ± 5.25
TA	46.12	47.81 ± 1.81	49.86 ± 2.99	52.23 ± 2.63
TE	50.02	52.57 ± 1.91	54.02 ± 2.65	55.75 ± 2.72
TH	1.53	4.56 ± 4.87	6.08 ± 4.88	5.87 ± 4.14
TL	69.23	72.34 ± 2.25	72.63 ± 2.43	73.55 ± 2.25
TR	65.78	69.37 ± 2.24	69.53 ± 2.07	72.33 ± 2.85
UR	40.77	58.48 ± 6.51	63.38 ± 4.88	66.49 ± 4.64
VI	64.67	68.77 ± 3.54	69.64 ± 3.63	71.08 ± 3.28
YO	35.48	53.55 ± 6.19	58.22 ± 5.47	65.46 ± 7.10
ZH	13.95	32.84 ± 7.10	40.34 ± 5.32	48.49 ± 4.30

Table 13: Zero- (column K=0) and few- (columns K>0) shot cross-lingual transfer results (%) on NER test set.




Figure 6: Early stopped 1-shot transfer (EN → DE) learning curve. The English-trained model overfits the 1-shot bucket quickly, showing decreasing dev performance during training.

	NER			POS		
	K=1	K=2	K=4	K=1	K=2	K=4
AF	4.54	8.75	13.44	4.97	6.11	7.90
AR	0.65	0.95	1.57	3.51	4.49	5.30
BG	0.98	2.19	3.23	-	-	-
BN	0.39	0.77	0.80	-	-	-
DE	8.75	13.20	20.61	9.36	15.33	21.48
EL	1.45	1.84	3.59	1.96	2.87	3.04
ES	6.29	10.59	19.66	10.00	17.53	22.63
ET	4.80	5.96	11.24	5.81	9.22	13.17
EU	3.77	5.55	12.31	2.60	3.45	4.69
FA	0.27	0.44	1.01	0.37	0.37	0.41
FI	5.61	9.05	15.66	4.59	7.03	8.78
FR	6.26	10.83	19.01	15.60	25.23	37.39
HE	0.86	1.90	3.23	1.22	1.93	2.26
HI	0.95	1.16	1.99	0.44	0.27	0.51
HU	5.07	9.19	14.35	3.18	3.92	4.15
ID	5.34	9.82	16.94	9.39	13.78	21.75
IT	7.89	10.94	21.27	11.99	16.15	21.35
JA	1.75	2.02	2.14	2.60	3.68	5.00
JV	2.49	3.05	3.44	-	-	-
KA	1.99	4.00	5.78	-	-	-
KK	0.89	1.22	2.11	-	-	-
KO	1.48	1.54	3.32	2.33	3.85	5.67
ML	0.36	1.04	1.30	-	-	-
MR	0.53	0.56	0.71	0.24	0.24	0.24
MS	4.86	7.44	13.70	-	-	-
MY	0.21	0.36	0.42	-	-	-
NL	7.18	10.65	20.14	7.94	11.42	16.79
PT	6.29	11.00	19.13	8.88	13.38	20.13
RU	1.60	2.34	3.77	4.15	6.11	9.32
SW	5.90	8.10	12.37	-	-	-
TA	0.65	1.54	2.08	1.32	1.28	1.62
TE	0.77	0.80	1.19	0.20	0.20	0.20
TH	1.63	1.87	2.08	-	-	-
TL	4.83	8.96	14.98	-	-	-
TR	4.89	8.48	16.43	2.09	2.26	3.01
UR	0.30	0.27	0.68	0.74	1.35	2.16
VI	4.33	8.39	13.41	1.62	2.16	2.90
YO	1.90	2.58	2.88	-	-	-
ZH	1.81	1.99	2.14	3.04	4.86	7.33

Table 14: Lexical overlap (per-mille) of target languages with EN for NER and POS using different K-shot buckets.