

Soft Language Clustering for Multilingual Model Pre-training

Jiali Zeng^{1*}, Yufan Jiang³, Yongjing Yin², Yi Jing⁴, Fandong Meng¹,
Binghuai Lin³, Yunbo Cao³, Jie Zhou¹

¹Pattern Recognition Center, WeChat AI, Tencent Inc

²Westlake University ³Tencent Cloud Xiaowei ⁴Northeastern University

{lemonzeng, fandongmeng, withtomzhou}@tencent.com

yinyongjing@westlake.edu.cn

Abstract

Multilingual pre-trained language models have demonstrated impressive (zero-shot) cross-lingual transfer abilities, however, their performance is hindered when the target language has distant typology from source languages or when pre-training data is limited in size. In this paper, we propose XLM-P, which contextually retrieves prompts as flexible guidance for encoding instances conditionally. Our XLM-P enables (1) lightweight modeling of language-invariant and language-specific knowledge across languages, and (2) easy integration with other multilingual pre-training methods. On the tasks of XTREME including text classification, sequence labeling, question answering, and sentence retrieval, both base- and large-size language models pre-trained with our proposed method exhibit consistent performance improvement. Furthermore, it provides substantial advantages for low-resource languages in unsupervised sentence retrieval and for target languages that differ greatly from the source language in cross-lingual transfer¹.

1 Introduction

Multilingual pre-trained language models (mPLMs) such as mBERT (Devlin et al., 2019), mBART (Liu et al., 2020), XLM-R (Conneau et al., 2020) and mT5 (Xue et al., 2021) have lately produced notable advancements in a number of downstream NLP tasks. In particular, the use of mPLMs significantly enhances few-shot fine-tuning and makes possible efficient zero-shot cross-lingual transfer (Hu et al., 2020). Essentially, an ideal mPLM should satisfy two properties: alignment between language pairs, which has been widely studied in the literature (Chi et al., 2022; Ouyang et al., 2021; Chi et al., 2021a); and a good trade-off between high-resource and low-resource

languages, which remains largely unexplored despite the success of mPLMs.

In this paper, we focus on the second property, specially the potential for model performance to suffer when a large number of languages are added. This can occur due to restricted model capacity or computational limitations, resulting in underrepresented languages being allocated less capacity (Conneau et al., 2020). Furthermore, the model’s coverage of world’s languages remains inadequate, limiting the range of language technology applications it can support (Ansell et al., 2021). A typical solution for the coverage-performance trade-off in multilingual learning is to assign additional model parameters to specific languages, such as language identity embeddings (Conneau and Lample, 2019), adaptors (Houlsby et al., 2019; Üstün et al., 2022; Ansell et al., 2021), and language-aware layers (Zhang et al., 2021). However, it is impractical for multilingual pre-training to maintain a separate component for each language, which can lead to more complicated and challenging optimization, especially for low-resource languages.

We propose to approach the above language-aware components from a different perspective. In linguistic typology, some patterns such as nominative-accusative alignment have broad global distributions, whereas others like morphology are more specific and detailed (Donohue and Wichmann, 2008). To take advantage of this, we introduce XLM-P, which uses a set of compact embeddings to represent soft clustering of the language patterns beyond language identity. We refer these embeddings as *prompts*, due to their similarity to prompt tuning (Lester et al., 2021). Concretely, we build a key-value prompt pool and use the attention mechanism to look up the prompts for each input. The retrieved prompts are then prepended to the input embeddings, and serve as categorization information to adapt the model weights conditionally. This allows for more efficient and effective mul-

*Corresponding author.

¹Code and model are available at <https://github.com/lemon0830/XLMP.git>

tilingual learning by leveraging the patterns and similarities across languages rather than maintaining separate components for each language.

We evaluate the proposed XLM-P on *Cross-Lingual Natural Language Understanding* tasks and *Cross-Lingual Sentence Retrieval* tasks of the XTREME benchmark, and the consistent improvement in performance demonstrates its effectiveness. In addition, we conduct empirical analyses to investigate the underlying reasons of the improvement of XLM-P. The advantages of XLM-P can be summed up as follows:

- The prompt pool and instance-wise prompt retrieval are lightweight and only result in 0.35% and 0.23% increase in parameters for the base and large models, respectively. When fine-tuning on downstream tasks, the prompt module can be easily added or removed as needed.
- Our XLM-P divides the prompts into general and specific ones without any explicit supervision. The dynamically retrieved instance-wise prompts tame the sentence encoding, thus enhancing the capability of multilingual pre-trained models.
- The prompt module is model-agnostic and can be outfitted with the other frameworks (e.g., encoder-decoder style PLMs) and multilingual pre-training objectives (e.g., contrastive learning used in this paper).

Overall, XLM-P is a versatile and efficient approach for improving multilingual pre-training.

2 Related Work

2.1 Cross-lingual LM pre-training

Trained by the masked language modeling (MLM) loss with a shared multilingual vocabulary, multilingual BERT (Devlin et al., 2019) achieves promising results in cross-lingual natural language understanding tasks (Hu et al., 2020), which has attracted increasing attention to improve the cross-lingual transferability. XLM-R (Conneau et al., 2020) increase the model capacity and use large-scale monolingual training data. In addition to monolingual data, XLM (Conneau and Lample, 2019) performs MLM on bilingual parallel corpus, while ALM (Yang et al., 2020) constructs code-switched sequences. In respect of training objectives, a series of studies have explored various pre-training

tasks to enhance the models’ transferability (Huang et al., 2019; Ouyang et al., 2021; Chi et al., 2021b, 2022, 2021a; Luo et al., 2021).

Compared to the above methods, our XLM-P exploits a small number of compact prompt vectors to tame the sentence encoding. Moreover, the prompt module is model-agnostic, and can be combined with the above methods to achieve further improvement.

2.2 Language-aware Components

To alleviate the issue of the ‘curse of multilinguality’, various language-aware components have been proposed, which can allocate additional capacity for individual languages especially under-represented languages. Conneau and Lample (2019) use language identity embeddings to explicitly guide the model. Ansell et al. (2021) present MAD-G, which contextually generates language adapters from language representations. Üstün et al. (2022) propose a single hypernet-work that unifies multi-task and multilingual learning with efficient adaptation. On multilingual neural machine translation, Philip et al. (2020) trained language-specific adapters. Zhang et al. (2021) use conditional routing to select shared and language-specific parameters. Stickland et al. (2021) use language-agnostic task adapters for fine-tuning BART and mBART to bilingual and multilingual MT.

We differ from them in that we do not use any language indicators, and regard the prompts as a bottleneck for storing a small number of discriminative features. The prompts are encoded by the large language model along with the input, which learns language-invariant and language-specific features via the deep modular interaction.

2.3 Prompt-based Tuning

Our work builds upon the recent results showing the effectiveness of adapting PLMs to downstream tasks conditioning on lightweight prompt vectors (Brown et al., 2020; Lester et al., 2021; Schick and Schütze, 2021; Sanh et al., 2021). Differentiable prompts (Li and Liang, 2021; Lester et al., 2021; Gu et al., 2022) show the power of adapting pre-trained language model to multiple downstream tasks by simply prepending a few learnable parameters to the input. More recently, the effectiveness of prompting has been investigated in multilingual (or cross-lingual) tasks, which is largely unexplored despite the success of prompting in English (Zhao

and Schütze, 2021; Zhou et al., 2022; Huang et al., 2022a).

Instead of exploring prompting in fine-tuning, we propose to adopt dynamic retrieval of prompts, which jointly optimized with the mPLM, as a method of soft language clustering to enhance multilingual pre-training.

3 Method

Figure 1 depicts the addition of a *Prompt Pool* to the transformer-based language model. Before feeding an input to the model, we perform *Instance-wise Prompt Retrieval* and convert the input to a prompt-wrapped one. Both the prompts and the model are jointly optimized on multilingual data using Masked Language Modeling, which trains the model to make predictions based on both context and clustering information. At the fine-tuning stage, we experiment with two strategies: standard fine-tuning and prompt-based fine-tuning.

3.1 Prompt Pool

In our proposed framework, we use a prompt pool to store fine-grained patterns sharing across languages as well as language-specific knowledge. Formally, the prompt pool is defined as:

$$\mathbf{P} = \{P_1, P_2, \dots, P_M\}, \quad (1)$$

where M is the total number of prompts in the prompt pool, and $P_j \in \mathbb{R}^{L_p \times D}$ is a single prompt with L_p vectors, whose dimension is the same as the embedding size D of the mPLM. We associate each prompt as value to a learnable key: $\{(k_1, P_1), (k_2, P_2), \dots, (k_M, P_M)\}$, where $k_i \in \mathbb{R}^{D_k}$. And we denote the set of all keys $\mathbf{K} = \{k_j\}_{j=1}^M$.

3.2 Instance-wise Prompt Retrieval

Ideally, we would like to let the input itself to decide which prompts to choose through query-key matching. Formally, given an input sentence $\mathbf{X} = \{x_0, \dots, x_n\}$, we first map each token x_i to a real-valued vector e_i by an embedding layer. The sequence of embeddings is then concatenated with the embedding of “[CLS]” token $\mathbf{E} = \{e_{cls}, e_0, \dots, e_n\}$. Next, we operate a pooling strategy² on the embedding sequence to obtain the input representation:

$$r = \text{Pool}(\mathbf{E}). \quad (2)$$

²In this paper, we use *max* pooling.

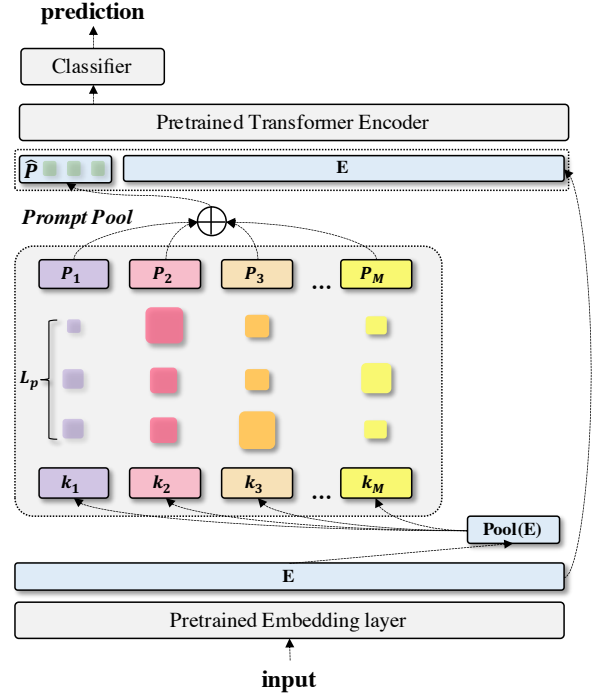


Figure 1: **Model architecture of XLM-P.** For each input, we look up the prompts in the key-value prompt pool $\{(k_j, P_j)\}$, where $P_j \in \mathbb{R}^{L_p \times D}$ and L_p is the number of vectors for each prompt. The retrieved prompts are prepended to the input embeddings E , and serve as soft categorization information across languages.

Next, we obtain the prompts for the input using the attention mechanism:

$$\hat{P} = \sum \alpha_j P_j, \quad (3)$$

$$\alpha_j = \text{softmax}((rW)k_j^T), \quad (4)$$

where W is a trainable parameter. Finally, we prepend the retrieved prompt \hat{P} to the input embeddings $\hat{\mathbf{E}} = [\hat{P}, \mathbf{E}]$, and feed the extended sequence $\hat{\mathbf{E}}$ to the model to get the contextualized sentence representations $\mathbf{H} = \{h_{P_0}, \dots, h_{P_{L_p}}, h_{cls}, h_0, \dots, h_n\}$.

Notably, unlike language identity embeddings used by Conneau and Lample (2019), this soft manner allows the model to capture fine-grained language-invariant and language-specific features from massive monolingual data without explicitly assigning a cue to each input.

3.3 Prompt-based Masked Language Modeling for Pre-training

Our XLM-P is trained to predict the masked tokens using the prompt-augmented contextualized representations. Following Devlin et al. (2019), we randomly mask 15% of the tokens in a monolingual sentence. With probabilities of 0.8, 0.1, and 0.1,

we replace each masked token with a special token “[MASK]”, a random token, or the unchanged token, respectively. The monolingual MLM loss is defined as:

$$\mathcal{L}_{\text{MLM}} = - \sum_{j \in M_x} \log p(x_j | \mathbf{H}, \mathbf{X}_{\setminus M_x}) \quad (5)$$

where $\mathbf{X}_{\setminus M_x}$ is the masked version of input \mathbf{X} . Since the prompt module is agnostic to training objectives, we can enhance the model with additional objectives, as demonstrated in Section 4.2.

3.4 XLM-P for Downstream Applications

Benefiting from the plug-and-play property of the proposed prompt module, we can choose either plug in or plug out it on demand. We explore two fine-tuning strategies to use our XLM-P on downstream applications, i.e., *Standard Fine-tuning* and *Prompt-based Fine-tuning*, and the detail is presented as follows.

Standard Fine-tuning. In this setting, we unplug the prompt module from the pre-trained model. Concretely, we simply feed the embedding features $\mathbf{E} = \{e_{cls}, e_0, \dots, e_n\}$ excluding the prompt embeddings to the model, which is identical to XLM-R.

Prompt-based Fine-tuning. We can keep the prompt module and conduct the retrieval during fine-tuning, as we do in training. For token-level prediction tasks (e.g., structured prediction and question answering), we remove the hidden states of the prompt after encoding and take the rest $\mathbf{H}_{\setminus \hat{p}} = \{h_{cls}, h_1, \dots, h_n\}$ as the input of task-specific classifiers. For sentence-level classification tasks, we apply a pooling operation on all of the hidden states $\mathbf{H}_{\setminus X} = \{h_{P_0}, \dots, h_{P_{L_p}}, h_{cls}\}$ to obtain the sentence-level representation used for classification.

4 Experiments

4.1 Settings

Pre-training. To train XLM-P, we extract a subset from CC-100 (Conneau et al., 2020) which involves monolingual data in 50 languages. We use XLM-R (Conneau et al., 2020) as the backbone. The XLM-P_{base} model has 12 layers with 768 hidden units and 12 attention heads, and the XLM-P_{large} model has 24 layers with 1024 hidden units and 16 attention heads. We set $M=256$ and $L_p=4$ for both base model and large model, and XLM-P introduces 983,040 and 1,310,720 parameters to

the original pre-trained model, merely accounting for 0.35% and 0.23% of the total parameters, respectively. The detail of the pre-training settings can be found in Appendix A.

Evaluation. We evaluate our model on XTREME (Hu et al., 2020), which is designed to assess the the cross-lingual generalization capabilities of pre-trained language models, with a specific focus on *Cross-lingual Natural Language Understanding* and *Cross-lingual Sentence Retrieval*. There are seven tasks for cross-lingual natural language understanding, which can be grouped into three categories: 1) Structured prediction: part-of-speech tagging (POS) on the Universal Dependencies v2.5 (Daniel Zeman and et al., 2019), and named entity recognition (NER) on the WikiAnn (Pan et al., 2017) dataset; 2) Question answering: cross-lingual question answering on MLQA (Lewis et al., 2020) and XQuAD (Artetxe et al., 2020), and gold passage of typologically diverse question answering (TyDiQA-GoldP, Clark et al. (2020)); 3) Sentence classification: cross-lingual natural language inference (XNLI, Conneau et al. (2018)), and cross-lingual paraphrase adversaries from word scrambling (PAWS-X, Yang et al. (2019)). The aim of the cross-lingual sentence retrieval task is to retrieve relevant sentences across languages, and we use the Tatoeba (Artetxe and Schwenk, 2019) dataset.

Baselines. As baselines, we employ the multilingual pre-trained language models listed below: 1) MBERT (Devlin et al., 2019) is pre-trained with MLM and next sentence prediction on Wikipedia in 104 languages; 2) XLM (Conneau and Lample, 2019) is pre-trained with MLM on 100 languages and translation language modeling (TLM) on 14 language pairs; 3) MT5 (Xue et al., 2021) is the multilingual version of T5 pre-trained with text-to-text tasks; 4) XLM-E (Chi et al., 2022) is trained with two pre-training tasks: namely multilingual replaced token detection (MRTD) and translation replaced token detection (TRTD). We report XLM-E (-TRTD) for fair comparison, which does not use parallel data during pre-training; and 5) VECO (Luo et al., 2021) is a unified cross-lingual language model for both NLU and NLG. 6) InfoXLM (Chi et al., 2021a) is jointly pre-trained with a cross-lingual contrastive learning task.

Task	Structured Prediction		Question Answering			Classification	
	POS	NER	XQuAD	MLQA	TyDiQA	XNLI	PAWS-X
#Languages	33	40	11	7	9	15	7
Metrics	F1	F1	F1/EM	F1/EM	F1/EM	Acc	Acc
MBERT _{base} †	70.3	62.2	64.5/49.4	61.4/44.2	59.7/43.9	65.4	81.9
XLM _{base} †	71.3	61.2	59.8/44.3	48.5/32.6	43.6/29.1	69.1	80.9
MT5 _{base} †	-	55.7	67.0/49.0	64.6/45.0	57.2/41.2	75.4	86.4
XLM-E (-TRTD) _{base} †	74.2	62.7	74.3/58.2	67.8/49.7	57.8/40.6	75.1	87.1
VECO _{large} †	75.1	65.7	77.3/61.8	71.7/53.2	67.6/49.1	79.9	88.7
InfoXLM _{large} †	-	-	-	73.6/55.2	-	81.4	-
<i>Backbones</i>							
XLM-R _{base} †	75.6	61.8	71.9/56.4	65.1/47.2	55.4/38.3	75.0	84.9
XLM-R _{large} †	72.6	65.4	76.6/60.8	71.6/53.2	65.1/45.0	79.2	86.4
<i>Standard Fine-tuning</i>							
XLM-P _{base}	74.3	63.8	75.3/60.4	67.4/49.4	58.5/41.6	75.4	86.6
XLM-P _{large}	76.9	68.1	79.0/63.9	72.4/53.5	72.1/55.0	81.1	88.7
<i>Prompt-based Fine-tuning</i>							
XLM-P _{base}	73.9	63.6	75.8/61.5	68.7/50.1	59.3/42.8	75.1	86.0
XLM-P _{large}	77.0	68.5	79.2/64.4	73.7/56.4	72.7/55.7	81.2	88.9

Table 1: **Experimental results on XTREME cross-lingual natural language understanding tasks.** The methods with † denote that we directly report scores from the corresponding work. We run 5 times with different random seeds and report the averaged results.

4.2 Main Results

Cross-Lingual Natural Language Understanding. Following Hu et al. (2020), we adopt the zero-shot transfer setting for evaluation, in which the models are fine-tuned on English training data but evaluated on all the target languages. Rather than selecting a single model for each language, we use only one model for evaluation. The detail of the hyper-parameters used for fine-tuning can be found in Appendix B.

The results, which are averaged across all the target languages and five runs with different random seeds, are illustrated in Table 1. Compared to the XLM-R based models, the XLM-P based models achieve significantly better performances. Besides, our XLM-P_{base} consistently outperforms the baselines MBERT_{base}, MT5_{base}, and XLM-E (-TRTD)_{base}, which are pre-trained without any parallel corpora. Moreover, XLM-P_{large} brings notable improvements over all the baselines on most of the tasks. Concretely, XLM-P models perform better on the structure prediction and the question answering tasks, while preserving competitive results on the sentence classification tasks. The overall experimental results demonstrate that multilingual pre-training can benefit from the guidance of our proposed prompt module.

Surprisingly, there appears to be minimal differ-

ence in performance between standard fine-tuning and prompt-based fine-tuning. This can be explained by the fact that only English training data was used during fine-tuning, and the prompt embeddings were not specifically optimized for the task at hand. Furthermore, recent studies have highlighted the challenges of prompt tuning in cross-lingual natural language understanding and have shown that this area is gaining more attention (Qi et al., 2022; Huang et al., 2022b). We plan to explore this further in future work. Finally, we compare our model with XLM-R under the translate-train-all setting and the results are reported in Appendix C. These results further support the effectiveness of our model.

Cross-Lingual Sentence Retrieval. Following Chi et al. (2021a) and Hu et al. (2020), we use 14 and 36 languages of the parallel corpora for evaluation, respectively. For sentence representations, we take the average of hidden states in a specific layer, and we use the 10-th layer for XLM-P. Then, we induce translation pairs using a nearest neighbor search with cosine similarity. As illustrated in Table 2, XLM-P achieves 64.0 and 61.7 accuracy scores on Tatoeba-14, and 63.8 and 61.0 accuracy scores on Tatoeba-36 in the directions of en → xx and xx → en, respectively, which outperforms XLM-R significantly. The improvement

Model	Tatoeba-14		Tatoeba-36	
	en \rightarrow xx	xx \rightarrow en	en \rightarrow xx	xx \rightarrow en
XLM-R †	59.5	57.6	55.6	53.4
XLM-E †	74.4	72.3	65.0	62.3
-TRTD †	55.8	55.1	46.4	44.6
InfoXLM †	80.6	77.8	68.6	67.3
-XLCO †	64.6	65.3	50.9	53.5
XLM-P	64.0	61.7	63.8	61.0
XLM-P+	73.2	77.2	76.4	69.0

Table 2: **Experimental results on cross-lingual sentence retrieval task Tatoeba in terms of accuracy (%)**. Methods with † denote that we directly report the scores from (Chi et al., 2022).

possibly due to that the sentence representations obtained from XLM-P encoding extra shared features across languages learned by our prompts. Moreover, under the setting of pre-training on multilingual monolingual corpus, XLM-P performs greatly better than XLM-E (-TRTD) and competes with InfoXLM (-XLCO).

Notably, InfoXLM outperforms our XLM-P due to the benefit of the additional cross-lingual contrastive objective. Our prompt module is designed to be compatible with other multilingual pre-training techniques, and to validate this, we added a simple dropout-based InfoNCE objective (van den Oord et al., 2018; He et al., 2020) to our XLM-P model. More detailed is introduced in Appendix D. We post-train XLM-P on monolingual corpora in 50 languages with both MLM loss and InfoNCE loss. The resulting model, XLM-P+, gives a significant improvement over XLM-P and performs better than XLM-E and InfoXLM on Tatoeba-36, but slightly worse on Tatoeba-14 against InfoXLM. This is due to the fact that we did not use the parallel training data utilized by InfoXLM and XLM-E³. In a nutshell, the results show that our proposed method can be effectively integrated with other pre-training objectives.

Finally, in Figure 2, we illustrate the effect of layer selection on sentence representations. The accuracy score is calculated by taking the average of all the 36 language pairs in xx \rightarrow en directions. The figure shows that all the models exhibit a parabolic trend across layers. Different from XLM-R that achieves the highest accuracy of 56.5 at the 7-th

³There are promising training strategies for pre-training mPLM using parallel corpus with our XLM-P, such as the interaction between retrieved prompts of sentence pairs. This is an intriguing research question that merits further investigation in future studies.

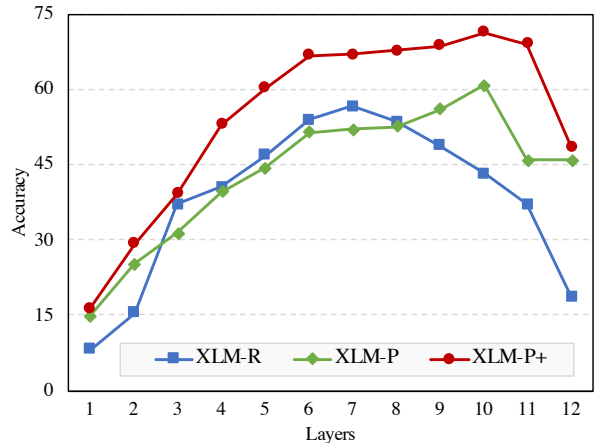


Figure 2: **Evaluation results on Tatoeba cross-lingual sentence retrieval over different layers**. For each layer, the accuracy score is averaged over all the 36 language pairs in the xx \rightarrow en direction.




Figure 3: **Effects of Prompt Capacity**. Left: Average score w.r.t. prompt pool size M given $L_p = 4$. Right: Average score w.r.t. prompt length L_p given $M = 256$. Dashed lines indicate the results of XLM-R_{small} under standard fine-tuning.

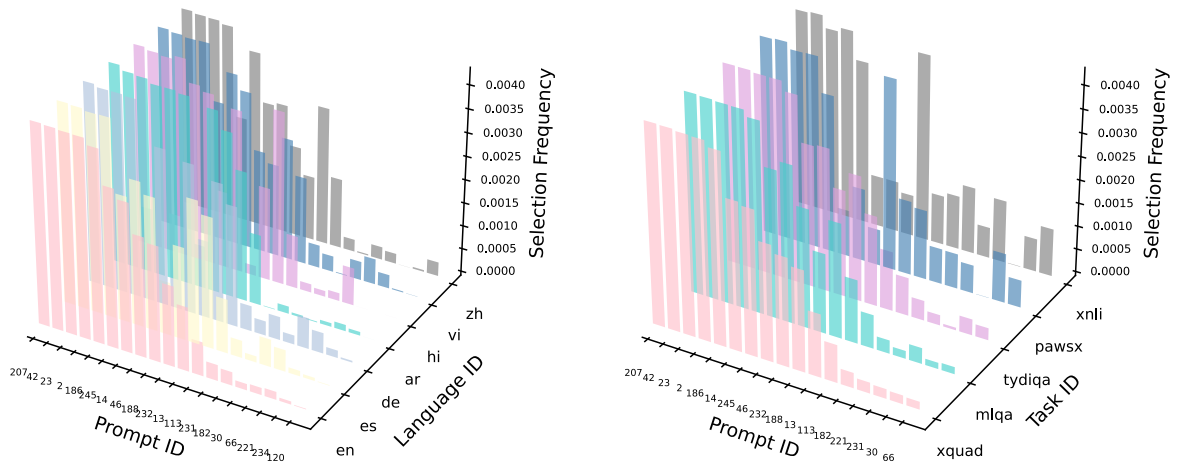
layer, the curve of XLM-P rises more steadily until it peaks at the 10-th layer. It can be observed that XLM-P outperforms XLM-R on the top layers, and XLM-P+ achieves notably higher average scores than XLM-R at all layers.

5 Analysis

We carry out a number of analyses in order to comprehend the design of our proposed XLM-P better.

5.1 Effect of Prompt Capacity

In XLM-P, the length of a simple prompt L_p and the size of the prompt pool M are two critical hyper-parameters that determine the total capacity of the learnable prompt embeddings. M determines the granularity of prompt selection and L_p decides the expressiveness of a single prompt. To investigate the effects of the two hyper-parameters, we train small-size XLM-R and XLM-P from scratch and keep the other setting unchanged (see Appendix A for more details). We evaluate the pre-trained



(a) Prompt selection histograms for MLQA in different languages (b) Prompt selection histograms for different tasks in English

Figure 4: **Prompt selection histograms.** We only display the top 16 selected prompts. The prompts retrieved can be grouped into language-shared and -specific prompts. Similar tasks tend to have similar patterns of prompt selection across different datasets, whereas tasks that are more irrelevant to each other prefer selecting more diverse prompts.

models on TydiQA and report the average of F1 and EM scores in Figure 3.

We can see that removing the prompt pool and only using a single prompt (i.e., pool size $M = 1$) results in a significant performance drop, suggesting that it is not sufficient to maintain a single shared prompt for all languages. Increasing the size of prompt pool shows a positive effect on performance, but excessive prompts degrade the performance. The result verifies our motivation, that the prompts are used to capture abstract clustering information. Too many prompts can dilute the effect, and thus negatively affect the generalization and transferability of representations.

5.2 Patterns of Retrieved Prompts

As depicted in the left portion of Figure 4, we present the prompt selection histograms for XLM-P on test sets of MLQA *in different languages*. The prompts retrieved can be easily distinguished into two categories: language-shared prompts (e.g., prompt 207, 42, 23 and 2) and language-specific prompts (e.g., prompt 14, 66, 221, 120). In addition, we display the histograms on test sets of *different tasks* in English in the right portion of Figure 4. In particular, the similar tasks tend to have similar patterns of prompt selection across different datasets (e.g., xquad and mlqa, both of which are question answering tasks), whereas the tasks that are more irrelevant to each other (e.g., xquad and pawsx, pawsx is a sentence classification task) prefer selecting more diverse prompts. This phenomenon echoes the effectiveness of our

XLM-P. Even without any explicit supervision, our model can learn to group the prompts into general and specific ones.

5.3 Visualization of Retrieved Prompts

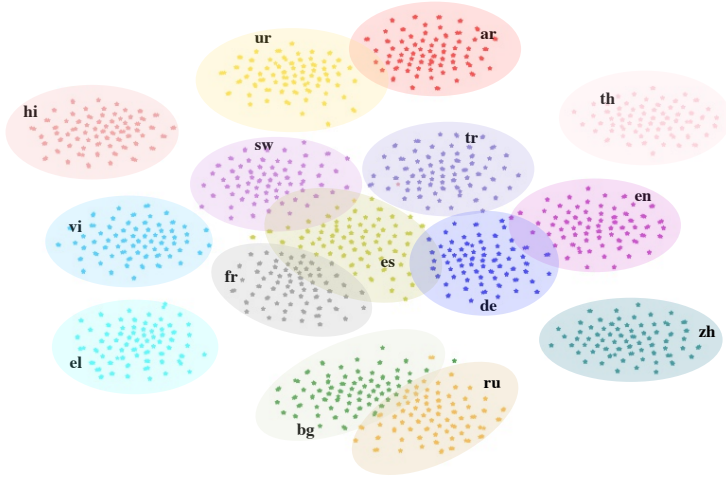
Based on the observation in Section 5.2, we further plot t-SNE visualizations (van der Maaten and Hinton, 2008) of prompts retrieved by sentences in 15 languages on XNLI test set⁴ in Figure 5 to support our hypothesis that prompts can serve as a representation of soft clustering for language.

We can find that the prompt representations retrieved by the sentences in different languages are evidently distinct, while the prompt representations retrieved by the sentences in similar languages (e.g., ar and ur, bg and ru) are closer. It implies that the prompts are able to capture the shared patterns across languages as well as language-specific patterns during multilingual pre-training.

5.4 Trade-off on High and Low Resource Languages

In Section 4.2, we display the average performance across all of the languages, and we present more details about the performance in this section. First, we choose the PAWS-X task with 7 languages including English (i.e., the language of the training set), three European languages (i.e., German, Spanish, French), and three Asian languages (i.e., Japanese, Korea, and Chinese). As shown in Table 4, compared with XLM-R, XLM-P achieves better or more competitive performance on the test sets in

⁴We sampled 200 sentences from each language.



Lang	Example
ar	لم أتحدث معه مرة أخرى.
ur	میں نے اسی دوبارہ بات نہیں کی
bg	Повече не съм говорил с него.
ru	Я больше с ним не разговаривал.
de	Ich habe nicht wieder mit ihm gesprochen.
en	I havent spoken to him again.
es	No he vuelto a hablar con él.
fr	Je ne lui ai pas parlé de nouveau
tr	Onunla bir daha konuşmadım.
sw	Sijaongea na yeye tena.
hi	मैंने उससे दोबारा बात नहीं की है।
th	ฉันไม่ได้คุยกับเขาอีกเลย
vi	Tôi đã không nói chuyện với anh ta nữa.
el	Δεν του μιλήσα ξανά.
zh	我还没有和他再次谈论。

Figure 5: **Visualization of prompts retrieved by sentences in different languages.** We shows one sentence in each language for a clearer understanding. The representations of prompts retrieved by sentences in various languages are distinct, while the representations of prompts retrieved by sentences in similar languages are more alike.

Model	Direction	High Resource							Low Resource							Avg.	Std.(↓)
		de	pt	nl	fr	es	ru	vi	eu	te	tl	bn	ta	sw	jv		
XLM-R †	xx → en	89.9	80.6	79.5	74.1	74.0	72.5	68.4	33.5	32.5	31.2	29.3	25.7	18.7	15.1	51.79	27.04
InfoXLM †	xx → en	93.9	84.7	80.8	79.4	88.2	83.8	89.6	36.7	53.0	42.1	49.6	53.7	39.5	13.2	63.44	25.28
XLM-P+	xx → en	93.3	87.3	88.9	82.8	87.4	82.0	82.8	60.4	58.5	59.1	47.8	49.2	39.3	35.2	68.14	20.28
XLM-R †	en → xx	89.4	80.6	79.5	74.1	74.0	72.5	68.4	33.5	32.5	31.2	29.3	25.7	18.7	15.1	51.75	26.98
InfoXLM †	en → xx	95.1	86.5	81.8	84.0	87.2	85.7	92.0	28.6	53.0	35.5	49.1	63.5	40.8	7.8	63.61	27.85
XLM-P+	en → xx	91.3	87.9	87.6	81.1	84.7	83.2	84.5	73.8	82.9	71.0	62.1	73.3	45.9	50.2	75.68	14.10

Table 3: **Experimental results on cross-lingual sentence retrieval task Tatoeba in terms of accuracy (%)**. Methods with † denote that we directly report the scores from (Chi et al., 2022). XLM-P+ yields similar benefits across different data scales, and greatly benefits low-resource languages (e.g., eu, te, and jv).

Model	Train	European Langs			Asian Langs			Avg.
		en	de	es	fr	ja	ko	
XLM-R	87.8	94.3	88.6	88.8	77.1	75.1	80.4	84.59
XLM-P	89.0	95.1	90.0	90.6	79.4	80.1	81.9	86.59

Table 4: **Experimental results on PAWS-X in terms of accuracy (%)**. We adopt the zero-shot transfer setting for evaluation, in which the models are fine-tuned on English training data but evaluated on all the target languages.

European languages, while obtaining huge gains on the test sets in Asian languages, which are different from English.

Then, we illustrate the results on Tatoeba in both the $xx \rightarrow en$ and $en \rightarrow xx$ directions in Table 3. Due to the limited space, we only select the seven languages with the highest scores and the seven languages with the lowest scores based on the performance of XLM-R on Tatoeba-36 in $xx \rightarrow en$ direction, which can be grouped as *High* resource

language and *Low* resource language pairs, respectively. We compare our XLM-P+ with XLM-R and InfoXLM, and report the mean and the standard deviation of accuracy. Both XLM-P+ and InfoXLM provide advantages in both directions. Specifically, XLM-P+ yields similar benefits across different data scales, and greatly benefits low-resource languages (e.g., eu, te, and jv). By contrast, InfoXLM performs well on high-resource languages but has a marginal impact on low-resource languages.

We argue that mPLMs still suffer from *insufficient modeling capacity*, and adding more languages can result in a decline in representation quality. Our proposed XLM-P can indeed alleviate this issue, especially for languages with limited data or languages that differ from the training data. Intuitively, the dynamically retrieved instance-wise prompts in XLM-P make the sentence encoding specific to the soft clustering information, thereby enhancing the model capability of multilingual pre-trained models.

6 Conclusion

This paper presents XLM-P, a new multilingual pre-trained language model equipped with contextually retrieved prompts. In particular, we prepend prompts-like learnable vectors to the input for modeling language interdependence and other potential sharing information. Compared with other language-aware components, the retrieved prompts are parameter-efficient and more flexible without the requirement of language detection. Experiments and analyses validate the effectiveness and robustness of XLM-P. In addition, our method is compatible with various existing multilingual pre-training objectives.

Limitations

In this paper, we simply prepend the retrieved prompt to the input embeddings before encoding. A well-designed method of combining prompts with the input embeddings, such as Prefix Tuning (Li and Liang, 2021), may result in additional enhancements. Finally, as observed in Section 4.2, prompt-based fine-tuning does not present obvious superiority over standard fine-tuning. Exploring the prompt tuning on cross-lingual natural language understanding is a challenging task that has recently gained attention (Qi et al., 2022; Huang et al., 2022b), and we leave it as future work.

Acknowledgements

We would like to thank all of the anonymous reviewers for the helpful comments.

References

- Alan Ansell, Edoardo Maria Ponti, Jonas Pfeiffer, Sebastian Ruder, Goran Glavaš, Ivan Vulić, and Anna Korhonen. 2021. [MAD-G: Multilingual adapter generation for efficient cross-lingual transfer](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4762–4781, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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.
- 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.

- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared 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 M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher 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 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.

- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2021a. [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. 2021b. [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, Bo Zheng, Saksham Singhal, Payal Bajaj, Xia Song, Xian-Ling Mao, Heyan Huang, and Furu Wei. 2022. [XLM-E: cross-lingual language model pre-training via ELECTRA](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 6170–6182. Association for Computational Linguistics.

- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. [TyDiQA: 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 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 7057–7067.
- 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.
- Mitchell Abrams Daniel Zeman, Joakim Nivre and et al. 2019. Universal dependencies v2.5.
- 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.
- Mark Donohue and Søren Wichmann. 2008. *The typology of semantic alignment*. Oxford University Press.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2022. [PPT: Pre-trained prompt tuning for few-shot learning](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8410–8423, Dublin, Ireland. Association for Computational Linguistics.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. 2020. [Momentum contrast for unsupervised visual representation learning](#). In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 9726–9735. Computer Vision Foundation / IEEE.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.
- 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](#). *CoRR*, abs/2003.11080.
- Haoyang Huang, Yaobo Liang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, and Ming Zhou. 2019. [Unicoder: A universal language encoder by pre-training with multiple cross-lingual tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2485–2494, Hong Kong, China. Association for Computational Linguistics.
- Lianzhe Huang, Shuming Ma, Dongdong Zhang, Furu Wei, and Houfeng Wang. 2022a. [Zero-shot cross-lingual transfer of prompt-based tuning with a unified multilingual prompt](#). *arXiv preprint arXiv:2202.11451*.
- Lianzhe Huang, Shuming Ma, Dongdong Zhang, Furu Wei, and Houfeng Wang. 2022b. [Zero-shot cross-lingual transfer of prompt-based tuning with a unified multilingual prompt](#). *CoRR*, abs/2202.11451.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). 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 4582–4597, Online. Association for Computational Linguistics.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Fuli Luo, Wei Wang, Jiahao Liu, Yijia Liu, Bin Bi, Songfang Huang, Fei Huang, and Luo Si. 2021. [VECO: Variable and flexible cross-lingual pre-training for language understanding and generation](#). 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 3980–3994, Online. Association for Computational Linguistics.
- Xuan Ouyang, Shuhuan Wang, Chao Pang, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2021. [ERNIE-M: enhanced multilingual representation by aligning cross-lingual semantics with monolingual corpora](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 27–38. Association for Computational Linguistics.
- 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.
- Jerin Philip, Alexandre Berard, Matthias Gallé, and Laurent Besacier. 2020. [Monolingual adapters for zero-shot neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4465–4470, Online. Association for Computational Linguistics.
- Kunxun Qi, Hai Wan, Jianfeng Du, and Haolan Chen. 2022. [Enhancing cross-lingual natural language inference by prompt-learning from cross-lingual templates](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1910–1923, Dublin, Ireland. Association for Computational Linguistics.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, Manan Dey, M. Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Stella Biderman, Leo Gao, Tali Bers, Thomas Wolf, and Alexander M. Rush. 2021. [Multitask prompted training enables zero-shot task generalization](#). *CoRR*, abs/2110.08207.
- Timo Schick and Hinrich Schütze. 2021. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Asa Cooper Stickland, Xian Li, and Marjan Ghazvininejad. 2021. [Recipes for adapting pre-trained monolingual and multilingual models to machine translation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 3440–3453. Association for Computational Linguistics.
- Ahmet Üstün, Arianna Bisazza, Gosse Bouma, Gertjan van Noord, and Sebastian Ruder. 2022. [Hyper-x: A unified hypernetwork for multi-task multilingual transfer](#). *CoRR*, abs/2205.12148.
- Aäron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. [Representation learning with contrastive predictive coding](#). *CoRR*, abs/1807.03748.
- Laurens van der Maaten and Geoffrey Hinton. 2008. [Visualizing data using t-sne](#). *Journal of Machine Learning Research*, 9(86):2579–2605.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Edouard Grave. 2020. [CCNet: Extracting high quality monolingual datasets from web crawl data](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4003–4012, Marseille, France. European Language Resources Association.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Jian Yang, Shuming Ma, Dongdong Zhang, Shuangzhi Wu, Zhoujun Li, and Ming Zhou. 2020. [Alternating language modeling for cross-lingual pre-training](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 9386–9393. AAAI Press.
- 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.
- Biao Zhang, Ankur Bapna, Rico Sennrich, and Orhan Firat. 2021. [Share or not? learning to schedule language-specific capacity for multilingual translation](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

Language				
af	ar	bg	bn	cs
de	el	en	es	et
eu	fr	fa	fi	fy
gu	gd	he	hi	hu
id	it	ja	jv	ka
kk	ko	lt	lv	ms
ml	my	mr	pl	pt
ne	nl	ru	ro	si
sw	ta	te	tr	th
tl	vi	ur	yo	zh

Table 5: **The language code of monolingual pre-training corpus.** Among them, the bolded 15 languages are the languages included in the pre-training data of small-sized XLM-P and small-sized XLM-R.

Mengjie Zhao and Hinrich Schütze. 2021. *Discrete and soft prompting for multilingual models*. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8547–8555, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Meng Zhou, Xin Li, Yue Jiang, and Lidong Bing. 2022. *Enhancing cross-lingual prompting with mask token augmentation*. *arXiv preprint arXiv:2202.07255*.

A Pre-training Details

For monolingual data, following XLM-R (Conneau et al., 2020) and Veco (Luo et al., 2021), we build a clean CommonCrawl Corpus using an open-source tool CCNet (Wenzek et al., 2020). We use monolingual data in 50 languages for base-sized and large-sized XLM-P and monolingual data in 15 languages for small-sized XLM-R and XLM-P. Table 5 reports the language codes for pre-training. Please ref (Luo et al., 2021) for the detailed data statistic of the monolingual pre-training corpus. Following Chi et al. (2021a) and Luo et al. (2021), we initialize the parameters of XLM-P with XLM-R (Conneau et al., 2020). We use the Adam optimizer (Kingma and Ba, 2015) with the learning rate 3e-4 for the base model and 1e-4 for the large model, respectively. The full set of pre-training hyperparameters for small-sized, base-sized and large-sized XLM-P are listed in Table 6. We conduct the pre-training experiments using 64 Nvidia A100-40GB GPUs with 8,192 batch size for base and large XLM-P.

B Hyperparameters for Fine-tuning

In Table 7, we present the hyperparameters for fine-tuning baselines and our XLM-P on the XTREME end tasks. For each task, the hyperparameters are searched on the joint validation set of all languages.

C Translate-Train-All Setting

In this section, we investigate another fine-tuning setting, *Translate-Train-All*, in which we fine tune a PLM on a mixed corpus consisting of golden training data in English and translated training data in other languages.

Table 8 presents the results on NER, TyDiQA, and PAWS-X. XLM-P_{large} outperforms XLM-R_{large} across all the tasks, confirming our model’s effective capacity for cross-lingual transfer.

D Detailed of Dropout-based InfoNCE

Specifically, we construct the “positive pairs” by passing the same sentence to the model twice, and take other sentences in the same mini-batch as “negatives”. We use the average of prompt hidden states and “CLS” hidden states of the latest layer as the sentence representation. The model is trained to predict the positive one among the samples as follows:

$$l_i = -\log \frac{e^{\text{sim}(v_i, v_i^+)/\tau}}{\sum_{k=1}^N e^{\text{sim}(v_i, v_k^+)/\tau}} \quad (6)$$

We post-train our XLM-P in monolingual corpora in 50 languages with the MLM loss and InfoNCE loss. The learning rate is 5e-5, the total number of training step is 100k, and the warmup steps is 10k.

Pre-training Hyperparameters	Large	Base	Small
Number of layers	24	12	4
Hidden Size	1024	768	768
FFN inner hidden size	4096	3072	3072
Attention heads	16	12	12
Attention head size	64	64	64
Embedding Size	1024	768	768
Mask percent	15%	15%	15%
Warmup steps	10k	10k	10k
Learning Rate	1e-4	5e-4	3e-4
Adam ϵ	1e-6	1e-6	1e-6
Adam β_1	0.9	0.9	0.9
Adam β_2	0.98	0.98	0.98
Attention Dropout	0.1	0.1	0.1
Dropout	0.1	0.1	0.1
Weight Decay	0.01	0.01	0.01
Max Sequence Length	512	512	512
Batch Size	8,192	8,192	2,048
Train Steps	240k	240k	125k
Total Parameters	561M	279M	222M

Table 6: The pre-training hyperparameters.

	POS	NER	XQuAD	MLQA	TyDiQA	XNLI	PAWS-X
Batch size	{8, 16, 32}	8	32	32	32	32	32
Learning rate	{1,2,3}e-5	{5,...,9}e-6	{2,3,4}e-5	{2,3,4}e-5	{2,3,4}e-5	{5,...,8}e-6	{10,20}e-6
Warmup	10%	10%	10%	10%	10%	12,500 steps	10%
Epochs	10	10	4	{2,3,4}	{10,20,40}	10	10

Table 7: Hyperparameters used for fine-tuning on the XTREME end tasks.

Model	NER	TyDiQA	PAWS-X	Avg.
XLM-R _{large}	87.8*	72.2/54.8	90.5	80.60
XLM-P _{large}	91.1*	74.2/58.1	91.2	82.82

Table 8: **Experimental results of using pseudo-parallel data.** Since Hu et al. (2020) do not release the translation data of NER task, we use the golden NER training data of XTREME for reproduction.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitations
- A2. Did you discuss any potential risks of your work?
We use the open data to pre-training our model.
- A3. Do the abstract and introduction summarize the paper's main claims?
Abstract & 1. Introduction
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
No response.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
No response.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
No response.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
No response.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
No response.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
No response.

C Did you run computational experiments?

4 Experiments

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
4.1 Settings, Appendix A, Appendix B, Appendix C

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

4.1 Settings, Appendix A, Appendix B, Appendix C

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

4.2 Main results

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Not applicable. Left blank.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.