

On Negative Interference in Multilingual Models: Findings and A Meta-Learning Treatment

Zirui Wang Zachary C. Lipton Yulia Tsvetkov
Carnegie Mellon University, Pittsburgh, USA
{ziruiw, zlipton, ytsvetko}@cs.cmu.edu

Abstract

Modern multilingual models are trained on concatenated text from multiple languages in hopes of conferring benefits to each (positive transfer), with the most pronounced benefits accruing to low-resource languages. However, recent work has shown that this approach can degrade performance on high-resource languages, a phenomenon known as *negative interference*. In this paper, we present the first systematic study of negative interference. We show that, contrary to previous belief, negative interference also impacts low-resource languages. While parameters are maximally shared to learn language-universal structures, we demonstrate that language-specific parameters do exist in multilingual models and they are a potential cause of negative interference. Motivated by these observations, we also present a meta-learning algorithm that obtains better cross-lingual transferability and alleviates negative interference, by adding language-specific layers as meta-parameters and training them in a manner that explicitly improves shared layers' generalization on all languages. Overall, our results show that negative interference is more common than previously known, suggesting new directions for improving multilingual representations.¹

1 Introduction

Advances in pretraining language models (Devlin et al., 2018; Liu et al., 2019; Yang et al., 2019) as general-purpose representations have pushed the state of the art on a variety of natural language tasks. However, not all languages enjoy large public datasets for pretraining and/or downstream tasks. Multilingual language models such as mBERT (Devlin et al., 2018) and XLM (Lample and Conneau, 2019) have been proven effective for cross-lingual

transfer learning by pretraining a single shared Transformer model (Vaswani et al., 2017) jointly on multiple languages. The goals of multilingual modeling are not limited to improving language modeling in low-resource languages (Lample and Conneau, 2019), but also include zero-shot cross-lingual transfer on downstream tasks—it has been shown that multilingual models can generalize to target languages even when labeled training data is only available in the source language (typically English) on a wide range of tasks (Pires et al., 2019; Wu and Dredze, 2019; Hu et al., 2020).

However, multilingual models are not equally beneficial for all languages. Conneau et al. (2019) demonstrated that including more languages in a single model can improve performance for low-resource languages but hurt performance for high-resource languages. Similarly, recent work (Johnson et al., 2017; Tan et al., 2019; Aharoni et al., 2019; Arivazhagan et al., 2019) in multilingual neural machine translation (NMT) also observed performance degradation on high-resource language pairs. In multi-task learning (Ruder, 2017), this phenomenon is known as *negative interference* or *negative transfer* (Wang et al., 2019), where training multiple tasks jointly hinders the performance on individual tasks.

Despite these empirical observations, little prior work analyzed or showed how to mitigate negative interference in multilingual language models. Particularly, it is natural to ask: (1) Can negative interference occur for low-resource languages also? (2) What factors play an important role in causing it? (3) Can we mitigate negative interference to improve the model's cross-lingual transferability?

In this paper, we take a step towards addressing these questions. We pretrain a set of monolingual and bilingual models and evaluate them on a range of downstream tasks to analyze negative interference. We seek to individually characterize the un-

¹Source code is available at <https://github.com/iedwardwangi/MetaAdapter>.

derlying factors of negative interference through a set of ablation studies and glean insights on its causes. Specifically, we examine if training corpus size and language similarity affect negative interference, and also measure gradient and parameter similarities between languages.

Our results show that negative interference can occur in both high-resource and low-resource languages. In particular, we observe that neither subsampling the training corpus nor adding typologically similar languages substantially impacts negative interference. On the other hand, we show that gradient conflicts and language-specific parameters do exist in multilingual models, suggesting that languages are fighting for model capacity, which potentially causes negative interference. We further test whether explicitly assigning language-specific modules to each language can alleviate negative interference, and find that the resulting model performs better within each individual language but worse on zero-shot cross-lingual tasks.

Motivated by these observations, we further propose to meta-learn these language-specific parameters to explicitly improve generalization of shared parameters on all languages. Empirically, our method improves not only within-language performance on monolingual tasks but also cross-lingual transferability on zero-shot transfer benchmarks. To the best of our knowledge, this is the first work to systematically study and remedy negative interference in multilingual language models.

2 Motivation

Multilingual transfer learning aims at utilizing knowledge transfer across languages to boost performance on low-resource languages. State-of-the-art multilingual language models are trained on multiple languages jointly to enable cross-lingual transfer through parameter sharing. However, languages are heterogeneous, with different vocabularies, morphosyntactic rules, and different pragmatics across cultures. It is therefore natural to ask, *is knowledge transfer beneficial for all languages in a multilingual model?* To analyze the effect of knowledge transfer from other languages on a specific language lg , we can compare multilingual models with the monolingual model trained on lg . For example, in Figure 1, we compare the performance on a named entity recognition (NER) task of monolingually-trained models vs. bilingual models (trained on lg and English) vs. state-of-




Figure 1: Comparing monolingual vs multilingual models on NER. Lower performance of multilingual models is likely an indicator of negative interference.

the-art XLM (Conneau et al., 2019). We can see that monolingual models outperform multilingual models on four out of six languages (See §3.3 for details). This shows that language conflicts may induce negative impacts on certain languages, which we refer to as *negative interference*. Here, we investigate the causes of negative interference (§3.3) and methods to overcome it (§4).

3 Investigating the Sources of Negative Interference in Multilingual Models

3.1 Methodology

To study negative interference, we compare multilingual models with monolingual baselines. Without loss of generality, we focus on analyzing bilingual models to minimize confounding factors. For two languages lg_1 and lg_2 , we pretrain a single bilingual model and two monolingual models. We then assess their performance on downstream tasks using two different settings. To examine negative interference, we evaluate both monolingual and multilingual models using the **within-language monolingual** setting, such that the pretrained model is finetuned and tested on the same language. For instance, if the monolingual model of lg_1 outperforms the bilingual model on lg_1 , we know that lg_2 induces negative impact on lg_1 in the bilingual model. Besides, since multilingual models are trained to enable cross-lingual transfer, we also report their performance on the **zero-shot cross-lingual transfer** setting, where the model is only finetuned on the source language, say lg_1 , and tested on the target language lg_2 .

We hypothesize that the following factors play important roles in causing negative interference and study each individually:

Training Corpus Size While prior work mostly report negative interference for high-resource lan-

	en	ar	fr	ru	hi	sw	te
corpus size	44.6	8.7	16.2	13.1	0.5	0.2	0.3
NER	✓	✓	✓	✓	✓	✓	✓
POS	✓	✓	✓	✓	✓	✓	✓
QA	✓	✓		✓		✓	✓
XNLI	✓	✓	✓	✓	✓	✓	

Table 1: Language training corpora statistics and downstream tasks availability. Corpus size measured in millions of sentences.

languages (Conneau et al., 2019; Arivazhagan et al., 2019), we hypothesize that it can also occur for languages with less resources. We study the impact of training data size per language on negative interference. We subsample a high-resource language, say lg_1 , to create a “low-resource version”. We then retrain the monolingual and bilingual models and compare with results of their high-source counterparts. Particularly, we test if reducing lg_1 ’s training size also reduces negative interference on lg_2 .

Language Similarity Language similarity has been shown important for effective transfer in multilingual models. Wu et al. (2019) shows that bilingual models trained on more similar language pairs result in better zero-shot transfer performance. We thus expect it to play a critical role in negative interference as well. For a specific language lg_1 , we pair it with languages that are closely and distantly related. We then compare these bilingual models’ performance on lg_1 to investigate if more similar languages cause less severe interference. In addition, we further add a third language lg_3 that is similar to lg_1 and train a trilingual model on lg_1 - lg_2 - lg_3 . We compare the trilingual model with the bilingual model to examine if adding lg_3 can mitigate negative interference on lg_1 .

Gradient Conflict Recent work (Yu et al., 2020) shows that gradient conflict between dissimilar tasks, defined as a negative cosine similarity between gradients, is predictive of negative interference in multi-task learning. Therefore, we study whether gradient conflicts exist between languages in multilingual models. In particular, we sample one batch for each language in the model and compute the corresponding gradients’ cosine similarity for every 10 steps during pretraining.

Parameter Sharing State-of-the-art multilingual models aim to share as many parameters as possible in the hope of learning a language-universal model for all languages (Wu et al., 2019). While prior studies measure the latent embedding

similarity between languages, we instead examine model parameters directly. The idea is to test whether model parameters are **language-universal** or **language-specific**. To achieve this, we prune multilingual models for each language using relaxed L_0 norm regularization (Louizos et al., 2017), and compare parameter similarities between languages. Formally, for a model $f(\cdot; \theta)$ parameterized by $\theta = \{\theta_i\}_{i=1}^n$ where each θ_i represents an individual parameter or a group of parameters, the method introduces a set of binary masks \mathbf{z} , drawn from some distribution $q(\mathbf{z}|\pi)$ parametrized by π , and learns a sparse model $f(\cdot; \theta \odot \mathbf{z})$ by optimizing:

$$\begin{aligned} \min_{\pi} \quad & \mathbb{E}_{q(\mathbf{z}|\pi)} \left[\frac{1}{N} \sum_{i=1}^N \mathcal{L}(f(x_i; \tilde{\theta}), y_i) + \lambda \|\tilde{\theta}\|_0 \right] \\ \text{s.t.} \quad & \tilde{\theta} = \theta \odot \mathbf{z}, \end{aligned} \quad (1)$$

where \odot is the Hadamard (elementwise) product, $\mathcal{L}(\cdot)$ is some task loss and λ is a hyper-parameter. We follow the work of (Louizos et al., 2017) and use the Hard Concrete distribution for the binary mask \mathbf{z} , such that the above objective is fully differentiable. Then, for each bilingual model, we freeze its pretrained parameter weights and learn binary masks \mathbf{z} for *each* language independently. As a result, we obtain two independent sets of mask parameters π which can be used to determine parameter importance. Intuitively, for each parameter group, it is language-universal if both languages consider it important (positive π values). On the other hand, if one language assigns positive value while the other assigns negative, it shows that the parameter group is language-specific. We compare them across languages and layers to analyze parameter similarity in multilingual models.

3.2 Experimental Setup

We focus on standard multilingual masked language modeling (MLM) used in mBERT and XLM. We first pretrain models and then evaluate their performance on four NLP benchmarks.

For pretraining, we mainly follow the setup and implementation of XLM (Lample and Conneau, 2019). We focus on monolingual and bilingual models for a more controllable comparison, which we refer to as **Mono** and **JointPair** respectively. In particular, we always include English (En) in bilingual models to compare on zero-shot transfer settings with prior work. Besides, we consider three

Model	NER (F1)							POS (F1)					
	ar	fr	ru	hi	sw	te	avg	ar	fr	ru	hi	te	avg
Within-language Monolingual													
Mono	89.2	88.0	87.8	89.1	85.1	82.1	86.9	92.7	76.2	96.7	97.0	94.5	91.4
JointPair	86.9	86.5	84.2	88.3	86.1	76.2	84.7	89.2	75.8	93.2	95.2	88.7	88.4
+ ffn	88.2	88.4	86.6	88.9	85.4	81.2	86.5	92.4	76.1	95.6	96.1	92.4	90.5
+ attn	87.3	86.8	84.1	88.5	84.9	77.4	84.8	91.8	75.4	94.4	95.3	90.9	89.6
+ adpt	87.8	86.8	84.5	87.7	86.3	77.0	85.0	91.7	75.6	94.0	95.2	91.5	89.6
+ share adpt	86.8	86.7	84.3	88.6	86.1	76.0	84.8	89.3	76.4	93.5	95.2	88.2	88.5
+ meta adpt	88.9	88.3	85.1	88.4	86.5	79.5	86.1	92.4	75.9	95.1	95.8	92.2	90.3
XLM	89.4	87.5	85.5	88.5	86.3	80.5	86.3	94.5	72.9	96.6	97.1	92.2	90.7
Zero-shot Cross-lingual													
JointPair	38.1	77.5	57.5	61.4	64.8	45.2	57.4	58.5	44.2	80.1	58.9	72.8	62.9
+ ffn	8.9	35.2	5.8	10.5	9.7	12.5	13.8	5.4	8.1	4.5	3.3	7.7	5.8
+ attn	15.4	39.4	10.2	9.9	13.4	11.6	16.7	6.2	4.5	7.5	4.8	6.9	6.0
+ adpt	37.2	75.5	59.2	61.0	64.4	44.7	57.0	57.0	43.5	81.6	58.2	73.5	62.8
+ share adpt	38.5	77.8	58.4	62.0	65.4	44.5	57.8	58.7	43.8	82.5	59.7	71.8	63.3
+ meta adpt	44.4	78.5	62.4	66.0	67.3	50.1	61.5	63.5	44.6	84.9	62.7	78.5	66.8
XLM	44.8	78.3	63.6	65.8	68.4	49.3	61.7	62.8	42.4	86.3	65.7	76.9	66.8

Table 2: NER and POS results. We observe negative interference when monolingual models outperform multilingual models. Besides, adding language-specific layers (e.g. ffn) mitigates interference but sacrifices transferability.

high-resource languages {Arabic (Ar), French (Fr), Russian (Ru)} and three low-resource languages {Hindi (Hi), Swahili (Sw), Telugu (Te)} (see Table 1 for their statistics). We choose these six languages based their data availability in downstream tasks. We use Wikipedia as training data with statistics shown in Table 1. For each model, we use BPE (Sennrich et al., 2016) to learn 32k subword vocabulary shared between languages. For multilingual models, we sample language proportionally to $P_i = (\frac{L_i}{\sum_j L_j})^{\frac{1}{T}}$, where L_i is the size of the training corpus for i -th language pair and T is the temperature. Each model is a standard Transformer (Vaswani et al., 2017) with 8 layers, 12 heads, 512 embedding size and 2048 hidden dimension for the feedforward layer. Notice that we specifically consider a smaller model capacity to be comparable with existing models with larger capacity but also include much more (over 100) languages. We use the Adam optimizer (Kingma and Ba, 2014) and exploit the same learning rate schedule as Lample and Conneau (2019). We train each model with 4 NVIDIA V100 GPUs with 32GB of memory. Using mixed precision, we fit a batch of 128 for each GPU and the total batch size is 512. Each epoch contains 10k steps and we train for 50 epochs.

For evaluation, we consider four downstream tasks: named entity recognition (NER), part-of-speech tagging (POS), question answering (QA), and natural language inference (NLI). (See Appendix A for finetuning details.)

NER We use the WikiAnn (Pan et al., 2017) dataset, which is a sequence labelling task built automatically from Wikipedia. A linear layer with softmax classifier is added on top of pretrained models to predict the label for each word based on its first subword. We report the F1 score.

POS Similar to NER, POS is also a sequence labelling task but with a focus on synthetic knowledge. In particular, we use the Universal Dependencies treebanks (Nivre et al., 2018). Task-specific layers are the same and we report F1, as in NER.

QA We choose to use the TyDiQA-GoldP dataset (Clark et al., 2020) that covers typologically diverse languages. Similar to popular QA dataset such as SQuAD (Rajpurkar et al., 2018), this is a span prediction task where task-specific linear classifiers are used to predict start/end positions of the answer. Standard metrics of F1 and Exact Match (EM) are reported.

NLI XNLI (Conneau et al., 2018) is probably the most popular cross-lingual benchmark. Notice that the original dataset only contains training data for English. Consequently, we only evaluate this task on the zero-shot transfer setting while we consider both settings for the rest of other tasks.

3.3 Results and Analysis

In Table 2 and 3, we report our results on NER, POS and QA together with XLM-100, which is trained on 100 languages and contains 827M parameters. In particular, we observe that monolin-

Model	ar	ru	sw	te	avg
Within-language Monolingual					
Mono	74.2	63.1	52.5	58.2	62.0
JointPair	71.3	58.2	52.8	52.2	58.6
+ ffn	73.4	61.2	51.4	57.5	60.9
+ attn	72.8	60.8	51.2	52.8	59.4
+ adpt	71.5	59.4	52.1	55.5	59.6
+ share adpt	71.0	58.5	52.8	53.9	59.1
+ meta adpt	73.0	61.8	54.5	56.2	61.4
XLM	74.3	62.5	58.7	55.4	62.7
Zero-shot Cross-lingual					
JointPair	54.1	43.2	41.5	21.5	40.1
+ ffn	2.2	0.0	4.4	0.0	1.7
+ attn	3.7	2.1	0.7	0.0	1.6
+ adpt	53.4	44.7	41.2	20.4	39.9
+ share adpt	54.3	44.8	42.2	22.7	41.0
+ meta adpt	57.5	45.8	43.0	23.1	42.4
XLM	59.4	47.3	42.3	16.3	41.3

Table 3: TyDiQA-GoldP results (F1). See Appendix C for full results.

gual models outperform bilingual models for all languages except Swahili on all three tasks. In fact, monolingual models even perform better than XLM on four out of six languages including *hi* and *te*, despite that XLM is much larger in model sizes and trained with much more resources. This shows that negative interference *can* occur on low-resource languages as well. While the negative impact is expected to be more prominent on high-resource languages, we demonstrate that it may occur for languages with resources fewer than commonly believed. The existence of negative interference confirms that state-of-the-art multilingual models cannot generalize equally well on all languages, and there is still a gap compared to monolingual models on certain languages.

We next turn to dissect negative interference by studying the four factors described in Section 3.1.

Training Corpus Size By comparing the validation perplexity on Swahili and Telugu in Figure 2, we find that while both monolingual models outperform bilingual models in the first few epochs, the Swahili model’s perplexity starts to *increase* and is eventually surpassed by the bilingual model in later epochs. This matches the intuition that monolingual models may overfit when training data size is small. To verify this, we subsample French and Russian to 100k sentences to create a “low-resource version” of them (denoted as fr_l/ru_l). As shown in Table 5, while the performance for both models drop compared to their “high-resource” counterparts, bilingual models indeed outperform mono-

Model	NER (F1)		POS (F1)	
	hi	te	hi	te
Within-language Monolingual				
JointPair	88.3	76.2	95.2	88.7
JointTri	87.8	76.4	95.3	88.7
Zero-shot Cross-lingual				
JointPair	61.4	45.2	58.9	72.8
JointTri	63.5	47.6	59.5	74.4

Table 4: Comparing trilingual models with bilingual models. This shows the effect of adding a third similar language to bilingual models.




Figure 2: Validation perplexity during pretraining.

lingual models for fr_l/ru_l , in contrast for fr/ru . This suggests that multilingual models can stimulate positive transfer for low-resource languages when monolingual models overfit. On the other hand, when we compare bilingual models on English, models trained using different sizes of fr/ru data obtain similar performance, indicating that the training size of the source language has little impact on negative interference on the target language (English in this case). While more training data usually implies larger vocabulary and more diverse linguistic phenomena, negative interference seems to arise from more fundamental conflicts contained in even small training corpus.

Language Similarity As illustrated by Table 5, the in-language performance on English drops as the paired language becomes more distantly related (French vs Russian). This verifies that transferring from more distant languages results in more severe negative interference.

It is therefore natural to ask if adding more similar languages can mitigate negative interference, especially for low-resource languages. We then train two trilingual models, adding Marathi to English-Hindi, and Kannada to English-Telugu. Compared to their bilingual counterparts (Table 4), trilingual models obtain similar within-language performance, which indicates that adding similar languages *cannot* mitigate negative interference.

Model	NER (F1)				POS (F1)				QA (F1/EM)	
	fr	fr _l	ru	ru _l	fr	fr _l	ru	ru _l	ru	ru _l
Within-language Performance on fr/ru										
Mono	88.0	81.7	87.8	82.4	76.2	68.5	96.7	88.7	63.1/49.2	47.2/29.5
JointPair	86.5	83.2	84.2	82.7	75.8	71.4	93.2	89.5	58.2/43.1	49.5/30.4
Within-language Performance on en										
JointPair	78.6	78.4	75.8	75.9	94.5	94.5	92.7	92.3	61.7/49.8	62.1/50.2

Table 5: Evaluating effects of training corpus sizes on negative interference.




Figure 3: Gradients similarity throughout training. “En-En” refers to gradients of two English batches within the Ar-En model, while “Ar-En” and “Fr-En” refer to gradients of two batches, one from each language, within Ar-En and Fr-En models respectively.

However, they do improve zero-shot cross-lingual performance. One possible explanation is that even similar languages can fight for language-specific capacity but they may nevertheless benefit the generalization of the shared knowledge.

Gradient Conflict In Figure 3, we plot the gradient cosine similarity between Arabic-English and French-English in their corresponding bilingual models over the first 25 epochs. We also plot the similarity within English, measured using two independently sampled batches². Specifically, gradients between two different languages are indeed less similar than those within the same language. The gap is more evident in the early few epochs, where we observe negative gradient similarities for Ar-En and Fr-En while those for En-En are positive. In addition, gradients in Ar-En are less similar than those in Fr-En, indicating that distant language pair can cause more severe gradient conflicts. These results confirm that gradient conflict exists in multilingual models and is correlated to per language performance, suggesting it may introduce optimization challenge that results in negative interference.

Parameter Sharing The existence of gradient

²Notice that we use gradient accumulation to sample an effectively larger batch of 4096 sentences to calculate the gradient similarity.

conflicts may imply that languages are fighting for capacity. Thus, we next study how language-universal these multilingual parameters are. Figure 4a shows the cosine similarity of mask parameters π across different layers. We observe that within-language similarity (En-En) is near perfect, which validates the pruning method’s robustness. The trend shows that model parameters are better shared in the bottom layers than the upper ones. Besides, it also demonstrates that parameters in multi-head attention layers obtain higher similarities than those in feedforward layers, suggesting that attention mechanism might be more language-universal. We additionally inspect π parameters with the highest absolute values and plot those values for Ar (Figure 4b), together with their En counterparts. A more negative value indicates that the parameter is more likely to be pruned for that language and vice versa. Interestingly, while many parameters with positive values (on the right) are language-universal as both languages assign very positive values, parameters with negative values (on the left) are mostly language-specific for Ar as En assigns positive values. We observe similar patterns for other languages as well. These results demonstrate that language-specific parameters do exist in multilingual models.

Having language-specific capacity in shared parameters is sub-optimal. It is less transferable and thus can hinder cross-lingual performance. Moreover, it may also take over capacity budgets for other languages and degrade their within-language performance, i.e., causing negative interference. A natural next question is whether explicitly adding language-specific capacity into multilingual models can alleviate negative interference. We thus train variants of bilingual models that contain language-specific components for each language. Particularly, we consider adding language-specific feedforward layers, attention layers, and residual adapter layers (Rebuffi et al., 2017; Houlsby et al., 2019), denoted as ffn, attn and adpt respectively.




Figure 4: **Left:** Parameter similarity across layers. **Middle:** Normalized pruning variables of highest absolute values for Ar in Ar-En model. 10 parameter groups with most negative values are shown on the left and 10 with most positive values are shown on the right. **Right:** Average MLM training loss after the warm-up stage.

For each type of component, we create two separate copies in each Transformer layer, one designated for each language, while the rest of the network remains unchanged. As shown in Table 2 and 3, adding language-specific capacity does mitigate negative interference and improve monolingual performance. We also find that language-specific feedforward layers obtain larger performance gains compared to attention layers, consistent with our prior analysis. However, these gains come at a cost of cross-lingual transferability, such that their zero-shot performance drops tremendously. Our results suggest a tension between addressing interference versus improving transferability. In the next section, we investigate how to address negative interference in a manner that can improve performance on *both* within-language tasks and cross-lingual benchmarks.

4 Mitigating Negative Interference via Meta Learning

4.1 Proposed Method

In the previous section, we demonstrated that while explicitly adding language-specific components can alleviate negative interference, it can also hinder cross-lingual transferability. We notice that a critical shortcoming of language-specific capacity is that they are **agnostic** of the rest of other languages, since by design they are trained on the designated language only. They are thus more likely to overfit and can induce optimization challenges for shared capacity as well. Inspired by recent work in meta learning (Flennerhag et al., 2019) that utilizes meta parameters to improve gradient geometry of the base network, we propose a novel meta-learning formulation of multilingual models that exploits language-specific parameters to im-

prove generalization of shared parameters.

For a model with some predefined language-specific parameters $\phi = \{\phi_i\}_{i=1}^L$, where ϕ_i is designated for the i -th language, and shared parameters θ , our solution is to treat ϕ as meta parameters and θ as base parameters. Ideally, we want ϕ to store non-transferable language-specific knowledge to resolve conflicts and improve generalization of θ in all languages (a.k.a. mitigate negative interference and improve cross-lingual transferability). Therefore, we train ϕ based on the following principle: *if θ follows the gradients on training data for a given ϕ , the resulting θ should obtain a good validation performance on all languages.* This implies a bilevel optimization problem (Colson et al., 2007) formally written as:

$$\begin{aligned} \min_{\phi} \quad & \frac{1}{L} \sum_{i=1}^L \mathcal{L}_{\text{val}}^i(\theta^*, \phi_i) \\ \text{s.t.} \quad & \theta^* = \arg \min_{\theta} \frac{1}{L} \sum_{i=1}^L \mathcal{L}_{\text{train}}^i(\theta, \phi_i), \end{aligned} \quad (2)$$

where $\mathcal{L}_{\text{val}}^i$ and $\mathcal{L}_{\text{train}}^i$ denote the training and the validation MLM loss for the i -th language. Since directly solving this problem can be prohibitive due to the expensive inner optimization, we approximate θ^* by adapting the current $\theta^{(t)}$ using a single gradient step, similar to techniques used in prior meta-learning methods (Finn et al., 2017). This results in a two-phase iterative training process shown in Algorithm 1 (See Appendix B).

To be specific, at each training step t on the i -th language during pretraining, we first adapt a gradient step on θ to obtain a new θ' and update ϕ_i

Algorithm 1 Training XLM with Meta Language-specific Layers

- 1: **Input:** Training data
 - 2: **Output:** The converged model $\{\theta^*, \phi^*\}$
 - 3: Initialize model parameters $\{\theta^{(0)}, \phi^{(0)}\}$
 - 4: **while** not converged **do**
 - 5: Sample language i
 - 6: Update language-specific parameters as:
 $\phi_i^{(t+1)} \leftarrow \text{GradientUpdate}(\phi_i^{(t)}, \nabla_{\phi_i^{(t)}} \frac{1}{L} \sum_{j=1}^L \mathcal{L}_{\text{val}}^j(\theta^{(t)} - \beta \nabla_{\theta^{(t)}} \mathcal{L}_{\text{train}}^i(\theta^{(t)}, \phi_i^{(t)}, \phi_j^{(t)}))$
 - 7: Update shared parameters as:
 $\theta^{(t+1)} \leftarrow \text{GradientUpdate}(\theta^{(t)}, \nabla_{\theta^{(t)}} \mathcal{L}_{\text{train}}(\theta^{(t)}, \phi^{(t+1)}))$
 - 8: **end while**
-

based on the θ' 's validation MLM loss:

$$\begin{aligned} \phi_i^{(t+1)} &= \phi_i^{(t)} - \alpha \nabla_{\phi_i^{(t)}} \frac{1}{L} \sum_{j=1}^L \mathcal{L}_{\text{val}}^j(\theta', \phi_j^{(t)}) \\ \theta' &= \theta^{(t)} - \beta \nabla_{\theta^{(t)}} \mathcal{L}_{\text{train}}^i(\theta^{(t)}, \phi_i^{(t)}), \end{aligned} \quad (3)$$

where α and β are learning rates. Notice that θ' is a function of $\phi_i^{(t)}$ and thus this optimization requires computing the gradient of gradient. Particularly, by applying chain rule to the gradient of $\phi_i^{(t)}$, we can observe that it contains a higher-order term:

$$\left[\nabla_{\phi_i^{(t)}, \theta^{(t)}}^2 \mathcal{L}_{\text{train}}^i(\theta^{(t)}, \phi_i^{(t)}) \right] \cdot \left[\nabla_{\theta'} \frac{1}{L} \sum_{j=1}^L \mathcal{L}_{\text{val}}^j(\theta', \phi_j^{(t)}) \right] \quad (4)$$

This is important, since it shows that ϕ_i can obtain information from other languages through higher-order gradients. In other words, language-specific parameters are **not** agnostic of other languages anymore without violating the language-specific requirement. This is because, in Eq. 3, while $\nabla_{\theta^{(t)}}$ is based on the i -th language only, the validation loss is computed for all languages. Finally, in the second phase, we update θ based on the new $\phi^{(t+1)}$:

$$\theta^{(t+1)} = \theta^{(t)} - \beta \nabla_{\theta^{(t)}} \mathcal{L}_{\text{train}}(\theta^{(t)}, \phi^{(t+1)}) \quad (5)$$

4.2 Evaluation

While our method is generic, we evaluate it applied on bilingual models with adapter networks. Adapters have been effectively utilized in multilingual models (Bapna et al., 2019), and we choose them for practical consideration of limiting per-language capacity. Unlike prior works that finetune adapters for adaptation, we train them jointly with shared parameters during pretraining. We follow Houlisby et al. (2019) and insert language-specific adapters after attention and feedforward layers. We

leave a more thorough investigation of how to better pick language-specific structures for future work. For downstream task evaluation, we finetune all layers. Notice that computing the gradient of gradient in Eq. 3 doubles the memory requirement. In practice, we utilize the finite difference approximation (Appendix B).

By evaluating their performance on the zero-shot transfer settings (Table 2, 3 and 6), we observe that our method, denoted as meta adpt, consistently improves the performance over JointPair baselines, while ordinary adapters (adpt) perform worse than JointPair. This shows that, the proposed method can effectively utilize the added language-specific adapters to improve generalization of shared parameters across languages. At the same time, our method also mitigates negative interference and outperforms JointPair on within-language performance, closing the gap with monolingual models. In particular, it performs better than ordinary adapters in both settings. We hypothesize that this is because it alleviates language conflicts during training and thus converges more robustly. For example, we plot training loss in the early stage in Figure 4c, which shows that ordinary adapters converge slower than JointPair due to overfitting of language-specific adapters while meta adapters converge much faster. For ablation studies, we also report results for JointPair trained with adapters shared between two languages, denoted as share adpt. Unlike language-specific adapters that can hinder transferability, shared adapters improve both within-language and cross-lingual performance with the extra capacity. However, meta adapters still obtain better performance. These results show that mitigating negative interference can improve multilingual representations.

Model	ar	fr	ru	hi	sw	avg
JointPair	67.1	73.5	69.2	61.5	62.3	66.7
+ ffn	42.5	51.4	40.7	36.2	34.8	41.1
+ attn	48.5	50.7	41.2	33.3	35.1	41.8
+ adpt	67.8	73.7	69.5	62.2	59.7	66.6
+ share adpt	67.9	73.4	70.0	61.8	62.2	67.1
+ meta adpt	68.5	74.8	70.2	64.5	61.5	67.9
XLM	68.2	75.2	72.3	65.4	58.1	67.8

Table 6: XNLI results (Accuracy).

5 Related Work

Unsupervised multilingual language models such as mBERT (Devlin et al., 2018) and XLM (Lample and Conneau, 2019; Conneau et al., 2019) work surprisingly well on many NLP tasks without parallel training signals (Pires et al., 2019; Wu and Dredze, 2019). A line of follow-up work (Wu et al., 2019; Artetxe et al., 2019; Karthikeyan et al., 2020) study what contributes to the cross-lingual ability of these models. They show that vocabulary overlap is not required for multilingual models, and suggest that abstractions shared across languages emerge automatically during pretraining. Another line of research investigate how to further improve these shared knowledge, such as applying post-hoc alignment (Wang et al., 2020b; Cao et al., 2020) and utilizing better calibrated training signal (Mulleire et al., 2019; Huang et al., 2019). While prior work emphasize how to share to improve transferability, we study multilingual models from a different perspective of how to unshare to resolve language conflicts.

Our work is also related to transfer learning (Pan and Yang, 2010) and multi-task learning (Ruder, 2017). In particular, prior work have observed (Rosenstein et al., 2005) and studied (Wang et al., 2019) negative transfer, such that transferring knowledge from source tasks can degrade the performance in the target task. Others show it is important to remedy negative transfer in multi-source settings (Ge et al., 2014; Wang and Carbonell, 2018). In this work, we study negative transfer in multilingual models, where languages contain heavily unbalanced training data and exhibit complex inter-task relatedness.

In addition, our work is related to methods that measure similarity between cross-lingual representations. For example, existing methods utilize statistical metrics to examine cross-lingual embeddings such as singular vector canonical correlation analysis (Raghu et al., 2017; Kudugunta et al., 2019), eigenvector similarity (Søgaard et al., 2018),

and centered kernel alignment (Kornblith et al., 2019; Wu et al., 2019). While these methods focus on testing latent representations, we directly compare similarity of neural network structures through network pruning. Finally, our work is related to meta learning, which sets a meta task to learn model initialization for fast adaptation (Finn et al., 2017; Gu et al., 2018; Flennerhag et al., 2019), data selection (Wang et al., 2020a), and hyperparameters (Baydin et al., 2018). In our case, the meta task is to mitigate negative interference.

6 Conclusion

We present the first systematic study of negative interference in multilingual models and shed light on its causes. We further propose a method and show it can improve cross-lingual transferability by mitigating negative interference. While prior efforts focus on improving sharing and cross-lingual alignment, we provide new insights and a different perspective on unsharing and resolving language conflicts.

Acknowledgments

We want to thank Jaime Carbonell for his support on the early stage of this project. We also would like to thank Zihang Dai, Graham Neubig, Orhan Firat, Yuan Cao, Jiateng Xie, Xinyi Wang, Ruochen Xu and Yiheng Zhou for insightful discussions. Lastly, we thank anonymous reviewers for their valuable feedbacks.

References

- Roei Aharoni, Melvin Johnson, and Orhan Firat. 2019. [Massively multilingual neural machine translation](#). *CoRR*, abs/1903.00089.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George Foster, Colin Cherry, et al. 2019. [Massively multilingual neural machine translation in the wild: Findings and challenges](#). *arXiv preprint arXiv:1907.05019*.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2019. [On the cross-lingual transferability of monolingual representations](#). *arXiv preprint arXiv:1910.11856*.
- Ankur Bapna, Naveen Arivazhagan, and Orhan Firat. 2019. [Simple, scalable adaptation for neural machine translation](#). *arXiv preprint arXiv:1909.08478*.

- Atilim Gunes Baydin, Robert Cornish, David Martinez Rubio, Mark Schmidt, and Frank Wood. 2018. Online learning rate adaptation with hypergradient descent.
- Steven Cao, Nikita Kitaev, and Dan Klein. 2020. Multilingual alignment of contextual word representations. In *International Conference on Learning Representations*.
- Jonathan H Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. Tydi qa: A benchmark for information-seeking question answering in typologically diverse languages. *arXiv preprint arXiv:2003.05002*.
- Benoît Colson, Patrice Marcotte, and Gilles Savard. 2007. An overview of bilevel optimization. *Annals of operations research*, 153(1):235–256.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Alexis Conneau, Guillaume Lample, Ruty Rinott, Adina Williams, Samuel R Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. Xnli: Evaluating cross-lingual sentence representations. *arXiv preprint arXiv:1809.05053*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- 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*, pages 1126–1135. JMLR. org.
- Sebastian Flennerhag, Andrei A Rusu, Razvan Pascanu, Hujun Yin, and Raia Hadsell. 2019. Meta-learning with warped gradient descent. *arXiv preprint arXiv:1909.00025*.
- Liang Ge, Jing Gao, Hung Ngo, Kang Li, and Aidong Zhang. 2014. On handling negative transfer and imbalanced distributions in multiple source transfer learning. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 7(4):254–271.
- Jiatao Gu, Yong Wang, Yun Chen, Victor OK Li, and Kyunghyun Cho. 2018. Meta-learning for low-resource neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3622–3631.
- 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. *arXiv preprint arXiv:1902.00751*.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization. *arXiv preprint arXiv:2003.11080*.
- 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. *arXiv preprint arXiv:1909.00964*.
- Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. 2017. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics*, 5:339–351.
- K Karthikeyan, Zihan Wang, Stephen Mayhew, and Dan Roth. 2020. Cross-lingual ability of multilingual bert: An empirical study. In *International Conference on Learning Representations*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. 2019. Similarity of neural network representations revisited. In *International Conference on Machine Learning*.
- Sneha Reddy Kudugunta, Ankur Bapna, Isaac Caswell, Naveen Arivazhagan, and Orhan Firat. 2019. Investigating multilingual nmt representations at scale. *arXiv preprint arXiv:1909.02197*.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *arXiv preprint arXiv:1901.07291*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Christos Louizos, Max Welling, and Diederik P Kingma. 2017. Learning sparse neural networks through l_0 regularization. *arXiv preprint arXiv:1712.01312*.
- Phoebe Mulcaire, Jungo Kasai, and Noah A Smith. 2019. Polyglot contextual representations improve crosslingual transfer. 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 3912–3918.
- Joakim Nivre, Mitchell Abrams, Željko Agić, Lars Ahrenberg, Lene Antonsen, Maria Jesus Aranzabe, Gashaw Arutie, Masayuki Asahara, Luma Ateyah, Mohammed Attia, et al. 2018. Universal dependencies 2.2.

- Sinno Jialin Pan and Qiang Yang. 2010. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359.
- 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.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. How multilingual is multilingual bert? *arXiv preprint arXiv:1906.01502*.
- Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. 2017. Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. In *Advances in Neural Information Processing Systems*, pages 6076–6085.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. *arXiv preprint arXiv:1806.03822*.
- Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. 2017. Learning multiple visual domains with residual adapters. In *Advances in Neural Information Processing Systems*, pages 506–516.
- Michael T Rosenstein, Zvika Marx, Leslie Pack Kaelbling, and Thomas G Dietterich. 2005. To transfer or not to transfer. In *NIPS 2005 workshop on transfer learning*, volume 898, pages 1–4.
- Sebastian Ruder. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725.
- Anders Søgaard, Sebastian Ruder, and Ivan Vulić. 2018. On the limitations of unsupervised bilingual dictionary induction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 778–788.
- Xu Tan, Yi Ren, Di He, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2019. Multilingual neural machine translation with knowledge distillation. *arXiv preprint arXiv:1902.10461*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Xinyi Wang, Yulia Tsvetkov, and Graham Neubig. 2020a. Balancing training for multilingual neural machine translation. *arXiv preprint arXiv:2004.06748*.
- Zirui Wang and Jaime Carbonell. 2018. Towards more reliable transfer learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 794–810.
- Zirui Wang, Zihang Dai, Barnabás Póczos, and Jaime Carbonell. 2019. Characterizing and avoiding negative transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11293–11302.
- Zirui Wang, Jiateng Xie, Ruochen Xu, Yiming Yang, Graham Neubig, and Jaime Carbonell. 2020b. Cross-lingual alignment vs joint training: A comparative study and a simple unified framework. In *International Conference on Learning Representations*.
- Shijie Wu, Alexis Conneau, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Emerging cross-lingual structure in pretrained language models. *arXiv preprint arXiv:1911.01464*.
- Shijie Wu and Mark Dredze. 2019. Beto, bentz, becas: The surprising cross-lingual effectiveness of bert. *arXiv preprint arXiv:1904.09077*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5754–5764.
- Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. 2020. Gradient surgery for multi-task learning. *arXiv preprint arXiv:2001.06782*.

A Fine-tuning Details

Notice that XNLI only has training data in available in English so we only evaluate zero-shot cross-lingual performance on it. Following (Hu et al., 2020), we finetune the model for 10 epochs for NER and POS, 2 epochs for QA and 200 epochs for XNLI. For NER, POS and QA, we search the following hyperparameters: batch size $\{16, 32\}$; learning rate $\{2e-5, 3e-5, 5e-5\}$. We use English dev set for zero-shot cross-lingual setting and the target language dev set for within-language monolingual setting. For XNLI, we search for: batch size $\{4, 8\}$; encoder learning rate $\{1e-6, 5e-6, 2e-5\}$; classifier learning rate $\{5e-6, 2e-5, 5e-5\}$. For models with language-specific components, we test freezing these components or finetuning them together. We discover that finetuning the whole network always yields better results. For all experiments, we save checkpoint after each epoch.

B Method Details

Let z_i be the output of the i -th layer of dimension d . The residual adapter network (Rebuffi et al., 2017; Houlsby et al., 2019; Bapna et al., 2019) is a bottleneck layer that first projects z_i to an inner layer with dimension b :

$$h_i = g(W_i^z z_i) \quad (6)$$

where $W_i^z \in \mathbb{R}^{d \times b}$ and g is some activation function such as *relu*. It is then projected back to the original input dimension d with a residual connection:

$$o_i = W_i^h h_i + z_i \quad (7)$$

where $W_i^h \in \mathbb{R}^{b \times d}$. In our experiments, we fix $b = \frac{1}{4}d$. For a bilingual model of lg_1 and lg_2 , we inject two language-specific adapters after each attention and feedforward layer, one for each language. For example, if the input text is in lg_1 , the network will be routed to adapters designated for lg_1 . The rest of the network and training protocol remain unchanged.

The injected adapter layers mimic the warp layers interleaved between base network layers in Flennerhag et al. (2019). Warp layers are meta parameters that aim to improve the performance of the base network. They precondition base network gradients to obtain better gradient geometry. In our experiments, we treat language-specific adapters as meta parameters to improve generalization of the shared network. The algorithm is outlined in

Algorithm 1. The adapters are updated according to Eq 3, which doubles the memory requirement. In particular, the high-order term in Eq 4 requires computing the gradient of gradient. In practice, we approximate this term using the finite difference approximation as:

$$\frac{\nabla_{\phi_i^{(t)}} \mathcal{L}_{\text{train}}^i(\theta^+, \phi_i^{(t)}) - \nabla_{\phi_i^{(t)}} \mathcal{L}_{\text{train}}^i(\theta^-, \phi_i^{(t)})}{2\epsilon} \quad (8)$$

where $\theta^\pm = \theta^{(t)} \pm \epsilon \nabla_{\theta'} \frac{1}{L} \sum_{j=1}^L \mathcal{L}_{\text{val}}^j(\theta', \phi_j^{(t)})$ and ϵ is a small scalar. We use the same value for learning rates α and β in Eq 3, to be consistent with standard learning rate schedule used in XLM (Lample and Conneau, 2019).

C Extra Results

We show the full results on the TyDiQA-GoldP dataset in Table 7.

Model	ar	ru	sw	te	avg
Within-language Monolingual					
Mono	74.2/62.5	63.1/49.2	52.5/37.4	58.2/41.0	62.0/47.5
JointPair	71.3/58.1	58.2/43.1	52.8/39.0	52.2/36.4	58.6/44.2
+ ffn	73.4/61.2	61.2/45.8	51.4/34.3	57.5/40.5	60.9/45.5
+ attn	72.8/61.0	60.8/45.4	51.2/34.0	52.8/36.8	59.4/44.3
+ adpt	71.5/58.7	59.4/44.8	52.1/38.7	55.5/38.9	59.6/45.3
+ share adpt	71.0/57.8	58.5/43.2	52.8/39.0	53.9/37.2	59.1/44.3
+ meta adpt	73.0/61.4	61.8/46.7	54.5/40.0	56.2/39.5	61.4/36.4
XLM	74.3/63.2	62.5/48.7	58.7/40.4	55.4/38.3	62.7/47.7
Zero-shot Cross-lingual					
JointPair	54.1/39.5	43.2/27.5	41.5/22.2	21.5/14.7	40.1/26.0
+ ffn	2.2/1.5	0.0/0.0	4.4/3.7	0.0/0.0	1.7/1.3
+ attn	3.7/2.0	2.1/1.2	0.7/1.0	0.0/0.0	1.6/1.1
+ adpt	53.4/39.1	44.7/27.9	41.2/21.8	20.4/13.8	39.9/25.7
+ share adpt	54.3/39.6	44.8/27.8	42.2/22.9	22.7/15.6	41.0/26.5
+ meta adpt	57.5/40.8	45.8/28.8	43.0/24.2	23.1/17.7	42.4/27.9
XLM	59.4/41.2	47.3/29.8	42.3/22.0	16.3/7.2	41.3/25.1

Table 7: Full results on TyDiQA-GoldP (F1/EM).