

# MetaXL: Meta Representation Transformation for Low-resource Cross-lingual Learning

Mengzhou Xia<sup>§\*</sup> Guoqing Zheng<sup>†</sup> Subhabrata Mukherjee<sup>‡</sup> Milad Shokouhi<sup>‡</sup>  
Graham Neubig<sup>†</sup> Ahmed Hassan Awadallah<sup>‡</sup>

<sup>§</sup>Princeton University <sup>†</sup>Carnegie Mellon University <sup>‡</sup>Microsoft Research  
mengzhou@princeton.edu, {zheng, submukhe, milads}@microsoft.com  
gneubig@cs.cmu.edu, hassanam@microsoft.com

## Abstract

The combination of multilingual pre-trained representations and cross-lingual transfer learning is one of the most effective methods for building functional NLP systems for low-resource languages. However, for extremely low-resource languages without large-scale monolingual corpora for pre-training or sufficient annotated data for fine-tuning, transfer learning remains an under-studied and challenging task. Moreover, recent work shows that multilingual representations are surprisingly disjoint across languages (Singh et al., 2019), bringing additional challenges for transfer onto extremely low-resource languages. In this paper, we propose MetaXL, a meta-learning based framework that *learns to transform* representations judiciously from auxiliary languages to a target one and brings their representation spaces closer for effective transfer. Extensive experiments on real-world low-resource languages – without access to large-scale monolingual corpora or large amounts of labeled data – for tasks like cross-lingual sentiment analysis and named entity recognition show the effectiveness of our approach. Code for MetaXL is publicly available at [github.com/microsoft/MetaXL](https://github.com/microsoft/MetaXL).

## 1 Introduction

Recent advances in multilingual pre-trained representations have enabled success on a wide range of natural language processing (NLP) tasks for many languages. However, these techniques may not readily transfer onto extremely low-resource languages, where: (1) large-scale monolingual corpora are not available for pre-training and (2) sufficient labeled data is lacking for effective fine-tuning for downstream tasks. For example, multilingual BERT (mBERT) (Devlin et al., 2018) is pre-trained on 104 languages with many articles on

\*Most of the work was done while the first author was an intern at Microsoft Research.

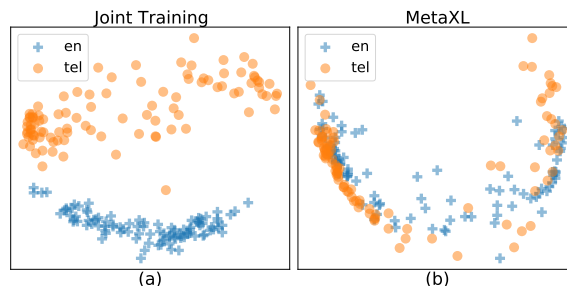


Figure 1: First two principal components of sequence representations (corresponding to [CLS] tokens) of Telugu and English examples from a jointly fine-tuned mBERT and a MetaXL model for the task of sentiment analysis. MetaXL pushes the source (EN) and target (TEL) representations closer to realize a more effective transfer. The Hausdorff distance between the source and target representations drops from 0.57 to 0.20 with F1 score improvement from 74.07 to 78.15.

Wikipedia and XLM-R (Conneau et al., 2020) is pre-trained on 100 languages with CommonCrawl Corpora. However, these models still leave behind more than 200 languages with few articles available in Wikipedia, not to mention the 6, 700 or so languages with no Wikipedia text at all (Artetxe et al., 2020). Cross-lingual transfer learning for these extremely low-resource languages is essential for better information access but under-studied in practice (Hirschberg and Manning, 2015). Recent work on cross-lingual transfer learning using pre-trained representations mainly focuses on transferring across languages that are already covered by existing representations (Wu and Dredze, 2019). In contrast, existing work on transferring to languages without significant monolingual resources tends to be more sparse and typically focuses on specific tasks such as language modeling (Adams et al., 2017) or entity linking (Zhou et al., 2019).

Building NLP systems in these settings is challenging for several reasons. First, a lack of sufficient annotated data in the target language prevents effective fine-tuning. Second, multilingual

pre-trained representations are not directly transferable due to language disparities. Though recent work on cross-lingual transfer mitigates this challenge, it still requires a sizeable monolingual corpus to train token embeddings (Artetxe et al., 2019). As noted, these corpora are difficult to obtain for many languages (Artetxe et al., 2020).

Additionally, recent work (Singh et al., 2019) shows that contextualized representations of different languages do not always reside in the same space but are rather partitioned into clusters in multilingual models. This representation gap between languages suggests that joint training with combined multilingual data may lead to sub-optimal transfer across languages. This problem is further exacerbated by the, often large, lexical and syntactic differences between languages with existing pre-trained representations and the extremely low-resource ones. Figure 1(a) provides a visualization of one such example of the disjoint representations of a resource-rich auxiliary language (English) and resource-scarce target language (Telugu).

We propose a meta-learning based method, MetaXL, to bridge this representation gap and allow for effective cross-lingual transfer to extremely low-resource languages. MetaXL *learns to transform* representations from auxiliary languages in a way that maximally facilitates transfer to the target language. Concretely, our meta-learning objective encourages transformations that increase the alignment between the gradients of the source-language set with those of a target-language set. Figure 1(b) shows that MetaXL successfully brings representations from seemingly distant languages closer for more effective transfer.

We evaluate our method on two tasks: named entity recognition (NER) and sentiment analysis (SA). Extensive experiments on 8 low-resource languages for NER and 2 low-resource languages for SA show that MetaXL significantly improves over strong baselines by an average of 2.1 and 1.3 F1 score with XLM-R as the multilingual encoder.

## 2 Meta Representation Transformation

### 2.1 Background and Problem Definition

The standard practice in cross-lingual transfer learning is to fine-tune a pre-trained multilingual language model  $f_\theta$  parameterized by  $\theta$ , (e.g. XLM-R and mBERT) with data from one or more auxiliary

languages<sup>1</sup> and then apply it to the target language. This is widely adopted in the zero-shot transfer setup where no annotated data is available in the target language. The practice is still applicable in the few-shot setting, in which case a small amount of annotated data in the target language is available.

In this work, we focus on cross-lingual transfer for extremely low-resource languages where only a small amount of unlabeled data and task-specific annotated data are available. That includes languages that are not covered by multilingual language models like XLM-R (e.g., Maori or Turkmen), or low-resource languages that are covered but with many orders of magnitude less data for pre-training (e.g., Telegu or Persian). We assume the only target-language resource we have access to is a small amount of task-specific labeled data.

More formally, given: (1) a limited amount of annotated task data in the target language, denoted as  $\mathcal{D}_t = \{(x_t^{(i)}, y_t^{(i)}); i \in [1, N]\}$ , (2) a larger amount of annotated data from one or more source language(s), denoted as  $\mathcal{D}_s = \{(x_s^{(j)}, y_s^{(j)}); j \in [1, M]\}$  where  $M \gg N$  and (3) a pre-trained model  $f_\theta$ , which is not necessarily trained on any monolingual data from the target language – our goal is to adapt the model to maximize the performance on the target language.

When some target language labeled data is available for fine-tuning, a standard practice is to jointly fine-tune (JT) the multilingual language model using a concatenation of the labeled data from both the source and target languages  $\mathcal{D}_s$  and  $\mathcal{D}_t$ . The representation gap (Singh et al., 2019) between the source language and target language in a jointly trained model brings additional challenges, which motivates our proposed method.

### 2.2 Representation Transformation

The key idea of our approach is to explicitly learn to transform source language representations, such that when training with these transformed representations, the parameter updates benefit performance on the target language the most. On top of an existing multilingual pre-trained model, we introduce an additional network, which we call the *representation transformation network* to model this transformation explicitly.

The representation transformation network models a function  $g_\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ , where  $d$  is the di-

<sup>1</sup>We also refer to auxiliary languages as source languages as opposed to target languages.

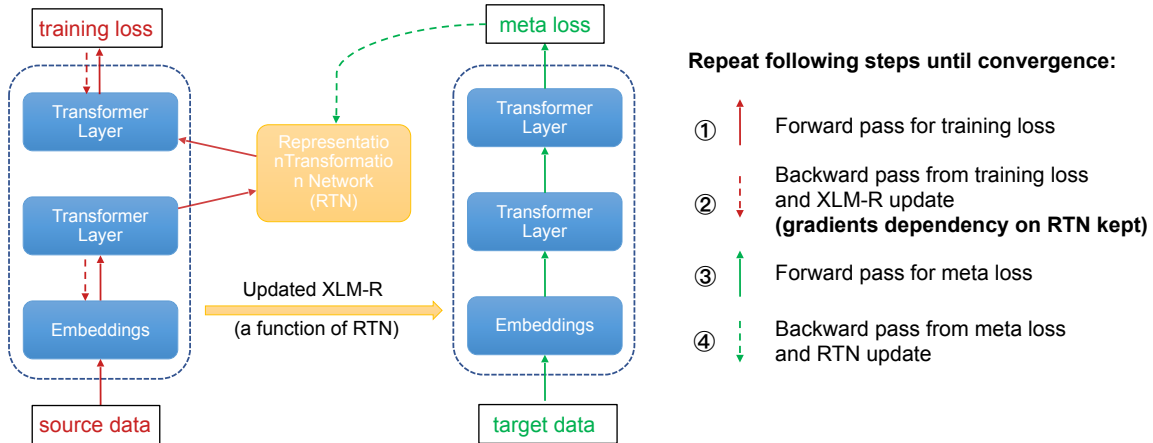


Figure 2: Overview of MetaXL. For illustration, only two Transformer layers are shown for XLM-R, and the representation transformation network is placed after the first Transformer layer. ① source language data passes through the first Transformer layer, through the current representation transformation network, and finally through the remaining layers to compute a training loss with the corresponding source labels. ② The training loss is back-propagated onto all parameters, but only parameters of XLM-R are updated. The updated weights of XLM-R are a function of the current representation transformation network due to gradient dependency (highlighted by the light-purple background of the updated XLM-R). ③ A batch of target language data passes through the updated XLM-R and the meta loss is evaluated with the corresponding labels. ④ The meta loss is back-propagated into the representation transformation network, since the meta-loss is in effect a function of weights from that network, and only the representation transformation network is updated.

mension of the representations. Conceptually, any network with proper input and output sizes is feasible. We opt to employ a two-layer feed-forward network, a rather simple architecture with the intention to avoid heavy parameter overhead on top of the pre-trained model. The input to the representation transformation network is representations from any layer of the pre-trained model. By denoting representations from layer  $i$  as  $h_i \in \mathbb{R}^d$ , we have a parameterized representation transformation network as follows:

$$g_\phi(h_i) = w_2^T (\text{ReLU}(w_1^T h_i + b_1)) + b_2 \quad (1)$$

where  $\phi = \{w_1, w_2, b_1, b_2 | w_1 \in \mathbb{R}^{d \times r}, w_2 \in \mathbb{R}^{r \times d}, b_1 \in \mathbb{R}^r, b_2 \in \mathbb{R}^d\}$  is the set of parameters of the representation transformation network. In practice, we set  $r$  to be bottlenecked, i.e.  $r < d$ , so the representation transformation network first compresses the input representation and then projects back onto the original dimension of the input representation.

As shown in Figure 2, by assuming that the base model has  $N$  layers, a source example  $(x_s, y_s) \in D_s$  passes through the first  $i$  layers, then through the representation transformation network, finally through the last  $N - i$  layers of the base model. We denote the final logits of this batch as  $f(x_s; \theta, \phi)$ ,

encoded by both the base model and the representation transformation network. In contrast, for a target example  $x_t, y_t \in D_t$ , we only pass it through the base model as usual, denoted as  $f(x_t; \theta)$ .

Ideally, suppose that we have a representation transformation network that could properly transform representations from a source language to the target language. In that case, the source data can be almost equivalently seen as target data on a representation level. Unfortunately, we cannot train such a representation transformation network in a supervised manner without extensive parallel data.

Architecturally, the representation transformation network adopts a similar structure to existing works on language and task adapters for cross-lingual and multi-task transfer (Pfeiffer et al., 2020b), a simple down- and up-projection of input representations. Nevertheless, beyond network architecture, the goal and training procedure of the two approaches are significantly different. Adapters are typically trained to encode task or language-specific information by fixing the rest of the model and updating the parameters of the adapters only. Adapters allow training parameter-efficient models that could be flexibly adapted to multiple languages and tasks. While in our proposed method, we use the representation trans-

---

**Algorithm 1** Training procedure for MetaXL

---

- Input:** Input data from the target language  $D_t$  and the source language  $D_s$
- 1: Initialize base model parameters  $\theta$  with pretrained XLM-R weights, initialize parameters of the representation transformation network  $\phi$  randomly
  - 2: **while** not converged **do**
  - 3:   Sample a source batch  $(x_s, y_s)$  from  $D_s$  and a target batch  $(x_t, y_t)$  from  $D_t$ ;
  - 4:   Update  $\theta$ :  $\theta^{(t+1)} = \theta^{(t)} - \alpha \nabla_{\theta} \mathcal{L}(x_s; \theta^{(t)}, \phi^{(t)})$
  - 5:   Update  $\phi$ :  $\phi^{(t+1)} = \phi^{(t)} - \beta \nabla_{\phi} \mathcal{L}(x_t; \theta^{(t)} - \alpha \nabla_{\theta} \mathcal{L}(x_s; \theta^{(t)}, \phi^{(t)}))$
  - 6: **end while**
- 

fer network at training time to adjust the training dynamics to maximally improve test-time performance on the target language. The optimization procedure and the function of the representation transformation network will be discussed in more detail in the next section.

### 2.3 Optimization

The training of the representation transformation network conforms to the following principle: *If the representation transformation network  $g_{\phi}$  effectively transforms the source language representations, such transformed representations  $f(x_s; \phi, \theta)$  should be more beneficial to the target task than the original representations  $f(x_s; \theta)$ , such that the model achieves a smaller evaluation loss  $\mathcal{L}_{D_t}$  on the target language.* This objective can be formulated as a bi-level optimization problem:

$$\begin{aligned} \min_{\phi} \mathcal{L}_{D_t}(f(x_t; \theta^*(\phi)), y_t) \\ \text{s.t. } \theta^*(\phi) = \arg \min_{\theta} \mathcal{L}_{D_s}(f(x_s; \phi, \theta), y_s) \end{aligned} \quad (2)$$

where  $\mathcal{L}(\cdot)$  is the task loss function. In this bi-level optimization, the parameters  $\phi$  of the representation transformation network are the meta parameters, which are only used at training time and discarded at test time. Exact solutions require solving for the optimal  $\theta^*$  whenever  $\phi$  gets updated. This is computationally infeasible, particularly when the base model  $f$  is complex, such as a Transformer-based language model. Similar to existing work involving such optimization problems (Finn et al., 2017; Liu et al., 2019; Shu et al., 2019; Zheng et al., 2021), instead of solving the optimal  $\theta^*$  for any given  $\phi$ , we adopt a one-step stochastic gradient descent update for  $\theta$  as an estimate to the optimal base model for a given  $\phi$ :

$$\theta' = \theta - \alpha \nabla_{\theta} \mathcal{L}_{D_s}(f(x_s; \phi, \theta), y_s) \quad (3)$$

where  $\mathcal{L}_{D_s}(x_s; \cdot)$  is the loss function of the lower problem in Equation 2 and  $\alpha$  is the corresponding

learning rate. Note that the resulting  $\theta'$  is in effect a function of  $\phi$ . We then evaluate the updated weights  $\theta'$  on data  $x_t$  from the target language for updating  $g_{\phi}$ :

$$\phi' = \phi - \beta \nabla_{\phi} \mathcal{L}_{D_t}(f(x_t; \theta'), y_t) \quad (4)$$

where  $\mathcal{L}_{D_t}(x_t; \cdot)$  is the loss function of the upper problem in Equation 2 and  $\beta$  is its corresponding learning rate. Note that the meta-optimization is performed over the parameters of the representation transformation network  $g_{\phi}$  whereas the objective is calculated solely using the updated parameters of the main architecture  $\theta'$ . By plugging Equation 3 into Equation 4, we can further expand the gradient term  $\nabla_{\phi} \mathcal{L}(f(x_t; \theta'), y_t)$ . We omit  $f$  and  $y$  in the following derivative for simplicity.

$$\begin{aligned} \nabla_{\phi} \mathcal{L}_{D_t}(x_t; \theta') \\ = \nabla_{\phi} \mathcal{L}_{D_t}(x_t; \theta - \alpha \nabla_{\theta} \mathcal{L}_{D_s}(x_s; \theta, \phi)) \\ = -\alpha \nabla_{\phi, \theta}^2 \mathcal{L}_{D_s}(x_s; \theta, \phi) \nabla_{\theta} \mathcal{L}_{D_t}(x_t; \theta') \\ = -\alpha \nabla_{\phi} (\nabla_{\theta} \mathcal{L}_{D_s}(x_s; \theta, \phi))^T \nabla_{\theta} \mathcal{L}_{D_t}(x_t; \theta') \end{aligned}$$

During training, we alternatively update  $\theta$  with Equation 3 and  $\phi$  with Equation 4 until convergence. We term our method MetaXL, for its nature to leverage Meta-learning for extremely low-resource cross(X)-Lingual transfer. Both Figure 2 and Algorithm 1 outline the procedure for training MetaXL.

## 3 Experiments

### 3.1 Data

We conduct experiments on two diverse tasks, namely, sequence labeling for Named Entity Recognition (NER) and sentence classification task for Sentiment Analysis (SA). For the NER task, we use the cross-lingual Wikiann dataset (Pan et al., 2017). For the sentiment analysis task, we use the English portion of Multilingual Amazon Reviews

Language	Code	Language Family	Related Language
Quechua	qu	Quechua	Spanish
Min Dong	cdo	Sino-Tibetan	Chinese
Ilocano	ilo	Austronesian	Indonesian
Mingrelian	xmf	Kartvelian	Georgian
Meadow Mari	mhr	Uralic	Russian
Maori	mi	Austronesian	Indonesian
Turkmen	tk	Turkic	Turkish
Guarani	gn	Tupian	Spanish

Table 1: Target language information on the NER task. The data set size of the these languages is 100.

Corpus (MARC) (Keung et al., 2020) as the high-resource language and product review datasets in two low-resource languages, Telugu and Persian (Gangula and Mamidi, 2018; Hosseini et al., 2018).

**WikiAnn** WikiAnn (Pan et al., 2017) is a multilingual NER dataset constructed with Wikipedia articles and anchor links. We use the train, development and test partitions provided in Rahimi et al. (2019). The dataset size ranges from 100 to 20k for different languages.

**MARC** The Multilingual Amazon Reviews Corpus (Keung et al., 2020) is a collection of Amazon product reviews for multilingual text classification. The dataset contains reviews in English, Japanese, German, French, Spanish, and Chinese with five-star ratings. Each language has 200k examples for training. Note that we only use its English dataset.

**SentiPers** SentiPers (Hosseini et al., 2018) is a sentiment corpus in Persian (fa) consisting of around 26k sentences of users’ opinions for digital products. Each sentence has an assigned quantitative polarity from the set of  $\{-2, -1, 0, 1, 2\}$ .

**Sentiraama** Sentiraama (Gangula and Mamidi, 2018) is a sentiment analysis dataset in Telugu (tel), a language widely spoken in India. The dataset contains example reviews in total, labeled as either positive or negative.

**Pre-processing** For SA, we use SentiPers and Sentiraama as target language datasets and MARC as the source language dataset. To unify the label space, we curate MARC by assigning negative labels to reviews rated with 1 or 2 and positive labels to those rated with 4 or 5. We leave out neutral reviews rated with 3. For SentiPers, we assign negative labels to reviews rated with -1 and -2 and positive labels to those rated with 1 or 2. For SentiPers, though the dataset is relatively large, we

mimic the low-resource setting by manually constructing a train, development, and test set with 100, 1000, and 1000 examples through sampling. For Sentiraama, we manually split the dataset into train, development, and test subsets of 100, 103, and 100 examples.<sup>2</sup>

### 3.2 Experimental Setup

**Base Model** We use mBERT<sup>3</sup> (Devlin et al., 2018) and XLM-R (Conneau et al., 2020) as our base models, known as the state-of-the-art multilingual pre-trained model. However, our method is generally applicable to all types of Transformer-based language models.

**Target Language** For NER, we use the same 8 low-resource languages as Pfeiffer et al. (2020c), summarized in Table 1. These languages have only 100 examples in the WikiAnn dataset and are not included for pre-training XLM-R. For SA, Persian and Telugu are the target languages. For both tasks under any setting, we only use a fixed number of 100 examples for each target language.

**Source Language** The selection of source languages is crucial for transfer learning. We experiment with two choices source languages on NER: English and a related language to the target language. The related language was chosen based on LangRank (Lin et al., 2019), a tool for choosing transfer languages for cross-lingual learning. A list of related languages used for each target is shown in Table 1. In absence of training data that fit our related-language criteria for the low-resource target languages in SA, we use only English as the source language.

**Tokenization** For all languages, either pre-trained with XLM-R or not, we use XLM-R’s default tokenizer for tokenizing. We tried with the approach where we train subword tokenizers for unseen languages similar to Artetxe et al. (2020) but obtained worse results than using the XLM-R tokenizer as is, due to the extremely small scale of target language data. We conjecture that the subword vocabulary that XLM-R learns is also beneficial to encode languages on which it is not even

<sup>2</sup>Details of data splits can be found at [github.com/microsoft/MetaXL](https://github.com/microsoft/MetaXL).

<sup>3</sup>XLM-R as a base model leads to significantly better results for both baselines and MetaXL than mBERT, thus we mainly present results with XLM-R in the main text. Detailed results on mBERT can be found in Appendix C

	Source	Method	qu	cdo	ilo	xmf	mhr	mi	tk	gn	average
(1)	-	target	57.14	37.72	61.32	59.07	55.17	76.27	55.56	48.89	56.39
(2)	English	JT	66.10	55.83	80.77	69.32	71.11	82.29	61.61	65.44	69.06
		MetaXL	68.67	55.97	77.57	73.73	68.16	88.56	66.99	69.37	<b>71.13</b>
(3)	Related	JT	79.65	53.91	78.87	79.67	66.96	87.86	64.49	70.54	72.74
		MetaXL	77.06	57.26	75.93	78.37	69.33	86.46	73.15	71.96	<b>73.69</b>

Table 2: F1 for NER across three settings where we, (1) only use the target language data; (2) use target language data along with 5k examples of English; (3) use the target language data along with 5k examples of a related language. JT stands for joint training and MetaXL stands for Meta Representation Transformation. We bold the numbers with a better average performance in each setting.

	Method	tel	fa
(1)	target only	86.87	82.58
(2)	JT	88.68	85.51
	MetaXL	<b>89.52</b>	<b>87.14</b>

Table 3: F1 for sentiment analysis on two settings using (1) only the target language data; (2) target language data along with 1k examples of English.

pre-trained on. We leave exploring the best tokenization strategy for leveraging pre-trained model on unseen language as future work.

## 4 Results and Analysis

### 4.1 Main Results

**NER** We present results of NER in Table 2, where we use 5k examples from English or a related language as source data. When we only use the annotated data of the target language to fine-tune XLM-R (*target*), we observe that the performance varies significantly across languages, ranging from 37.7 to 76.3 F1 score. Jointly fine-tuning XLM-R with target and source data (*JT*) leads to a substantial average gain of around 12.6 F1 score. Using the same amount of data from a related language (instead of English) is more effective, showing an average improvement of 16.3 F1 score over using target data only. Our proposed method, MetaXL, consistently outperforms the joint training baselines, leading to a significant average gain of 2.07 and 0.95 F1 score when paired with English or related languages, respectively.

**SA** We present results on the task of SA in Table 3, where we use 1K examples from English as source language data. We find that auxiliary data from source languages brings less but still signifi-

cant gains to the joint training baseline (*JT*) over using target language data only (*target only*), as in the NER task. In addition, MetaXL still outperforms joint training by around 0.9 and 1.6 F1 score on Telugu and Persian. These results support our hypothesis that MetaXL can transfer representations from other languages more effectively. That, in turn, contributes to the performance gain on the target task.

### 4.2 Source Language Data Size

To evaluate how MetaXL performs with different sizes of source language data, we perform experiments on varying the size of source data. For NER, we experiment with 5k, 10k, and 20k source examples. For SA, we test on 1k, 3k and 5k <sup>4</sup> source examples.

As observed from Table 4, MetaXL delivers consistent gains as the size of source data increases over the joint training model (except on fa when using 5k auxiliary data).<sup>5</sup> However, the marginal gain decreases as the source data size increases on NER. We also note that MetaXL continues to improve even when joint training leads to a minor performance drop for SA.

### 4.3 Placement of Representation Transformation Network

Previous works (Jawahar et al., 2019; Tenney et al., 2019) have observed that lower and intermediate layers encode surface-level and syntactic information, whereas top layers are more semantically focused. These findings suggest that the placement of the representation transformation network can potentially affect the effectiveness of transfer. To

<sup>4</sup>No significant gains were observed for any of the models when going beyond 5K examples.

<sup>5</sup>Please refer to Appendix C for full results.

NER (average)				SA (tel)				SA (fa)			
# en	JT	MetaXL	$\Delta$	# en	JT	MetaXL	$\Delta$	# en	JT	MetaXL	$\Delta$
5k	69.06	71.13	+2.07	1k	88.68	90.53	+1.85	1k	85.51	87.14	+1.63
10k	70.11	71.63	+1.52	3k	87.13	87.23	+0.10	3k	82.88	86.19	+3.31
20k	72.31	73.36	+1.05	5k	84.91	85.71	+0.80	5k	86.34	85.63	-0.71

Table 4: F1 on various source language transfer data sizes. # en denotes the number of English examples used for transfer.  $\Delta$  denotes the improvement of MetaXL over the joint training baseline. RTN is placed after 12th layer.

Method	NER	SA	
	Average	tel	fa
JT	69.06	88.68	85.51
MetaXL L0	70.02	89.52	85.41
MetaXL L6	70.27	86.00	85.80
MetaXL L12	<b>71.13</b>	<b>90.53</b>	<b>87.14</b>
MetaXL L0,12	69.00	84.85	86.64

Table 5: F1 when placing the transfer component at different positions on XLM-R. Under this setting, we use 5k English data for NER and 1K English data for SA. L stands for layer.

this end, we conducted experiments with representation transformation networks placed at various depths of the Transformer model.

Specifically, we experiment with placing the representation transformation network after the 0th (embedding layer), 6th and 12th layer (denoted by L0, L6, L12). We also experiment with placing two identical representation transformation networks after both the 0th and 12th layers.

As observed from Table 5, transformations at the 12th layer are consistently effective, suggesting that transformation at a higher and more abstract level results in better transfer for both tasks.<sup>6</sup> Transferring from lower layers leads to fewer gains for SA, coinciding with the fact that SA is more reliant on global semantic information. Transferring at multiple layers does not necessarily lead to higher performance, possibly because it results in increased instability in the bi-level optimization procedure.

#### 4.4 Joint Training with Representation Transformation Networks

There are two major differences between MetaXL and joint training: (1) source language data un-

Layer	Method	NER	SA	
		Average	tel	fa
-	JT	69.06	88.68	85.51
L0	JT w/ RTN	59.80	63.95	72.32
	MetaXL	<b>70.02</b>	<b>89.52</b>	85.41
L12	JT w/ RTN	67.18	83.75	70.40
	MetaXL	<b>71.13</b>	<b>90.53</b>	<b>87.14</b>

Table 6: F1 when joint training with and without the representation transformation network in XLM-R. In this setting, we use 5k English examples for NER and 1k English examples for SA. NER results are aggregated over 8 target languages. Bold denotes that MetaXL outperforms both JT and JT w/ RTN baselines.

dergoes transformation via an augmented representation transformation network; (2) we adopt a bi-level optimization procedure to update the base model and the representation transformation network. To verify that the performance gain from MetaXL is not attributed to increased model capacity, we conduct experiments on joint training using the representation transformation network. Specifically, the forward pass remains the same as MetaXL, whereas the backward optimization employs the standard stochastic gradient descent algorithm. We conduct experiments on placing the representation transformation network after the 0th layer or 12th layer and present results in Table 6<sup>7</sup>.

Interestingly, joint training with the representation transformation network deteriorates the model performance compared to vanilla joint training. Transferring after the 0th layer is even more detrimental than the 12th layer. This finding shows that Transformer models are rather delicate to subtle architectural changes. In contrast, MetaXL breaks the restriction, pushing the performance higher for both layer settings.

<sup>6</sup>Please refer to Appendix B.2 for full results.

<sup>7</sup>Please refer to Appendix B.3 for full results.




Figure 3: PCA visualization of token-level representations of Quechua and English from the joint training mBERT model on NER. With MetaXL, the Hausdorff distance drops from 0.60 to 0.53 and the F1 score increases from 60.25 to 63.76.

#### 4.5 Analysis of Transformed Representations

To verify that MetaXL does bring the source and target language spaces closer, we qualitatively and quantitatively demonstrate the representation shift with transformation. In particular, we collect representations of both the source and target languages from the joint training and the MetaXL models, with mBERT as the multilingual encoder, and present the 2-component PCA visualization in Figure 1 for SA and Figure 3 for NER. SA models are trained on Telugu paired with 5k English examples, and NER models are trained on Quechua paired with 5k English. From the figures, MetaXL merges the representations from two languages for SA, but the phenomenon is not as evident for NER.

Singh et al. (2019) quantitatively analyze mBERT representations with canonical correlation analysis (CCA). However, CCA does not suit our case as we do not have access to semantically aligned data for various languages. Thus we adopt Hausdorff distance as a metric that has been widely used in vision and NLP tasks (Huttenlocher et al., 1993; Dubuisson and Jain, 1994; Patra et al., 2019) to measure the distance between two distinct datasets. Informally, the Hausdorff distance measures the average proximity of data representations in the source language to the nearest ones in the target language, and vice versa. Given a set of representations of the source language  $\mathcal{S} = \{s_1, s_2, \dots, s_m\}$  and a set of representations of the target language  $\mathcal{T} = \{t_1, t_2, \dots, t_n\}$ , we compute the Hausdorff distance as follows:

$$\max\{\max_{s \in \mathcal{S}} \min_{t \in \mathcal{T}} d(s, t), \max_{t \in \mathcal{T}} \min_{s \in \mathcal{S}} d(s, t)\} \quad (5)$$

where cosine distance is used as the inner distance, i.e.,

$$d(s, t) \triangleq 1 - \cos(s, t) \quad (6)$$

For SA, we observe a drastic drop of Hausdorff distance from 0.57 to 0.20 and a substantial performance improvement of around 4 F1 score. For NER, we observe a minor decline of Hausdorff distance from 0.60 to 0.53 as the representations are obtained at the token level, leading to a significant performance gain of 3 F1 score. For NER, we observe a correlation of 0.4 between performance improvement and the reduction in representation distance. Both qualitative visualization and quantitative metrics confirm our hypothesis that MetaXL performs more effective transfer by bringing the representations from different languages closer.

#### 4.6 Additional Results on High-resource Languages

	fr	es	ru	zh
JT	76.50	72.87	71.14	60.62
MetaXL	72.43	70.38	71.08	58.81

Table 7: F1 on mBERT rich languages in a simulated low-resource setting.

Despite our experiments so far on extremely low-resource languages, given by few labeled data for fine-tuning and limited or no unlabeled data for pre-training, MetaXL is generally applicable to all languages. To better understand the scope of applying MetaXL to languages with varying resources, we perform experiments on 4 target languages that do not belong to our extremely low-resource category for the NER task, namely, Spanish (es), French (fr), Italian (it), Russian (ru) and Chinese (zh). These languages are typically considered high-resource with 20k labeled examples in the WikiAnn datasets and large amount of unlabeled data consumed by mBERT for pre-training. We use only 100 examples for all target languages to mimic the low-resource setting and use 5k English examples for transfer.

As shown in Table 7, we found slight performance drop using MetaXL for these high-resource languages. We conjecture that these languages have been learned quite well with the mBERT model during the pre-training phase, therefore, leaving less scope for effective representation transformation in



the low-resource setup. Nonetheless, this can be remedied with a back-off strategy by further fine-tuning the resulting model from MetaXL on the concatenated data from both source and target languages to match the performance of JT training. As high-resource languages are out of the scope of this paper, we leave further analysis and understanding of these scenarios for future work.

## 5 Related Work

**Unifying Language Spaces** MetaXL in essence brings the source and target representations closer. Previous works have shown that learning invariant representations across languages leads to better transfer. On the representation level, adversarial training is widely adopted to filter away language-related information (Xie et al., 2017; Chen et al., 2018). On the form level, Xia et al. (2019) show that replacing words in a source language with the correspondence in the target language brings significant gains in low-resource machine translation.

**Adapters** Adapter networks are designed to encode task (Houlsby et al., 2019; Stickland and Murray, 2019; Pfeiffer et al., 2020a), domain (Bapna and Firat, 2019) and language (Pfeiffer et al., 2020c) specific information to efficiently share parameters across settings. Though RTN in MetaXL is similar to adapter networks in architecture, in contrast to adapter networks, it plays a more explicit role in transforming representations across languages to bridge the representation gap. More importantly, MetaXL trains the representation transformation network in a meta-learning based paradigm, significantly different from how adapters are trained.

**Meta Learning** MetaXL falls into the category of meta learning for its goal to *learn to transform* under the guidance of the target task. Related techniques have been used in Finn et al. (2017), which aims to learn a good initialization that generalizes well to multiple tasks and is further extended to low-resource machine translation (Gu et al., 2018) and low-resource natural language understanding tasks (Dou et al., 2019). The bi-level optimization procedure is widely adopted spanning across neural architecture search (Liu et al., 2019), instance re-weighting (Ren et al., 2018; Shu et al., 2019), learning from pseudo labels (Pham et al., 2020) and mitigating negative inference in multilingual systems (Wang et al., 2020). MetaXL is the first to

meta learn a network that explicitly transforms representations for cross-lingual transfer on extremely low-resource languages.

## 6 Conclusions and Future Work

In this paper, we study cross-lingual transfer learning for extremely low-resource languages without large-scale monolingual corpora for pre-training or sufficient annotated data for fine-tuning. To allow for effective transfer from resource-rich source languages and mitigate the representation gap between multilingual pre-trained representations, we propose MetaXL to *learn to transform* representations from source languages that best benefits a given task on the target language. Empirical evaluations on cross-lingual sentiment analysis and named entity recognition tasks demonstrate the effectiveness of our approach. Further analysis on the learned transformations verify that MetaXL indeed brings the representations of both source and target languages closer, thereby, explaining the performance gains. For future work, exploring transfer from multiple source languages to further improve the performance and investigating the placement of multiple representation transformation networks on multiple layers of the pre-trained models are both interesting directions to pursue.

## Acknowledgements

We thank the anonymous reviewers for their constructive feedback, and Wei Wang for valuable discussions.

## Ethical Considerations

This work addresses cross-lingual transfer learning onto extremely low-resource languages, which is a less studied area in NLP community. We expect that progress and findings presented in this paper could advocate awareness of advancing NLP for extremely low-resource languages and help improve information access for such under-represented language communities.

The proposed method is somewhat compute-intensive as it requires approximating second-order gradients for updating the meta-parameters. This might impose negative impact on carbon footprint from training the described models. Future work on developing more efficient meta-learning optimization methods and accelerating meta-learning training procedure might help alleviate this concern.

## References

- Oliver Adams, Adam Makarucha, Graham Neubig, Steven Bird, and Trevor Cohn. 2017. Cross-lingual word embeddings for low-resource language modeling. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 937–947.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2019. On the cross-lingual transferability of monolingual representations. *arXiv preprint arXiv:1910.11856*.
- Mikel Artetxe, Sebastian Ruder, Dani Yogatama, Gorka Labaka, and Eneko Agirre. 2020. A call for more rigor in unsupervised cross-lingual learning. *arXiv preprint arXiv:2004.14958*.
- Ankur Bapna and Orhan Firat. 2019. [Simple, scalable adaptation for neural machine translation](#). 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 1538–1548, Hong Kong, China. Association for Computational Linguistics.
- Xilun Chen, Yu Sun, Ben Athiwaratkun, Claire Cardie, and Kilian Weinberger. 2018. Adversarial deep averaging networks for cross-lingual sentiment classification. *Transactions of the Association for Computational Linguistics*, 6:557–570.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard 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.
- 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*.
- 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.
- Zi-Yi Dou, Keyi Yu, and Antonios Anastasopoulos. 2019. [Investigating meta-learning algorithms for low-resource natural language understanding 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 1192–1197, Hong Kong, China. Association for Computational Linguistics.
- M-P Dubuisson and Anil K Jain. 1994. A modified hausdorff distance for object matching. In *Proceedings of 12th international conference on pattern recognition*, volume 1, pages 566–568. IEEE.
- 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.
- Rama Rohit Reddy Gangula and Radhika Mamidi. 2018. Resource creation towards automated sentiment analysis in telugu (a low resource language) and integrating multiple domain sources to enhance sentiment prediction. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Jiatao Gu, Yong Wang, Yun Chen, Victor O. K. 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, Brussels, Belgium. Association for Computational Linguistics.
- Julia Hirschberg and Christopher D Manning. 2015. Advances in natural language processing. *Science*, 349(6245):261–266.
- Pedram Hosseini, Ali Ahmadian Ramaki, Hassan Maleki, Mansoureh Anvari, and Seyed Abolghasem Mirroshandel. 2018. Sentipers: A sentiment analysis corpus for persian. *arXiv preprint arXiv:1801.07737*.
- 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 *International Conference on Machine Learning*, pages 2790–2799.
- Daniel P Huttenlocher, Gregory A. Klanderman, and William J Rucklidge. 1993. Comparing images using the hausdorff distance. *IEEE Transactions on pattern analysis and machine intelligence*, 15(9):850–863.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. What does bert learn about the structure of language? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657.
- Phillip Keung, Yichao Lu, György Szarvas, and Noah A. Smith. 2020. [The multilingual Amazon reviews corpus](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4563–4568, Online. Association for Computational Linguistics.

- Yu-Hsiang Lin, Chian-Yu Chen, Jean Lee, Zirui Li, Yuyan Zhang, Mengzhou Xia, Shruti Rijhwani, Junxian He, Zhisong Zhang, Xuezhe Ma, Antonios Anastasopoulos, Patrick Littell, and Graham Neubig. 2019. Choosing transfer languages for cross-lingual learning. In *The 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, Florence, Italy.
- Hanxiao Liu, Karen Simonyan, and Yiming Yang. 2019. **DARTS: Differentiable architecture search**. In *International Conference on Learning Representations*.
- 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.
- Barun Patra, Joel Ruben Antony Moniz, Sarthak Garg, Matthew R. Gormley, and Graham Neubig. 2019. **Bilingual lexicon induction with semi-supervision in non-isometric embedding spaces**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 184–193, Florence, Italy. Association for Computational Linguistics.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2020a. Adapterfusion: Non-destructive task composition for transfer learning. *arXiv preprint arXiv:2005.00247*.
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020b. **AdapterHub: A framework for adapting transformers**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 46–54, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020c. **MAD-X: An Adapter-Based Framework for Multi-Task Cross-Lingual Transfer**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7654–7673, Online. Association for Computational Linguistics.
- Hieu Pham, Qizhe Xie, Zihang Dai, and Quoc V Le. 2020. Meta pseudo labels. *arXiv preprint arXiv:2003.10580*.
- Afshin Rahimi, Yuan Li, and Trevor Cohn. 2019. Massively multilingual transfer for ner. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 151–164.
- Mengye Ren, Wenyan Zeng, Bin Yang, and Raquel Urtasun. 2018. Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning*, pages 4334–4343.
- Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. 2019. Meta-weight-net: Learning an explicit mapping for sample weighting. In *Advances in Neural Information Processing Systems*, pages 1919–1930.
- Jasdeep Singh, Bryan McCann, Richard Socher, and Caiming Xiong. 2019. Bert is not an interlingua and the bias of tokenization. In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 47–55.
- Asa Cooper Stickland and Iain Murray. 2019. Bert and pals: Projected attention layers for efficient adaptation in multi-task learning. In *International Conference on Machine Learning*, pages 5986–5995.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. Bert rediscovers the classical nlp pipeline. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601.
- Zirui Wang, Zachary C. Lipton, and Yulia Tsvetkov. 2020. **On negative interference in multilingual models: Findings and a meta-learning treatment**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4438–4450, Online. Association for Computational Linguistics.
- Shijie Wu and Mark Dredze. 2019. Beto, bentz, becas: The surprising cross-lingual effectiveness of bert. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 833–844.
- Mengzhou Xia, Xiang Kong, Antonios Anastasopoulos, and Graham Neubig. 2019. Generalized data augmentation for low-resource translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5786–5796.
- Qizhe Xie, Zihang Dai, Yulun Du, Eduard Hovy, and Graham Neubig. 2017. Controllable invariance through adversarial feature learning. In *Advances in Neural Information Processing Systems*, pages 585–596.
- Guoqing Zheng, Ahmed Hassan Awadallah, and Susan Dumais. 2021. Meta label correction for noisy label learning. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*.
- Shuyan Zhou, Shruti Rijhwani, and Graham Neubig. 2019. Towards zero-resource cross-lingual entity linking. *arXiv preprint arXiv:1909.13180*.

## A Hyper-parameters

We use a maximum sequence length of 200 and 256 for NER and AS respectively. We use a bottleneck dimension of  $r = 384$  and  $r = 192$  for the representation transformation network, same as Pfeiffer et al. (2020c). During the bi-level optimization process, we adopt a learning rate of  $3e-05$  for training the main architecture and tuned the learning rate on  $3e-5$ ,  $1e-6$  and  $1e-7$  for training the representation transformation network. We use a batch size of 16 for NER and 12 for AS, and train 20 epochs for each experiment on both tasks. We use a single NVIDIA Tesla V100 with a 32G memory size for each experiment. For each language, we pick the best model according to the validation performance after each epoch.

## B Detailed Results on Each Language

### B.1 Source Data Size

The full results of using 10k and 20k English examples as transfer data are presented in Table 9.

### B.2 Placement of RTN

The full results of placing the representation transformation network at different layers are presented in Table 10.

### B.3 Joint Training w/ RTN

The full results of joint training with the representation transformation network are presented in Table 11.

## C Additional Results on mBERT

We conduct experiments on mBERT (Devlin et al., 2019), which covers 104 languages with most Wikipedia articles. For a language that is not pre-trained with mBERT, we train its subword tokenizer with the task data. Further, we combine the vocabulary from the newly trained tokenizer with the original mBERT vocabulary. A similar

approach has been adopted in (Artetxe et al., 2020). Table 12 and Table 8 present results for NER and SA respectively where we finetune the tasks on mBERT. Note that the languages of SA are both covered by mBERT and XLM-R, while the languages of NER are not. Table 13 show MetaXL results on mBERT with various sizes of source data.

Nevertheless, our method consistently brings gains on both tasks. We observe an average of 2 F1 points improvement on NER and 2.0 F1 points improvement on SA. It shows that the improvement brought by our method is consistent across different language models.

	Method	tel	fa
(1)	target only	75.00	73.86
(2)	JT	75.13	74.81
	MetaXL	<b>77.36</b>	<b>76.69</b>

Table 8: F1 for sentiment analysis on mBERT on two settings using (1) only the target language data; (2) target language data along with 10k examples of English.

	Source	Method	qu	cdo	ilo	xmf	mhr	mi	tk	gn	average
(1)	-	target only	57.14	37.72	61.32	59.07	55.17	76.27	55.56	48.89	56.39
(2)	10k en	JT	71.49	50.21	76.19	73.39	66.36	89.34	66.04	67.89	70.11
		MetaXL	72.57	57.02	81.55	65.56	70.18	90.64	66.98	68.54	<b>71.63</b>
(3)	20k en	JT	73.19	53.93	88.78	71.49	62.56	90.80	68.29	69.44	72.31
		MetaXL	73.04	55.17	85.99	73.09	70.97	89.21	66.02	73.39	<b>73.36</b>

Table 9: Experiment results for NER on XLM-R across three settings where we, (1) only use the target language data; (2) use target language data along with 10k examples of English; (3) use target language data along with 20k examples of English. JT stands for joint training

Layer	Method	qu	cdo	ilo	xmf	mhr	mi	tk	gn	average
-	JT	66.1	55.83	80.77	69.32	71.11	82.29	61.61	65.44	69.06
L0	MetaXL	70.43	54.76	77.14	66.09	68.72	89.53	63.59	69.86	70.02
L6	MetaXL	65.53	56.67	78.5	72.0	68.75	88.05	65.73	66.96	70.27
L0,12	MetaXL	69.83	53.97	69.44	69.26	66.96	89.41	67.92	65.18	69.00

Table 10: Experiment results for NER on XLM-R with RTN placed across multiple layer settings. (All with 5k English examples)

Layer	Method	qu	cdo	ilo	xmf	mhr	mi	tk	gn	average
-	JT	66.10	55.83	80.77	69.32	71.11	82.29	61.61	65.44	69.06
L0	JT w/ RTN	50.81	45.67	60.09	58.91	63.83	81.71	65.37	52.02	59.80
L12	JT w/ RTN	64.41	50.2	73.83	63.87	68.7	85.88	71.92	58.6	67.18

Table 11: Experiment results for NER on XLM-R, Joint Training (JT) with RTN. (All with 5k English examples)

	Source	Method	qu	cdo	ilo	xmf	mhr	mi	tk	gn	average
(1)	-	target	58.44	26.77	63.39	32.06	53.66	82.90	52.53	46.01	51.97
(2)	English	JT	60.25	35.29	73.06	43.45	60.17	86.29	60.09	57.80	59.55
		MetaXL	63.76	38.63	76.36	45.14	60.63	88.96	64.81	54.13	<b>61.55</b>

Table 12: NER results on mBERT where we use 5k English examples as auxiliary data and place RTN after 12th layer.

NER (average)				SA (tel)				SA (fa)			
# en	JT	MetaXL	$\Delta$	# en	JT	MetaXL	$\Delta$	# en	JT	MetaXL	$\Delta$
5k	59.55	61.55	+2.00	100	75.12	77.36	+2.24	100	74.25	75.78	+1.53
10k	62.36	63.66	+1.30	1k	74.76	76.39	+1.63	1k	74.71	75.58	+0.87
20k	62.39	63.38	+0.99	5k	74.07	78.15	+4.08	5k	74.81	76.69	+1.88

Table 13: F1 on various source language transfer data sizes on mBERT. # en denotes the number of English examples used for transfer.  $\Delta$  denotes the improvement of MetaXL over the joint training baseline.