

---

# MetaXCR: Reinforcement-Based Meta-Transfer Learning for Cross-Lingual Commonsense Reasoning

---

**Jie He**

School of Informatics  
University of Edinburgh  
Edinburgh  
j.he@ed.ac.uk

**Yu Fu**

College of Intelligence and Computing  
Tianjin University  
Tianjin  
fuyu\_1998@tju.edu.cn

## Abstract

Commonsense reasoning (CR) has been studied in many pieces of domain and has achieved great progress with the aid of large datasets. Unfortunately, most existing CR datasets are built in English, so most previous work focus on English. Furthermore, as the annotation of commonsense reasoning is costly, it is impossible to build a large dataset for every novel task. Therefore, there are growing appeals for Cross-lingual Low-Resource Commonsense Reasoning, which aims to leverage diverse existed English datasets to help the model adapt to new cross-lingual target datasets with limited labeled data. In this paper, we propose a multi-source adapter for cross-lingual low-resource Commonsense Reasoning (MetaXCR). In this framework, we first extend meta learning by incorporating multiple training datasets to learn a generalized task adapters across different tasks. Then, we further introduce a reinforcement-based sampling strategy to help the model sample the source task that is the most helpful to the target task. Finally, we introduce two types of cross-lingual meta-adaption methods to enhance the performance of models on target languages. Extensive experiments demonstrate MetaXCR is superior over state-of-the-arts, while being trained with fewer parameters than other work.

## 1 Introduction

Commonsense reasoning (CR) is a basic skill of humans to deal with ordinary situations that involve reasoning about physical and social world [8]. To endow computers with human-like commonsense reasoning capabilities is hence one of major goals of artificial intelligence. In natural language processing (NLP), different from well-established tasks such as machine translation, natural language generation, question answering, commonsense reasoning interweaves with many other tasks (e.g., conversation generation [40], machine reading comprehension [24] and exhibits different forms (e.g., question answering [18, 33], co-reference resolution [30]). To facilitate commonsense reasoning research in NLP, a wide variety of CR datasets have been created recently [13, 31], covering different CR forms and aspects, such as : social interaction [13, 31], laws of nature [4], emotional responses [29].

Unfortunately, these datasets are all in English. Leveraging CR resources in the source language of English to enable cross-lingual knowledge transfer to low-resource target languages is therefore desirable for low-resource commonsense reasoning. Among this line of research, training multilingual pretrained language models (e.g., mBERT [10], XLM [7]) on the resources of the source language and fine-tuning on the target language data is a promising approach, which has been explored in open-domain QA [1], semantic role labeling [9], dialogue [32]. However, using multilingual pretrained language models for cross-lingual CR faces the following two challenges.

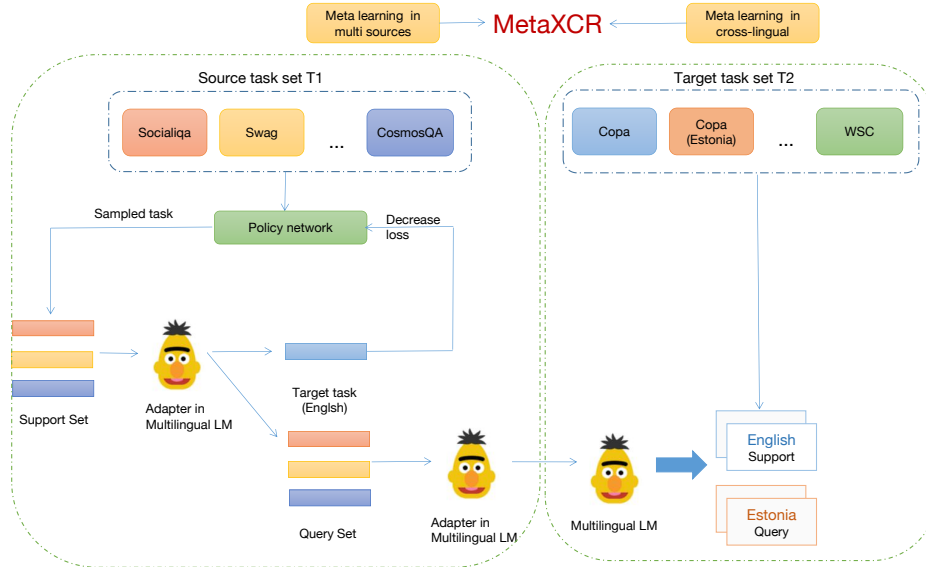


Figure 1: Our MetaXCR framework.

- Language bias in cross-lingual transfer. English-centric resources undoubtedly introduce language bias to commonsense reasoning modeling in other languages [26].
- Dataset selection in cross-lingual transfer. Previous studies have demonstrated that dataset selection in multi-source transfer learning will have a great impact on the target task, and inappropriate datasets may even hurt performance [2, 37].

In order to mitigate these issues, we propose a reinforcement-based meta-transfer learning framework for cross-lingual commonsense reasoning (MetaXCR), which performs cross-dataset and cross-lingual transfer learning to improve the model’s adaptability to the target language dataset. To address the issue of task (i.e., source dataset) sampling, we propose a reinforcement-based learning algorithm for adaptively sampling CR tasks of the source language, which are conducive to the target language CR task. Specifically, as shown in Figure 1, we first randomly sample a batch of source tasks to form a support set for pseudo-update of the adapter, and then sample a batch of instances from the target CR task to form a query set to meta-update the adapter. In the sampling process, we train an LSTM network to predict the sampling probability of the chosen source tasks via the proposed reinforcement learning algorithm. Using the loss decrease of the pseudo model on the target task as reward to the policy network, we enable the model to automatically learn a strategy that is beneficial to its sampling. For the source to target language adaptation, we sample from the source dataset corresponding to the target task to create a support set, and from the dataset of the target task to build a query set, which are used to update parameters of modules of MetaXCR other than the adapter.

To summarize, our contributions are three-fold:

- We propose MetaXCR, to the best of our knowledge, the first work that uses meta-learning and reinforcement learning to for low-resource commonsense reasoning.
- We propose adapters to solve the cross-lingual and cross-dataset challenges, which enables the model to explore the interaction and combination of different CR tasks and languages. We propose a reinforcement learning algorithm for sampling source tasks in meta-transfer learning of commonsense reasoning and compare with other sampling strategies.
- Experiments show that our MetaXCR has achieved state-of-the-art on cross-lingual low-resource commonsense reasoning datasets, XCOPA [26], only requiring 1.6% of the training parameters.

## 2 Related work

### Meta learning

Meta-learning, which is “learn to learn”, trying to enhance the generalization ability of models and dealing with the problem of rapid adaptation to low-resource tasks, is widely used in computer vision [16, 19, 21, 28, 34]. Recent years have gained an increasing attention of meta-learning in NLP. Yan et al. [37] propose a meta-learning method on the monolingual MCQA task to integrate multiple question answering tasks for low-resource QA. Unfortunately, they ignore correlations among different source tasks, and fine-tuning the entire model is not suitable for low-resource tasks as it is prone to overfitting. Xiao et al. [35] propose an approach to improve sampling in the meta-learning process. Different from their work where the goal is to train a generalization model, we are interested in exploring related tasks to train a better model for the target task. Chen and Shuai [6] also uses adapters to perform meta-training on summarization data from different corpora. However, they use rule-based sampling method to pre-select tasks for meta-training, and rules are task-specific. By contrast, our reinforcement-based sampling method can be applied to any tasks. To sum up, previous works in this aspect either use meta-learning to cross-lingual issues for the same task or to generalize over different tasks for the same language. While we propose a meta-learning framework for solving both cross-lingual and cross-task challenges in commonsense reasoning simultaneously.

### Commonsense reasoning

Recent years have witnessed that many commonsense reasoning data sets have been proposed. Gordon et al. [14] create COPA for causal inference while Rahman and Ng [27] present Winograd Scheme Challenge (WSC), a dataset testing commonsense reasoning in the form of anaphora resolution. Since the size of these datasets is too small, effective training can not be obtained until the recent emergence of pre-training methods has brought new improvements to these datasets [17]. Meanwhile, large-scale commonsense reasoning datasets have been curated [4, 18, 20, 30, 31, 33] to train neural CR models. More challenging CR datasets have also been constructed, e.g., HellaSwag [39] from Swag [38] via adversarial filtering, as powerful language models have either established new state-of-the-art performance [5]. RAINBOW [23], which uses multitask learning to provide a pre-trained commonsense reasoning model on top of various large-scale commonsense reasoning datasets, is similar to our work. However, the RAINBOW benchmark is evaluated only on English commonsense reasoning datasets while we focus on cross-lingual low-resource commonsense reasoning scenarios.

## 3 MetaXCR

**Problem Formulation: Cross-lingual Cross-task Low-Resource commonsense reasoning** is a task that requires a model to solve CR problems where the training data is not available or sufficient and the language of the target task is different from the language of source tasks. Let’s denote the labelled training data of source tasks as  $T_s = \{\{T_1^i\}_{i=1}^{n_1}, \dots, \{T_k^i\}_{i=1}^{n_k}\}$ , where  $T_k = \{T_k^i\}_{i=1}^{n_k}$  represents all tasks in the k-th source dataset. The training data of source datasets in this paper are all in English. The target evaluation tasks are  $T_t = \{\{T_\ell^i\}_{i=1}^{n_\ell}\}$ , where  $T_\ell = \{T_\ell^i\}_{i=1}^{n_\ell}$  denotes the target task in the  $\ell$ -th target language. For a single tasks in source dataset  $s_j \in \{1, \dots, k\}$  denote as  $T_{s_j}^i$ , it is composed by support batch denote as  $D_{\text{support}}^{s_j}$  and query batch denote as  $D_{\text{query}}^{s_j}$ . The target evaluation tasks are in CR forms different from those of the source tasks (e.g., Swag vs. WSC). As our goal is to deal with low-resource CR, the target evaluation tasks can be divided into two scenarios: zero-shot and few-shot. Few-shot means  $T_\ell$  has labelled instances that can be used for further fine-tuning. **Zero-shot** indicates that there are no labeled instances for  $T_\ell$  at all and that the model trained on the source tasks will be directly evaluated on the corresponding target task. Specially, We use  $f(\theta)$  to denote the trained multilingual model.

**General introduction to our method:** As shown in Figure 1, we first sample numbers of tasks  $\{T_{s_j}^i\}_{j=1}^M$  from the source task set  $T_s$ . The data of the sampled tasks are then used to update the adapter parameters via the proposed meta-learning described in section 3.2.1. As different source tasks have different contributions to the target task, we propose different sampling strategies choose desirable source tasks for the target task in the cross-task meta-learning process of our MetaXCR in section 3.2.3. As source tasks are all in English while our target tasks are in a range of different languages, we introduce cross-lingual meta-learning in section 3.2.2. In the following, we first introduce the proposed commonsense reasoning framework, and then elaborate on the cross-lingual and cross-task meta-transfer learning process and the reinforcement-based sampling strategy.

---

**Algorithm 1** MetaXCR

---

**Input:** Task distribution over source  $p(T_s)$ ; data distribution over target  $p(T_t)$ ; adapter A with parameters  $\theta_A$ ; pretrained learner B with parameters  $\theta_B$ ; policy network with parameters  $\phi$  learning rates in XMSMS  $\alpha_{ctml}, \alpha_{clml}, \beta_{ctml}, \beta_{clml}, \gamma$

- 1: freeze the pretrained learner parameters  $\theta_B$
- 2: **while** not done **do**
- 3:   Generate k-dim vector of sampling probabilities  $P_T = (P_{T_{s_1}}, P_{T_{s_2}}, \dots, P_{T_{s_k}})$  using  $f(\phi)$
- 4:   choose M source task set  $\{T_{s_j}\}_{j=1}^M$  according to Eq. (6)
- 5:   Sample task  $T_{s_j}^i$  from  $p(T_{s_j})$ , get  $\{T_{s_j}^i\}_{j=1}^M, s_j \in \{1, \dots, k\}$  for meta-learning
- 6:   **for** all  $\{T_{s_j}^i\}_{j=1}^M$  **do**
- 7:     using support set  $D_{\text{support}}^{s_j}$  in  $T_{s_j}^i$  and compute adapted parameters:
- 8:      $\theta_A^{s_j} = \theta_A - \alpha_{ctml} \nabla_{\theta_A} L_{D_{\text{support}}^{s_j}}(f(\theta_A))$
- 9:     using query set  $D_{\text{query}}^{s_j}$  in  $T_{s_j}^i$  and compute the update loss:
- 10:      $L_{s_j} = L_{D_{\text{query}}^{s_j}}(f(\theta_A^{s_j}))$
- 11:   **end for**
- 12:   Update  $\theta_A = \theta_A - \beta_{ctml} \nabla_{\theta_A} \sum_{j=1}^M L_{s_j}$
- 13: **end while**
- 14:
- 15: unfreeze the pretrained learner parameters  $\theta_B$  and freeze the trained parameters  $\theta_A$
- 16: **for** all language  $\ell$  **do**
- 17:   **while** not done **do**
- 18:     Sample target english task set  $T_{t_e}^i \sim p(T_{t_e})$
- 19:     Sample target language task set  $T_{t_\ell}^i \sim p(T_{t_\ell})$
- 20:     using support set  $D_{\text{support}}^{t_e}$  in  $T_{t_e}^i$  and compute adapted parameters:
- 21:      $\theta_B^{t_\ell} = \theta_B - \alpha_{clml} \nabla_{\theta_B} L_{D_{\text{support}}^{t_e}}(f(\theta_B))$
- 22:     using query set  $D_{\text{query}}^{t_\ell}$  and compute the update loss:
- 23:      $L_{t_\ell} = L_{D_{\text{query}}^{t_\ell}}(f(\theta_B^{t_\ell}))$
- 24:     Update  $\theta_B = \theta_B - \beta_{clml} \nabla_{\theta_B} L_{t_\ell}$
- 25:   **end while**
- 26: **end for**

---

### 3.1 Commonsense Reasoning Framework

**Base Model** We choose the state-of-the-art pretrained multilingual BERT (mBERT) [10] as the base commonsense reasoning model. Special token [CLS] is added at the beginning of each sentence to aggregate information. As multiple-choice selection tasks differ in the numbers of candidate answers (e.g., 2 candidate answers per question in COPA vs. 5 in CommonsenseQA), a classifier with a fixed number of classes is not a good fit for this scenario. We thus follow Sap et al. [31]’s method. First, we concatenate the context, question, and answer, using specific separator tokens. For BERT, the concatenation is formatted as [CLS] < context > [UNUSED] < question > [SEP] < answer<sub>i</sub> > [SEP]. Then, we use a multilayer perceptron to compute a score  $\hat{y}_i$  with the classifier token  $h_{\text{CLS}} \in \mathbb{R}^H$ :

$$\hat{y}_i = W_2 \tanh(W_1 h_{\text{CLS}} + b_1) \quad (W_1 \in \mathbb{R}^{H \times H}, b_1 \in \mathbb{R}^H, W_2 \in \mathbb{R}^{1 \times H})$$

Finally, we estimate the probability distribution over candidate answers using a softmax layer:  $\hat{y} = \text{softmax}([\hat{y}_1, \dots, \hat{y}_N])$ . The final answer predicted by the base model corresponds to the triple with the highest probability.

**Adapters** Inspired by Pfeiffer et al. [25], we use task-related adapters in mBERT to separate language and tasks to achieve better performance on cross-lingual CR. For each layer  $l$  in mBERT, we insert an adapter module which is a bottlenecked feed-forward network consisting of a down-layer  $f_{\theta_d}$  and an upproject layer  $f_{\theta_u}$ . A skip-connection from input to output is also established. The adapter (AP) can be formulated as:

$$\text{AP}(h) = f_{\theta_u}(\text{ReLU}(f_{\theta_d}(h)) + h) \quad (1)$$

Only parameters in adapters are updated during fine-tuning on a downstream task, which aim to capture knowledge that is task-specific but generalises across languages.

## 3.2 Cross-Lingual and Cross-Task Meta-Transfer Learning

### 3.2.1 Cross-Task Meta-Learning

CTML is presented in lines 1-13 in Algorithm 1. In CTML, for each meta-training iteration, we sample  $L$  source dataset. For each sampled source dataset, we need to randomly sample instances from it to obtain support and query batch to form a task denoted as  $T_{s_1}, T_{s_2}, \dots, T_{s_L}$  to construct meta-train tasks in meta learning. The learning rate for the inner and outer loop are different:  $\alpha$  denotes the learning rate for the inner loop, while  $\beta$  for the outer loop. The inner loop aims to learn the task information from support batch. For a task  $s_i \in \{s_1, \dots, s_L\}$ , The task specific parameters  $\theta^{s_i}$  are updated by

$$\theta^{s_i} = \theta - \alpha \nabla_{\theta} L_{D_{support}^{s_i}}(f(\theta)) \quad (2)$$

where  $\nabla_{\theta} L_{D_{support}^{s_i}}(f(\theta))$  is computed as the gradient of the cost function  $L_{D_{support}^{s_i}}(f(\theta))$  with respect to model parameter  $\theta$ .

In the outer loop,  $L_{D_{query}^{s_i}}(f(\theta^{s_i}))$  is calculated with respect to  $\theta^{s_i}$ , to update the meta model on the corresponding query batch  $D_{query}^{s_i}$ .

It is worth noting that  $f(\theta^{s_i})$  is an implicit function of  $\theta$ . As the second-order Hessian gradient matrix requires expensive computation, we employ a first-order approximation (FOMAML), which ignores second-order derivatives:

$$\theta = \theta - \beta \frac{1}{L} \sum_{i=1}^L \nabla_{\theta} L_{D_{query}^{s_i}}(f(\theta^{s_i})) \quad (3)$$

We keep running the CTML module until the meta-model converges. After performing CTML, the CLML module will be applied upon the learnt meta representations for the final transfer learning on the target language.

### 3.2.2 Cross-Lingual Meta-Learning

CLML is formulated in lines 15-26 in Algorithm 1. Following MAML [12], in each iteration, we sample a support and query batch from a same task, but with the source and target language respectively. In order to make representations learned before well adapted to different target languages, we sample labelled instances for the support set of tasks in CLML from the training data of source tasks in the source language (i.e., English in our experiments). For the query set in CLML, we sample from the target task data  $T_{t_\ell}$  in the target language. The entire procedure of CLML is similar to CTML. For each target language, we use support batches from the source language for the inner-loop update and query batches from the target language for the outer-loop update. In addition, since only the task adapter is trained in CTML, task representations can be learned. Therefore, in CLML, we fix learned task representations and only train mBERT itself. In order to prevent overfitting, we only train the last layer of mBERT.

## 3.3 Target-adapted meta sampling

For MAML, it is crucial to properly define meta-train tasks. An intuitive strategy is to use tasks from all datasets during every meta-training step. However, this is not a good choice since different source dataset might have different data distributions. A single undesirable task in a dataset might optimize the model towards a wrong direction. We therefore propose a reinforced strategy to choose source tasks for meta-training. In our experiments, we compared the proposed sampling strategy to other heuristic sampling strategies.

### 3.3.1 Heuristic Sampling and Motivation for Reinforced Sampling

For each MAML step, we sample data to form a task  $T_{s_j}^i$  for every source dataset  $s_j$ . Our goal is to train good initialization parameters for the multilingual model on the given target dataset. We can calculate a general loss  $L_{original}$  on the original model  $f(\theta)$  using sampled data from the target dataset. After optimizing the  $f(\theta)$  with  $D_{support}^{s_j}$  of  $s_j$  according to Eq. (2), we can get task-specific model  $f(\theta^{s_j})$  and get task-specific loss  $L_{s_j}$  using the same sampled data with general loss.

To achieve the goal, we can use the countdown of the  $L_{s_j}$  as a guiding signal to sample a corresponding dataset and can be formulated as:

$$P_i^{s_j} = \frac{D_{s_j}}{\sum_{j=1}^k D_{s_j}} \quad (4)$$

where  $s_j$  denote the task from the  $j$ -th source dataset,  $i$  denotes the MAML step.  $D_{s_j}$  is the countdown of  $L_{s_j}$  which can be seen as how good the model might be after being trained on the chosen task.

A simple way is to just use the  $P_i^{s_j}$  as the probability to choose the corresponding task to update the model through Eq. (3). However, this heuristic sampling strategy might not consistent with the distribution of the entire source dataset and ignore the long-range dependencies among a sequence of source dataset sampling decisions. Inspired by [36], we use a LSTM-based network to catch the long-term sampling dependencies to form the meta-training tasks. Since we do not have any annotated labels to train the LSTM-based network, we use REINFORCE, a policy gradient algorithm, for our proposed reinforcement-based meta sampling.

### 3.3.2 RL algorithm

Let  $f_\phi(\cdot)$  denote the network trained by reinforcement learning, and  $\phi$  be parameters to be tuned. For the REINFORCE algorithm, the gradient can be calculated as:

$$\nabla_\phi J(\phi) = E_{\tau \sim f_\phi(\cdot)} [R(\tau) \nabla_\phi \log f_\phi(i)] \quad (5)$$

where  $\tau$  is the trajectory and  $R(\tau)$  is the total reward of the trajectory.  $i$  denotes action in the trajectory  $\tau$ . During REINFORCE training, we use the reward of the total trajectory to guide how to update the network.

we use a baseline to reduce the high variance of the traditional REINFORCE algorithm and let the reward not always be non-negative. To get the expectation, we sample limited numbers of trajectories and use the sampled trajectories to estimate the gradient:

$$\nabla_\phi J(\phi) \approx \frac{1}{N} \sum_{n=1}^N \sum_{g=1}^G \nabla_\phi (R(\tau^n) - b) \log f_\phi(i_g^n) \quad (6)$$

where  $b$  is the baseline value,  $\tau^n$  is the  $n$ -th sampled trajectory,  $i_g^n$  denote the  $g$ -th action in the  $n$ -th trajectory.

### 3.3.3 Our Sampling Approach

We use the same network architecture [36] together with a feedforward and attention layer. For REINFORCE training, We use the  $L_{\text{original}}$  as the baseline, and  $(L_{\text{original}} - L_{s_j})$  as the reward of choose the  $j$ -th source dataset in the MAML steps denote as  $D_j$ . This reward not only fit our motivation to guide the MAML toward a good initialization point for target dataset but also incorporate the baseline in the final reward. At every MAML step  $s$ , we feed the  $(D_1^{s-1}, D_2^{s-1}, \dots, D_k^{s-1})$  into the policy network together with the probabilities  $(P_1^{s-1}, P_2^{s-1}, \dots, P_k^{s-1})$ , which is estimated in the previous step of the policy network. Then we can obtain a probability distribution over source datasets denoted as  $(P_1^s, P_2^s, \dots, P_k^s)$ . For training the feed-forward and attention layer, we calculate every source datasets reward  $\{D_j\}_{j=1}^k$  at every step.

As showed in Eq. (6), we use the sampled trajectories to get the gradient to update the policy network. It might fast converage to a local minnum and the policy will become a deterministic policy. To alleviate this issue, we incorporate  $\epsilon - greedy$  technique into the sampling process as [11].

It regards sampling process as a progressive process and sampling in the front affects the probability of sampling in the back. The most important part is to incorporate the  $\epsilon - greedy$  technique into the sampling process, to calculate the log of trajectory which is required in Eq. (6):

$$\log f_\phi(i^n) = \log \left( \prod_{g=1}^G \left( \frac{\epsilon}{k-g+1} + \frac{(1-\epsilon) * P_{i_g^s}^s}{\sum_{g=1}^G P_{i_g^s}^s - \sum_{z=1}^{g-1} P_{i_z^s}^s} \right) \right) \quad (7)$$

By setting  $\epsilon$ , we can control dataset sampling probability estimation.

The  $\epsilon - greedy$  techniques can be seen as the trade-off between exploration and exploitation. Large  $\epsilon$  indicate a high probability towards randomly sampling, which results in a high exploration rate.

Methods	en	et	id	it	sw	ta	th	tr	vi	zh	avg
<b>Zero-Shot Learning</b>											
M-BERT	62.2	55.2	57	57	51	52.2	51	53.2	59.2	64.4	56.24
MAD- $X^{mbert}$	60.6	50.6	-	-	-	-	-	<b>58.2</b>	58.2	<b>64.8</b>	-
Seq	61.6	52.2	57.6	54.4	<b>53.4</b>	52.6	52.2	51.8	58	63.8	55.76
Multi	65.2	53	57	<b>58</b>	<b>53.4</b>	54	<b>53.8</b>	53	56.2	61	56.46
Ours	<b>66.6</b>	<b>54.6</b>	<b>59</b>	54.2	49.8	<b>54.6</b>	<b>53.8</b>	54	<b>60.4</b>	60.8	<b>56.78</b>
<b>Few-Shot Learning</b>											
M-BERT	63.2	52.2	58.2	57.2	51	<b>57.2</b>	52.6	54.6	57.8	52.4	55.64
MAD- $X^{mbert}$	65.4	51.2	-	-	-	-	-	57	<b>60.2</b>	63.4	-
Seq	64	52	55.4	56.8	<b>52</b>	51.2	51.2	<b>58.6</b>	58.4	62	56.16
Multi	62.08	49.6	56.8	<b>60.8</b>	50.6	54.2	53	51.4	55	61.4	55.56
Ours	<b>68</b>	<b>54.4</b>	<b>59.8</b>	56.6	50.6	55.2	<b>54.6</b>	53.6	58.8	<b>64.2</b>	<b>57.58</b>

Table 1: Accuracy results on the XCOPA test set for zero- and few-shot MetaXCR for different target languages.

## 4 Experiments

### 4.1 Datasets

#### 4.1.1 Source Datasets

We used six commonsense reasoning datasets as our source datasets. The detail description can be found in Appendix A.

#### 4.1.2 Target Dataset

We used XCOPA as our target dataset. XCOPA [26] comprises 100 annotated instances in the validation set and 500 instances in the test set, which are translations in multiple target languages from the English COPA validation and test set. COPA [14] is a dataset that aims to benchmark causal reasoning in a binary classification setting. Since Haitian Creole (HT) and Quechua (QU) are not covered by mBERT, we didn’t report results on these two languages.

### 4.2 Baselines

We compared our method against the following baselines.

- **Sequential fine-tuning:** This baseline sequentially updates model parameters by fine-tuning models on each task. The order of tasks for fine-tuning is predefined and at each step the model is initialized with parameters learned from the previous fine-tuning.
- **Multi-task learning:** All tasks are trained simultaneously in hoping that shared representations across tasks are learned to enable the model to generalize to the target task.
- **M-BERT:** This comes from Ponti et al. [26], which use socialiaqa to fine-tune mBERT, and then continue to fine-tune mBERT on the XCOPA data of the target language.
- **Mad- $X^{mbert}$ :** This comes from Pfeiffer et al. [25], a state-of-the-art model on XCOPA. It uses a task adapter and a language adapter in each layer of BERT, and uses socialiaqa to train the task adapter, and texts in different languages to train the language adapter through the masking method. Here we reproduce their model with the original mBERT as the backbone.

## 5 Results

**Zero-Shot Learning:** We carried out experiments with MetaXCR in a zero-shot setting, where only CTML was performed as introduced in section 3.2.1. We report the impact of CTML for each target language as a difference in accuracy with and without meta-learning on top of the baseline method on the test set (Table 4.2). Overall, we observe that our zero-shot MetaXCR outperforms the baselines and results reported by Ponti et al. [26]. It establishes new state-of-the-art performance on several different target languages for zero-shot cross-lingual CR.

**Few-Shot Learning:** For few-shot learning, we used the checkpoint model of MetaXCR after CTML was trained and then fine-tuned the checkpoint model on the development set (400 instances) of




Figure 2: The test accuracy curve of the COPA for three sampling strategies

Tasks	Copa	
	Dev	Test
FIX	67.8	65.2
SAMPLE	65	64
FIRST	66.6	64.8
LAST	67.2	62.8
UNIFORM	68.8	66.6
Our	<b>70.2</b>	<b>67.4</b>

Table 3: The accuracy of different Sampling strategies

corresponding target languages. The fine-tuned model was finally evaluated on the test set. Few-shot learning results in Table 1 demonstrate that our method also achieves state-of-the-art results. Compared with the zero-shot setting, MetaXCR obtains significant improvements on 7 languages (except et, sw, ta) (from +0.8% to +3.4%) over M-BERT. Additionally, the performance of M-BERT, sequential fine-tuning and multi-task learning has dropped on most languages after few-shot learning. In contrast, our few-shot MetaXCR is better than zero-shot MetaXCR, suggesting that it benefits from the proposed cross-lingual meta-learning method.

## 6 Analyses

### 6.1 Single-task vs Multi-task Adapter

We use multiple source tasks for cross-task meta-transfer learning. We therefore want to know the effects of each task (e.g., Swag, CosmosQA etc.) and their different combinations on the target task (i.e., COPA). We conducted experiments to evaluate these effects, which are shown in Table 2. Clearly, using multiple source tasks is much better than using single source task.

### 6.2 Sampling Strategy

In this section, we explore and compare five different sampling strategies to organize the source dataset at MAML steps. 1. Sampling task uniformly (UNIFORM). 2. Using the size of the source dataset as the probability (FIX). 3. Using the softmax probability of the difference  $\{D_j\}_{j=1}^k$  (SAMPLE). 4. Using the top numbers of source based on difference (FIRST). 5. Using the last numbers of source based on difference (LAST). 6. Our reinforcement-based sampling method. All results are showed in the Table 3. From the table, we can know FIRST has an improvement over LAST and SAMPLE. In the Figure 2, we show the results of these three sampling strategies. It is obvious that the FIRST have an improvement over the LAST and SAMPLE on the whole curve, indicating the metric can guide the choose of tasks. And our proposed sampling method surpasses all others methods in the final result, which proves that by incorporating the long-term information, we can organize difference sources well and get a better result.

Tasks	COPA	
	Dev	Test
$\alpha$ NLI	58.2	56.4
CommonsenseQA	63	61.6
CosmosQA	59.8	62.4
Hellaswag	55.6	53.8
SocialQA	62.2	63.8
Swag	61.4	55.6
Three	64.4	65.2
Four	68	65.8
Five	66.6	64.2
Six	<b>68.8</b>	<b>66.6</b>

Table 2: Single-task adapter vs. multi-task adapter on COPA. For experiments with multiple source tasks, we randomly select  $n$  ( $n = 3-6$ ) tasks from all source datasets.

## 7 Conclusion

In this work, we propose a novel method named reinforcement-based meta transfer for low-resource cross-lingual cross-task commonsense reasoning on low resource setting. Our method considers multiple sources meta learning and target lingual adaptation into a unified framework, which is able to learn a general representation from multiple sources and alleviate the discrepancy between source and target by using the policy network. We demonstrate the superiority of our methods on both



supervised setting and unsupervised adaptation settings over the state-of-the-arts. In our future work, we explore to extend this approach for other low resource tasks in NLP.

## Acknowledgments and Disclosure of Funding

we would like to thank the anonymous reviewers for their insightful comments.

## References

- [1] Akari Asai, Jungo Kasai, Jonathan H. Clark, Kenton Lee, Eunsol Choi, and Hannaneh Hajishirzi. XOR QA: Cross-lingual open-retrieval question answering. In *NAACL-HLT*, 2021.
- [2] Trapit Bansal, Rishikesh Jha, and Andrew McCallum. Learning to few-shot learn across diverse natural language classification tasks. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5108–5123, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics. doi: 10.18653/v1/2020.coling-main.448. URL <https://www.aclweb.org/anthology/2020.coling-main.448>.
- [3] Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Scott Wen-tau Yi, and Yejin Choi. Abductive commonsense reasoning. *International Conference on Learning Representations*, 2020.
- [4] Yonatan Bisk, Rowan Zellers, Ronan Le bras, Jianfeng Gao, and Yejin Choi. Piqa: Reasoning about physical commonsense in natural language. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7432–7439, Apr. 2020. doi: 10.1609/aaai.v34i05.6239. URL <https://ojs.aaai.org/index.php/AAAI/article/view/6239>.
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf>.
- [6] Yi-Syuan Chen and Hong-Han Shuai. Meta-transfer learning for low-resource abstractive summarization. *CoRR*, abs/2102.09397, 2021. URL <https://arxiv.org/abs/2102.09397>.
- [7] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.747. URL <https://www.aclweb.org/anthology/2020.acl-main.747>.
- [8] Ernest Davis and Gary Marcus. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Commun. ACM*, 58(9):92–103, August 2015. ISSN 0001-0782. doi: 10.1145/2701413. URL <https://doi.org/10.1145/2701413>.
- [9] Angel Daza and Anette Frank. X-SRL: A parallel cross-lingual semantic role labeling dataset. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3904–3914, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.321. URL <https://www.aclweb.org/anthology/2020.emnlp-main.321>.

- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://www.aclweb.org/anthology/N19-1423>.
- [11] Yue Dong, Yikang Shen, Eric Crawford, Herke van Hoof, and Jackie Chi Kit Cheung. Banditsum: Extractive summarization as a contextual bandit. *arXiv preprint arXiv:1809.09672*, 2018.
- [12] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135. PMLR, 06–11 Aug 2017. URL <http://proceedings.mlr.press/v70/finn17a.html>.
- [13] Maxwell Forbes, Jena D. Hwang, Vered Shwartz, Maarten Sap, and Yejin Choi. Social chemistry 101: Learning to reason about social and moral norms. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 653–670, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.48. URL <https://www.aclweb.org/anthology/2020.emnlp-main.48>.
- [14] Andrew Gordon, Zornitsa Kozareva, and Melissa Roemmele. SemEval-2012 task 7: Choice of plausible alternatives: An evaluation of commonsense causal reasoning. In *\*SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 394–398, Montréal, Canada, 7-8 June 2012. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/S12-1052>.
- [15] Edward Grefenstette, Brandon Amos, Denis Yarats, Phu Mon Htut, Artem Molchanov, Franziska Meier, Douwe Kiela, Kyunghyun Cho, and Soumith Chintala. Generalized inner loop meta-learning. *arXiv preprint arXiv:1910.01727*, 2019.
- [16] Jianzhu Guo, Xiangyu Zhu, Chenxu Zhao, Dong Cao, Zhen Lei, and Stan Z. Li. Learning meta face recognition in unseen domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [17] Pengcheng He, Xiaodong Liu, Weizhu Chen, and Jianfeng Gao. A hybrid neural network model for commonsense reasoning. In *Proceedings of the First Workshop on Commonsense Inference in Natural Language Processing*, pages 13–21, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-6002. URL <https://www.aclweb.org/anthology/D19-6002>.
- [18] Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Cosmos QA: Machine reading comprehension with contextual commonsense reasoning. 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 2391–2401, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1243. URL <https://www.aclweb.org/anthology/D19-1243>.
- [19] Muhammad Abdullah Jamal and Guo-Jun Qi. Task agnostic meta-learning for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [20] Bill Yuchen Lin, Wangchunshu Zhou, Ming Shen, Pei Zhou, Chandra Bhagavatula, Yejin Choi, and Xiang Ren. CommonGen: A constrained text generation challenge for generative commonsense reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1823–1840, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.165. URL <https://www.aclweb.org/anthology/2020.findings-emnlp.165>.

- [21] Bo Liu, Hao Kang, Haoxiang Li, Gang Hua, and Nuno Vasconcelos. Few-shot open-set recognition using meta-learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [22] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.
- [23] Nicholas Lourie, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. UNICORN on RAINBOW: A universal commonsense reasoning model on a new multitask benchmark. *CoRR*, abs/2103.13009, 2021. URL <https://arxiv.org/abs/2103.13009>.
- [24] Todor Mihaylov and Anette Frank. Knowledgeable reader: Enhancing cloze-style reading comprehension with external commonsense knowledge. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 821–832, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1076. URL <https://www.aclweb.org/anthology/P18-1076>.
- [25] Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 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, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.617. URL <https://www.aclweb.org/anthology/2020.emnlp-main.617>.
- [26] Edoardo Maria Ponti, Goran Glavaš, Olga Majewska, Qianchu Liu, Ivan Vulić, and Anna Korhonen. XCOPA: A multilingual dataset for causal commonsense reasoning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2362–2376, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.185. URL <https://www.aclweb.org/anthology/2020.emnlp-main.185>.
- [27] Altaf Rahman and Vincent Ng. Resolving complex cases of definite pronouns: The Winograd schema challenge. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 777–789, Jeju Island, Korea, July 2012. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/D12-1071>.
- [28] Jathushan Rajasegaran, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Mubarak Shah. itaml : An incremental task-agnostic meta-learning approach. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [29] Hannah Rashkin, Maarten Sap, Emily Allaway, Noah A. Smith, and Yejin Choi. Event2Mind: Commonsense inference on events, intents, and reactions. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 463–473, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1043. URL <https://www.aclweb.org/anthology/P18-1043>.
- [30] Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8732–8740, Apr. 2020. doi: 10.1609/aaai.v34i05.6399. URL <https://ojs.aaai.org/index.php/AAAI/article/view/6399>.
- [31] Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. Social IQa: Commonsense reasoning about social interactions. 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 4463–4473, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1454. URL <https://www.aclweb.org/anthology/D19-1454>.
- [32] Sebastian Schuster, Sonal Gupta, Rushin Shah, and Mike Lewis. Cross-lingual transfer learning for multilingual task oriented dialog. 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 3795–3805, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1380. URL <https://www.aclweb.org/anthology/N19-1380>.

- [33] Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. CommonsenseQA: A question answering challenge targeting commonsense knowledge. 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 4149–4158, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1421. URL <https://www.aclweb.org/anthology/N19-1421>.
- [34] Guangting Wang, Chong Luo, Xiaoyan Sun, Zhiwei Xiong, and Wenjun Zeng. Tracking by instance detection: A meta-learning approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [35] Yubei Xiao, Ke Gong, Pan Zhou, Guolin Zheng, Xiaodan Liang, and Liang Lin. Adversarial meta sampling for multilingual low-resource speech recognition. *CoRR*, abs/2012.11896, 2020. URL <https://arxiv.org/abs/2012.11896>.
- [36] Yubei Xiao, Ke Gong, Pan Zhou, Guolin Zheng, Xiaodan Liang, and Liang Lin. Adversarial meta sampling for multilingual low-resource speech recognition. *arXiv preprint arXiv:2012.11896*, 2020.
- [37] Ming Yan, Hao Zhang, Di Jin, and Joey Tianyi Zhou. Multi-source meta transfer for low resource multiple-choice question answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7331–7341, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.654. URL <https://www.aclweb.org/anthology/2020.acl-main.654>.
- [38] Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. SWAG: A large-scale adversarial dataset for grounded commonsense inference. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 93–104, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1009. URL <https://www.aclweb.org/anthology/D18-1009>.
- [39] Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. HellaSwag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1472. URL <https://www.aclweb.org/anthology/P19-1472>.
- [40] Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. Commonsense knowledge aware conversation generation with graph attention. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 4623–4629. International Joint Conferences on Artificial Intelligence Organization, 7 2018. doi: 10.24963/ijcai.2018/643. URL <https://doi.org/10.24963/ijcai.2018/643>.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
  - (b) Did you describe the limitations of your work? [Yes]
  - (c) Did you discuss any potential negative societal impacts of your work? [Yes]
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [N/A]

- (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No]
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? [Yes]
  - (b) Did you mention the license of the assets? [N/A]
  - (c) Did you include any new assets either in the supplemental material or as a URL? [No]
  - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

## A Source datasets

- SocialIQA [31] is a collection of 38K questions for probing social situations and interactions. Each question is paired with four candidate answers.
- $\alpha$ NLI [3] contains over 20K commonsense narrative contexts and 200K explanations. It requires to identify the best explanation from two candidates.
- CosmosQA [18] is a four-choice reading comprehension dataset, focusing on understanding narrative and interpreting likely causes and effects of events.
- CommonsenseQA [33] contains 12,247 question-answer pairs. Each question has five candidate answers, aiming at investigating question answering with prior knowledge.
- Swag [38] requires to reason about situations and to anticipate what might happen next, with 113K four-choice questions.
- HellaSwag [39] uses pre-trained language models to perform adversarial filtering to collect 70K challenging examples.

## B Experimental settings

All our experiments were conducted on a single NVIDIA TiTan 24GB GPU with PyTorch. We adopted mBERT (bert-base-multilingual-cased)<sup>1</sup> as the backbone. The hidden size of the adapter is 128. Due to the memory limitation, the maximal input length of mBERT was set to 128.

Our implementation of MetaXCR was based on higher<sup>2</sup> [15]. For CTML and CLML, unless otherwise specified, a CTML meta-batch contained 3 tasks while CLML meta-batch 1 task. The batch size for each task was 12. The base model and meta-learner were both optimized with Adamw [22] optimizer, and the learning rate was set to 0.00015 and 0.0002 respectively. The number of inner gradient steps was 1, and we used accuracy convergence as a criteria for stopping.

<sup>1</sup><https://github.com/huggingface/transformers>

<sup>2</sup><https://github.com/facebookresearch/higher>

## C Comparing Cross-lingual adaptation and mono-lingual adaptation

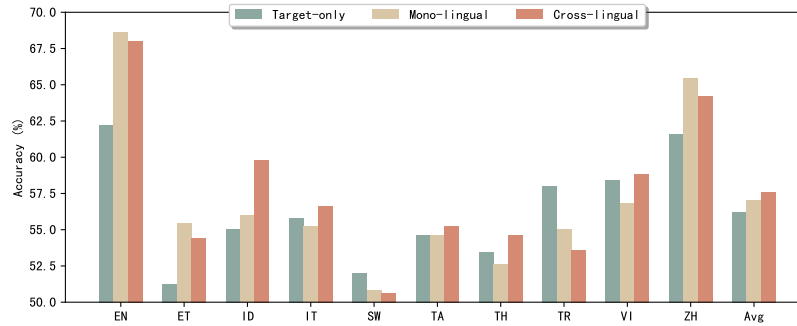


Figure 3: The accuracy results for Target-only, mono-lingual, cross-lingual training method.

In this section, we have constructed two baselines to compare our CLML with other traditional methods. The target-only method performs fine-tuning on target language development sets and mono-lingual meta learning adaptation uses the target language development set as support task and query task instead of using English data.

The experimental results are shown in the Figure 3. We observe that the cross-lingual meta adaptation method performs best in 9 languages, while the target-only method performs the worst. On the one hand, because of direct fine-tune, over-fitting problem tend to occur. On the other hand, due to the existence of the aforementioned anglo-centric bias, the effect of fine-tune is not good. For mono-lingual meta learning adaptation, there is also a certain gap with cross-lingual meta learning adaptation, which is in line with our conjecture that the model can benefit more from learning how to transfer English knowledge to other languages.