

# INFOXLM: An Information-Theoretic Framework for Cross-Lingual Language Model Pre-Training

Zewen Chi<sup>†‡\*</sup>, Li Dong<sup>‡</sup>, Furu Wei<sup>‡</sup>, Nan Yang<sup>‡</sup>, Saksham Singhal<sup>‡</sup>, Wenhui Wang<sup>‡</sup>  
Xia Song<sup>‡</sup>, Xian-Ling Mao<sup>†</sup>, Heyan Huang<sup>†</sup>, Ming Zhou<sup>‡</sup>

<sup>†</sup>Beijing Institute of Technology

<sup>‡</sup>Microsoft Corporation

{czw,maoxl,hhy63}@bit.edu.cn

{lidong1,fuwei,nanya,saksingh,wenwan,xiaso}@microsoft.com

## Abstract

In this work, we present an information-theoretic framework that formulates cross-lingual language model pre-training as maximizing mutual information between multilingual-multi-granularity texts. The unified view helps us to better understand the existing methods for learning cross-lingual representations. More importantly, inspired by the framework, we propose a new pre-training task based on contrastive learning. Specifically, we regard a bilingual sentence pair as two views of the same meaning and encourage their encoded representations to be more similar than the negative examples. By leveraging both monolingual and parallel corpora, we jointly train the pretext tasks to improve the cross-lingual transferability of pre-trained models. Experimental results on several benchmarks show that our approach achieves considerably better performance. The code and pre-trained models are available at <https://aka.ms/infoxlm>.

## 1 Introduction

Learning cross-lingual language representations plays an important role in overcoming the language barrier of NLP models. The recent success of cross-lingual language model pre-training (Devlin et al., 2019; Conneau and Lample, 2019; Conneau et al., 2020a; Chi et al., 2020; Liu et al., 2020) significantly improves the cross-lingual transferability in various downstream tasks, such as cross-lingual classification, and question answering.

State-of-the-art cross-lingual pre-trained models are typically built upon multilingual masked language modeling (MMLM; Devlin et al. 2019; Conneau et al. 2020a), and translation language modeling (TLM; Conneau and Lample 2019). The goal of both pretext tasks is to predict masked tokens given input context. The difference is that

MMLM uses monolingual text as input, while TLM feeds bilingual parallel sentences into the model. Even without explicit encouragement of learning universal representations across languages, the derived models have shown promising abilities of cross-lingual transfer.

In this work, we formulate cross-lingual pre-training from a unified information-theoretic perspective. Following the mutual information maximization principle (Hjelm et al., 2019; Kong et al., 2020), we show that the existing pretext tasks can be viewed as maximizing the lower bounds of mutual information between various multilingual-multi-granularity views.

Specifically, MMLM maximizes mutual information between the masked tokens and the context in the same language while the anchor points across languages encourages the correlation between cross-lingual contexts. Moreover, we present that TLM can maximize mutual information between the masked tokens and the parallel context, which implicitly aligns encoded representations of different languages. The unified information-theoretic framework also inspires us to propose a new cross-lingual pre-training task, named as cross-lingual contrast (XLCO). The model learns to distinguish the translation of an input sentence from a set of negative examples. In comparison to TLM that maximizes token-sequence mutual information, XLCO maximizes sequence-level mutual information between translation pairs which are regarded as cross-lingual views of the same meaning. We employ the momentum contrast (He et al., 2020) to realize XLCO. We also propose the mixup contrast and conduct the contrast on the universal layer to further facilitate the cross-lingual transferability.

Under the presented framework, we develop a cross-lingual pre-trained model (INFOXLM) to leverage both monolingual and parallel corpora. We jointly train INFOXLM with MMLM, TLM

\*Contribution during internship at Microsoft Research. Contact person: Li Dong and Furu Wei.

and XLCO. We conduct extensive experiments on several cross-lingual understanding tasks, including cross-lingual natural language inference (Conneau et al., 2018), cross-lingual question answering (Lewis et al., 2020), and cross-lingual sentence retrieval (Artetxe and Schwenk, 2019). Experimental results show that INFOXLM outperforms strong baselines on all the benchmarks. Moreover, the analysis indicates that INFOXLM achieves better cross-lingual transferability.

## 2 Related Work

### 2.1 Cross-Lingual LM Pre-Training

Multilingual BERT (mBERT; Devlin et al. 2019) is pre-trained with the multilingual masked language modeling (MMLM) task on the monolingual text. mBERT produces cross-lingual representations and performs cross-lingual tasks surprisingly well (Wu and Dredze, 2019). XLM (Conneau and Lample, 2019) extends mBERT with the translation language modeling (TLM) task so that the model can learn cross-lingual representations from parallel corpora. Unicoder (Huang et al., 2019) tries several pre-training tasks to utilize parallel corpora. ALM (Yang et al., 2020) extends TLM to code-switched sequences obtained from translation pairs. XLM-R (Conneau et al., 2020a) scales up MMLM pre-training with larger corpus and longer training. LaBSE (Feng et al., 2020) learns cross-lingual sentence embeddings by an additive translation ranking loss.

In addition to learning cross-lingual encoders, several pre-trained models focus on generation. MASS (Song et al., 2019) and mBART (Liu et al., 2020) pretrain sequence-to-sequence models to improve machine translation. XNLG (Chi et al., 2020) focuses on the cross-lingual transfer of language generation, such as cross-lingual question generation, and abstractive summarization.

### 2.2 Mutual Information Maximization

Various methods have successfully learned visual or language representations by maximizing mutual information between different views of input. It is difficult to directly maximize mutual information. In practice, the methods resort to a tractable lower bound as the estimator, such as InfoNCE (Oord et al., 2018), and the variational form of the KL divergence (Nguyen et al., 2010). The estimators are also known as contrastive learning (Arora et al., 2019) that measures the representation similarities

between the sampled positive and negative pairs. In addition to the estimators, various view pairs are employed in these methods. The view pair can be the local and global features of an image (Hjelm et al., 2019; Bachman et al., 2019), the random data augmentations of the same image (Tian et al., 2019; He et al., 2020; Chen et al., 2020), or different parts of a sequence (Oord et al., 2018; Henaff, 2020; Kong et al., 2020). Kong et al. (2020) show that learning word embeddings or contextual embeddings can also be unified under the framework of mutual information maximization.

## 3 Information-Theoretic Framework for Cross-Lingual Pre-Training

In representation learning, the learned representations are expected to preserve the information of the original input data. However, it is intractable to directly model the mutual information between the input data and the representations. Alternatively, we can maximize the mutual information between the representations from different views of the input data, e.g., different parts of a sentence, a translation pair of the same meaning.

In this section, we start from a unified information-theoretic perspective, and formulate cross-lingual pre-training with the mutual information maximization principle. Then, under the information-theoretic framework, we propose a new cross-lingual pre-training task, named as cross-lingual contrast (XLCO). Finally, we present the pre-training procedure of our INFOXLM.

### 3.1 Multilingual Masked Language Modeling

The goal of multilingual masked language modeling (MMLM; Devlin et al. 2019) is to recover the masked tokens from a randomly masked sequence. For each input sequence of MMLM, we sample a text from the monolingual corpus for pre-training. Let  $(c_1, x_1)$  denote a monolingual text sequence, where  $x_1$  is the masked token, and  $c_1$  is the corresponding context. Intuitively, we need to maximize their dependency (i.e.,  $I(c_1; x_1)$ ), so that the context representations are predictive for masked tokens (Kong et al., 2020).

For the example pair  $(c_1, x_1)$ , we construct a set  $\mathcal{N}$  that contains  $x_1$  and  $|\mathcal{N}| - 1$  negative samples drawn from a proposal distribution  $q$ . According to the InfoNCE (Oord et al., 2018) lower bound, we

have:

$$I(c_1; x_1) \geq \frac{E}{q(\mathcal{N})} \left[ \log \frac{f_{\theta}(c_1, x_1)}{\sum_{x' \in \mathcal{N}} f_{\theta}(c_1, x')} \right] + \log |\mathcal{N}| \quad (1)$$

where  $f_{\theta}$  is a function that scores whether the input  $c_1$  and  $x_1$  is a positive pair.

Given context  $c_1$ , MMLM learns to minimize the cross-entropy loss of the masked token  $x_1$ :

$$\mathcal{L}_{\text{MMLM}} = -\log \frac{\exp(g_{\theta_T}(c_1)^{\top} g_{\theta_E}(x_1))}{\sum_{x' \in \mathcal{V}} \exp(g_{\theta_T}(c_1)^{\top} g_{\theta_E}(x'))} \quad (2)$$

where  $\mathcal{V}$  is the vocabulary,  $g_{\theta_E}$  is a look-up function that returns the token embeddings,  $g_{\theta_T}$  is a Transformer that returns the final hidden vectors in position of  $x_1$ . According to Equation (1) and Equation (2), if  $\mathcal{N} = \mathcal{V}$  and  $f_{\theta}(c_1, x_1) = \exp(g_{\theta_T}(c_1)^{\top} g_{\theta_E}(x_1))$ , we can find that MMLM maximizes a lower bound of  $I(c_1; x_1)$ .

Next, we explain why MMLM can implicitly learn cross-lingual representations. Let  $(c_2, x_2)$  denote a MMLM instance that is in different language as  $(c_1, x_1)$ . Because the vocabulary, the position embedding, and special tokens are shared across languages, it is common to find anchor points (Pires et al., 2019; Dufter and Schütze, 2020) where  $x_1 = x_2$  (such as subword, punctuation, and digit) or  $I(x_1, x_2)$  is positive (i.e., the representations are associated or isomorphic). With the bridge effect of  $\{x_1, x_2\}$ , MMLM obtains a v-structure dependency “ $c_1 \rightarrow \{x_1, x_2\} \leftarrow c_2$ ”, which leads to a negative co-information (i.e., interaction information)  $I(c_1; c_2; \{x_1, x_2\})$  (Tsujishita, 1995). Specifically, the negative value of  $I(c_1; c_2; \{x_1, x_2\})$  indicates that the variable  $\{x_1, x_2\}$  enhances the correlation between  $c_1$  and  $c_2$  (Fano, 1963).

In summary, although MMLM learns to maximize  $I(c_1, x_1)$  and  $I(c_2, x_2)$  in each language, we argue that the task encourages the cross-lingual correlation of learned representations. Notice that for the setting without word-piece overlap (Artetxe et al., 2020; Conneau et al., 2020b; K et al., 2020), we hypothesize that the information bottleneck principle (Tishby and Zaslavsky, 2015) tends to transform the cross-lingual structural similarity into isomorphic representations, which has similar bridge effects as the anchor points. Then we can explain how the cross-lingual ability is spread out as above. We leave more discussions about the setting without word-piece overlap for future work.

### 3.2 Translation Language Modeling

Similar to MMLM, the goal of translation language modeling (TLM; Conneau and Lample 2019) is also to predict masked tokens, but the prediction is conditioned on the concatenation of a translation pair. We try to explain how TLM pre-training enhances cross-lingual transfer from an information-theoretic perspective.

Let  $c_1$  and  $c_2$  denote a translation pair of sentences, and  $x_1$  a masked token taken in  $c_1$ . So  $c_1$  and  $x_1$  are in the same language, while  $c_1$  and  $c_2$  are in different ones. Following the derivations of MMLM in Section 3.1, the objective of TLM is maximizing the lower bound of mutual information  $I(c_1, c_2; x_1)$ . By re-writing the above mutual information, we have:

$$I(c_1, c_2; x_1) = I(c_1; x_1) + I(c_2; x_1 | c_1) \quad (3)$$

The first term  $I(c_1; x_1)$  corresponds to MMLM, which learns to use monolingual context. In contrast, the second term  $I(c_2; x_1 | c_1)$  indicates cross-lingual mutual information between  $c_2$  and  $x_1$  that is not included by  $c_1$ . In other words,  $I(c_2; x_1 | c_1)$  encourages the model to predict masked tokens by using the context in a different language. In conclusion, TLM learns to utilize the context in both languages, which implicitly improves the cross-lingual transferability of pre-trained models.

### 3.3 Cross-Lingual Contrastive Learning

Inspired by the unified information-theoretic framework, we propose a new cross-lingual pre-training task, named as cross-lingual contrast (XLCO). The goal of XLCO is to maximize mutual information between the representations of parallel sentences  $c_1$  and  $c_2$ , i.e.,  $I(c_1, c_2)$ . Unlike maximizing token-sequence mutual information in MMLM and TLM, XLCO targets at cross-lingual sequence-level mutual information.

We describe how the task is derived as follows. Using InfoNCE (Oord et al., 2018) as the lower bound, we have:

$$I(c_1; c_2) \geq \frac{E}{q(\mathcal{N})} \left[ \log \frac{f_{\theta}(c_1, c_2)}{\sum_{c' \in \mathcal{N}} f_{\theta}(c_1, c')} \right] + \log |\mathcal{N}| \quad (4)$$

where  $\mathcal{N}$  is a set that contains the positive pair  $c_2$  and  $|\mathcal{N}| - 1$  negative samples. In order to maximize the lower bound of  $I(c_1; c_2)$ , we need to design the function  $f_{\theta}$  that measures the similarity between the input sentence and the proposal distribution

$q(\mathcal{N})$ . Specifically, we use the following similarity function  $f_\theta$ :

$$f_\theta(c_1, c_2) = \exp(g_\theta(c_1)^\top g_\theta(c_2)) \quad (5)$$

where  $g_\theta$  is the Transformer encoder that we are pre-training. Following (Devlin et al., 2019), a special token [CLS] is added to the input, whose hidden vector is used as the sequence representation. Additionally, we use a linear projection head after the encoder in  $g_\theta$ .

**Momentum Contrast** Another design choice is how to construct  $\mathcal{N}$ . As shown in Equation (4), a large  $|\mathcal{N}|$  improves the tightness of the lower bound, which has been proven to be critical for contrastive learning (Chen et al., 2020).

In our work, we employ the momentum contrast (He et al., 2020) to construct the set  $\mathcal{N}$ , where the previously encoded sentences are progressively reused as negative samples. Specifically, we construct two encoders with the same architecture which are the query encoder  $g_{\theta_Q}$  and the key encoder  $g_{\theta_K}$ . The loss function of XLCO is:

$$\mathcal{L}_{\text{XLCO}} = -\log \frac{\exp(g_{\theta_Q}(c_1)^\top g_{\theta_K}(c_2))}{\sum_{c' \in \mathcal{N}} \exp(g_{\theta_Q}(c_1)^\top g_{\theta_K}(c'))} \quad (6)$$

During training, the query encoder  $g_{\theta_Q}$  encodes  $c_1$  and is updated by backpropagation. The key encoder  $g_{\theta_K}$  encodes  $\mathcal{N}$  and is learned with momentum update (He et al., 2020) towards the query encoder. The negative examples in  $\mathcal{N}$  are organized as a queue, where a newly encoded example is added while the oldest one is popped from the queue. We initialize the query encoder and the key encoder with the same parameters, and pre-fill the queue with a set of encoded examples until it reaches the desired size  $|\mathcal{N}|$ . Notice that the size of the queue remains constant during training.

**Mixup Contrast** For each pair, we concatenate it with a randomly sampled translation pair from another parallel corpus. For example, consider the pairs  $\langle c_1, c_2 \rangle$  and  $\langle d_1, d_2 \rangle$  sampled from two different parallel corpora. The two pairs are concatenated in a random order, such as  $\langle c_1 d_1, c_2 d_2 \rangle$ , and  $\langle c_1 d_2, d_1 c_2 \rangle$ . The data augmentation of mixup encourages pre-trained models to learn sentence boundaries and to distinguish the order of multilingual texts.

**Contrast on Universal Layer** As a pre-training task maximizing the lower bound of sequence-level mutual information, XLCO is usually jointly learned with token-sequence tasks, such as MMLM, and TLM. In order to make XLCO more compatible with the other pretext tasks, we propose to conduct contrastive learning on the most universal (or transferable) layer in terms of MMLM and TLM.

In our implementations, we instead use the hidden vectors of [CLS] at layer 8 to perform contrastive learning for base-size (12 layers) models, and layer 12 for large-size (24 layers) models. Because previous analysis (Sabet et al., 2020; Dufter and Schütze, 2020; Conneau et al., 2020b) shows that the specific layers of MMLM learn more universal representations and work better on cross-lingual retrieval tasks than other layers. We choose the layers following the same principle.

The intuition behind the method is that MMLM and TLM encourage the last layer to produce language-distinguishable token representations because of the masked token classification. But XLCO tends to learn similar representations across languages. So we do not directly use the hidden states of the last layer in XLCO.

### 3.4 Cross-Lingual Pre-Training

We pretrain a cross-lingual model INFOXLM by jointly maximizing the lower bounds of three types of mutual information, including monolingual token-sequence mutual information (MMLM), cross-lingual token-sequence mutual information (TLM), and cross-lingual sequence-level mutual information (XLCO). Formally, the loss of cross-lingual pre-training in INFOXLM is defined as:

$$\mathcal{L} = \mathcal{L}_{\text{MMLM}} + \mathcal{L}_{\text{TLM}} + \mathcal{L}_{\text{XLCO}} \quad (7)$$

where we apply the same weight for the loss terms.

Both TLM and XLCO use parallel data. The number of bilingual pairs increases with the square of the number of languages. In our work, we set English as the pivot language following (Conneau and Lample, 2019), i.e., we only use the parallel corpora that contain English.

In order to balance the data size between high-resource and low-resource languages, we apply a multilingual sampling strategy (Conneau and Lample, 2019) for both monolingual and parallel data. An example in the language  $l$  is sampled with the probability  $p_l \propto (n_l/n)^{0.7}$ , where  $n_l$  is the number

of instances in the language  $l$ , and  $n$  refers to the total number of data. Empirically, the sampling algorithm alleviates the bias towards high-resource languages (Conneau et al., 2020a).

## 4 Experiments

In this section, we first present the training configuration of INFOXLM. Then we compare the fine-tuning results of INFOXLM with previous work on three cross-lingual understanding tasks. We also conduct ablation studies to understand the major components of INFOXLM.

### 4.1 Setup

**Corpus** We use the same pre-training corpora as previous models (Conneau et al., 2020a; Conneau and Lample, 2019). Specifically, we reconstruct CC-100 (Conneau et al., 2020a) for MMLM, which remains 94 languages by filtering the language code larger than 0.1GB. Following (Conneau and Lample, 2019), for the TLM and XLCO tasks, we employ 14 language pairs of parallel data that involves English. We collect translation pairs from MultiUN (Ziemski et al., 2016), IIT Bombay (Kunchukuttan et al., 2018), OPUS (Tiedemann, 2012), and WikiMatrix (Schwenk et al., 2019). The size of parallel corpora is about 42GB. More details about the pre-training data are described in the appendix.

**Model Size** We follow the model configurations of XLM-R (Conneau et al., 2020a). For the Transformer (Vaswani et al., 2017) architecture, we use 12 layers and 768 hidden states for INFOXLM (i.e., base size), and 24 layers and 1,024 hidden states for INFOXLM<sub>LARGE</sub> (i.e., large size).

**Hyperparameters** We initialize the parameters of INFOXLM with XLM-R. We optimize the model with Adam (Kingma and Ba, 2015) using a batch size of 2048 for a total of 150K steps for INFOXLM, and 200K steps for INFOXLM<sub>LARGE</sub>. The same number of training examples are fed to three tasks. The learning rate is scheduled with a linear decay with 10K warmup steps, where the peak learning rate is set as 0.0002 for INFOXLM, and 0.0001 for INFOXLM<sub>LARGE</sub>. The momentum coefficient is set as 0.9999 and 0.999 for INFOXLM and INFOXLM<sub>LARGE</sub>, respectively. The length of the queue is set as 131,072. The training procedure takes about 2.3 days  $\times$  2 Nvidia DGX-2 stations for INFOXLM, and 5 days  $\times$  16

Nvidia DGX-2 stations for INFOXLM<sub>LARGE</sub>. Details about the pre-training hyperparameters can be found in the appendix.

### 4.2 Evaluation

We conduct experiments over three cross-lingual understanding tasks, i.e., cross-lingual natural language inference, cross-lingual sentence retrieval, and cross-lingual question answering.

#### Cross-Lingual Natural Language Inference

The Cross-Lingual Natural Language Inference corpus (XNLI; Conneau et al. 2018) is a widely used cross-lingual classification benchmark. The goal of NLI is to identify the relationship of an input sentence pair. We evaluate the models under the following two settings. (1) Cross-Lingual Transfer: fine-tuning the model with English training set and directly evaluating on multilingual test sets. (2) Translate-Train-All: fine-tuning the model with the English training data and the pseudo data that are translated from English to the other languages.

**Cross-Lingual Sentence Retrieval** The goal of the cross-lingual sentence retrieval task is to extract parallel sentences from bilingual comparable corpora. We use the subset of 36 language pairs of the Tatoeba dataset (Artetxe and Schwenk, 2019) for the task. The dataset is collected from Tatoeba<sup>1</sup>, which is an open collection of multilingual parallel sentences in more than 300 languages. Following (Hu et al., 2020), we use the averaged hidden vectors in the seventh Transformer layer to compute cosine similarity for sentence retrieval.

**Cross-Lingual Question Answering** We use the Multilingual Question Answering (MLQA; Lewis et al. 2020) dataset for the cross-lingual QA task. MLQA provides development and test data in seven languages in the format of SQuAD v1.1 (Rajpurkar et al., 2016). We follow the fine-tuning method introduced in (Devlin et al., 2019) that concatenates the question-passage pair as the input.

### 4.3 Results

We compare INFOXLM with the following pre-trained Transformer models: (1) Multilingual BERT (MBERT; Devlin et al. 2019) is pre-trained with MMLM on Wikipedia in 102 languages; (2) XLM (Conneau and Lample, 2019) pretrains both MMLM and TLM tasks on Wikipedia in 100

<sup>1</sup><https://tatoeba.org/eng/>

Models	#M	en	fr	es	de	el	bg	ru	tr	ar	vi	th	zh	hi	sw	ur	Avg
<i>Fine-tune multilingual model on English training set (Cross-lingual Transfer)</i>																	
mBERT*	N	82.1	73.8	74.3	71.1	66.4	68.9	69.0	61.6	64.9	69.5	55.8	69.3	60.0	50.4	58.0	66.3
XLM (w/o TLM)*	N	83.7	76.2	76.6	73.7	72.4	73.0	72.1	68.1	68.4	72.0	68.2	71.5	64.5	58.0	62.4	71.3
XLM*	N	85.0	78.7	78.9	77.8	76.6	77.4	75.3	72.5	73.1	76.1	73.2	76.5	69.6	68.4	67.3	75.1
XLM (w/o TLM)*	1	83.2	76.7	77.7	74.0	72.7	74.1	72.7	68.7	68.6	72.9	68.9	72.5	65.6	58.2	62.4	70.7
UNICODER	1	85.4	79.2	79.8	78.2	77.3	78.5	76.7	73.8	73.9	75.9	71.8	74.7	70.1	67.4	66.3	75.3
XLM-R*	1	85.8	79.7	80.7	78.7	77.5	79.6	78.1	74.2	73.8	76.5	74.6	76.7	72.4	66.5	68.3	76.2
XLM-R (reimpl)	1	84.7	79.1	79.4	77.4	76.6	78.4	76.0	73.5	72.6	75.5	73.0	74.5	71.0	65.7	67.6	75.0
INFOXML	1	86.4	80.3	80.9	79.3	77.8	79.3	77.6	75.6	74.2	77.1	74.6	77.0	72.2	67.5	67.3	<b>76.5</b>
–XLCO	1	86.5	80.5	80.3	78.7	77.3	78.8	77.4	74.6	73.8	76.8	73.7	76.7	71.8	66.3	66.4	76.0
XLM-R <sub>LARGE</sub> *	1	89.1	84.1	85.1	83.9	82.9	84.0	81.2	79.6	79.8	80.8	78.1	80.2	76.9	73.9	73.8	80.9
XLM-R <sub>LARGE</sub> (reimpl)	1	88.9	83.6	84.8	83.1	82.4	83.7	80.7	79.2	79.0	80.4	77.8	79.8	76.8	72.7	73.3	80.4
INFOXML <sub>LARGE</sub>	1	89.7	84.5	85.5	84.1	83.4	84.2	81.3	80.9	80.4	80.8	78.9	80.9	77.9	74.8	73.7	<b>81.4</b>
<i>Fine-tune multilingual model on all training sets (Translate-Train-All)</i>																	
XLM (w/o TLM)*	1	84.5	80.1	81.3	79.3	78.6	79.4	77.5	75.2	75.6	78.3	75.7	78.3	72.1	69.2	67.7	76.9
XLM*	1	85.0	80.8	81.3	80.3	79.1	80.9	78.3	75.6	77.6	78.5	76.0	79.5	72.9	72.8	68.5	77.8
XLM-R*	1	85.4	81.4	82.2	80.3	80.4	81.3	79.7	78.6	77.3	79.7	77.9	80.2	76.1	73.1	73.0	79.1
XLM-R (reimpl)	1	85.0	81.0	81.9	80.6	79.7	81.4	79.5	77.7	77.3	79.5	77.5	79.1	75.3	72.2	70.9	78.6
INFOXML	1	86.5	82.6	83.0	82.3	81.3	82.4	80.6	79.5	78.9	81.0	78.9	80.7	77.8	73.3	71.6	<b>80.0</b>

Table 1: Evaluation results on XNLI cross-lingual natural language inference. We report test accuracy in 15 languages. The model number #M=N indicates the model selection is done on each language’s validation set (i.e., each language has a different model), while #M=1 means only one model is used for all languages. Results with “\*” are taken from [Conneau et al. \(2020a\)](#). “(reimpl)” is our reimplementation of fine-tuning, which is the same as INFOXML. Results of INFOXML and XLM-R (reimpl) are averaged over five runs. “–XLCO” is the model without cross-lingual contrast.

languages; (3) XLM-R ([Conneau et al., 2020a](#)) scales up MMLM to the large CC-100 corpus in 100 languages with much more training steps; (4) UNICODER ([Liang et al., 2020](#)) continues training XLM-R with MMLM and TLM. (5) INFOXML–XLCO continues training XLM-R with MMLM and TLM, using the same pre-training datasets with INFOXML.

### Cross-Lingual Natural Language Inference

Table 1 reports the classification accuracy on each test of XNLI under the above evaluation settings. The final scores on test set are averaged over five random seeds. INFOXML outperforms all baseline models on the two evaluation settings of XNLI. In the cross-lingual transfer setting, INFOXML achieves 76.5 averaged accuracy, outperforming XLM-R (reimpl) by 1.5. Similar improvements can be observed for large-size models. Moreover, the ablation results “–XLCO” show that cross-lingual contrast is helpful for zero-shot transfer in most languages. We also find that INFOXML improves the results in the translate-train-all setting.

**Cross-Lingual Sentence Retrieval** In Table 2 and Table 3, we report the top-1 accuracy scores of cross-lingual sentence retrieval with the base-size

models. The evaluation results demonstrate that INFOXML produces better aligned cross-lingual sentence representations. On the 14 language pairs that are covered by parallel data, INFOXML obtains 77.8 and 80.6 averaged top-1 accuracies in the directions of  $xx \rightarrow en$  and  $en \rightarrow xx$ , outperforming XLM-R by 20.2 and 21.1. Even on the 22 language pairs that are not covered by parallel data, INFOXML outperforms XLM-R on 16 out of 22 language pairs, providing 8.1% improvement in averaged accuracy. In comparison, the ablation variant “–XLCO” (i.e., MMLM+TLM) obtains better results than XLM-R in Table 2, while getting worse performance than XLM-R in Table 3. The results indicate that XLCO encourages the model to learn universal representations even on the language pairs without parallel supervision.

**Cross-Lingual Question Answering** Table 4 compares INFOXML with baseline models on MLQA, where we report the F1 and the exact match (EM) scores on each test set. Both INFOXML and INFOXML<sub>LARGE</sub> obtain the best results against the four baselines. In addition, the results of the ablation variant “–XLCO” indicate that the proposed cross-lingual contrast is beneficial on MLQA.

Models	Direction	ar	bg	zh	de	el	fr	hi	ru	es	sw	th	tr	ur	vi	Avg
XLM-R	xx → en	36.8	67.6	60.7	89.9	53.7	74.1	54.2	72.5	74.0	18.7	38.3	61.1	36.6	68.4	57.6
INFOXLM	xx → en	59.0	78.6	86.3	93.9	62.1	79.4	87.1	83.8	88.2	39.5	84.9	83.3	73.0	89.6	<b>77.8</b>
-XlCo	xx → en	42.9	65.5	69.5	91.1	55.6	76.4	71.6	74.9	74.8	20.5	68.1	69.8	51.6	81.8	65.3
XLM-R	en → xx	38.6	69.9	60.3	89.4	57.3	74.3	49.3	73.0	74.6	14.4	58.4	64.0	36.9	72.5	59.5
INFOXLM	en → xx	68.6	78.6	86.4	95.1	72.6	84.0	88.3	85.7	87.2	40.8	91.2	84.7	73.3	92.0	<b>80.6</b>
-XlCo	en → xx	45.4	64.0	69.3	88.1	56.5	72.3	69.6	73.6	71.5	22.1	79.7	64.3	48.2	79.8	64.6

Table 2: Evaluation results on Tatoeba cross-lingual sentence retrieval. We report the top-1 accuracy of 14 language pairs that are covered by parallel data.

Models	Direction	af	bn	et	eu	fi	he	hu	id	it	jv	ja	ka	kk	ko	ml	mr	nl	fa	pt	ta	te	tl	Avg
XLM-R	xx → en	55.2	29.3	49.3	33.5	66.7	53.9	61.6	70.8	68.2	15.1	57.2	41.4	40.3	51.6	56.5	46.0	79.5	68.0	80.6	25.7	32.5	31.2	50.6
INFOXLM	xx → en	48.6	49.6	38.3	36.7	65.7	62.9	61.7	79.9	72.2	13.2	78.3	57.4	49.2	74.5	76.6	72.0	80.8	82.2	84.7	53.7	53.0	42.1	<b>60.6</b>
-XlCo	xx → en	33.1	33.5	25.9	20.8	48.4	49.1	46.1	68.5	60.4	12.2	60.6	38.6	35.1	60.6	57.8	49.1	72.2	66.0	75.3	36.5	38.0	25.5	46.1
XLM-R	en → xx	55.0	27.9	50.2	32.5	72.9	63.2	67.1	71.9	68.0	9.8	58.2	52.0	41.7	58.3	60.8	42.1	78.9	69.6	82.1	33.2	38.9	29.7	52.9
INFOXLM	en → xx	51.8	49.1	35.2	28.6	65.6	66.5	61.7	80.1	72.8	7.8	80.4	61.9	50.6	79.6	78.7	68.1	81.8	82.8	86.5	63.5	53.0	35.5	<b>61.0</b>
-XlCo	en → xx	28.1	23.5	19.0	12.6	45.2	49.7	40.8	62.8	57.5	3.4	58.2	38.9	31.3	61.0	57.5	37.2	67.8	66.4	75.0	43.0	31.6	17.9	42.2

Table 3: Evaluation results on Tatoeba cross-lingual sentence retrieval. We report the top-1 accuracy scores of 22 language pairs that are not covered by parallel data.

#### 4.4 Analysis and Discussion

To understand INFOXLM and the cross-lingual contrast task more deeply, we conduct analysis from the perspectives of cross-lingual transfer and cross-lingual representations. Furthermore, we perform comprehensive ablation studies on the major components of INFOXLM, including the cross-lingual pre-training tasks, mixup contrast, the contrast layer, and the momentum contrast. To reduce the computation load, we use INFOXLM15 in our ablation studies, which is trained on 15 languages for 100K steps.

**Cross-Lingual Transfer Gap** Cross-lingual transfer gap (Hu et al., 2020) is the difference between the performance on the English test set and the averaged performance on the test sets of all other languages. A lower cross-lingual transfer gap score indicates more end-task knowledge from the English training set is transferred to other languages. In Table 5, we compare the cross-lingual transfer gap scores of INFOXLM with baseline models on MLQA and XNLI. Note that we do not include the results of XLM because it is pre-trained on 15 languages or using #M=N. The results show that INFOXLM reduces the gap scores on both MLQA and XNLI, providing better cross-lingual transferability than the baselines.

**Cross-Lingual Representations** In addition to cross-lingual transfer, learning good cross-lingual representations is also the goal of cross-lingual pre-




Figure 1: Evaluation results of different layers on Tatoeba cross-lingual sentence retrieval.

training. In order to analyze how the cross-lingual contrast task affects the alignment of the learned cross-lingual representations, we evaluate the representations of different middle layers on the Tatoeba test sets of the 14 languages that are covered by parallel data. Figure 1 presents the averaged top-1 accuracy of cross-lingual sentence retrieval in the direction of xx → en. INFOXLM outperforms XLM-R on all of the 12 layers, demonstrating that our proposed task improves the cross-lingual alignment of the learned representations. From the results of XLM-R, we observe that the model suffers from a performance drop in the last few layers. The reason is that MMLM encourages the representations of the last hidden layer to be similar to token embeddings, which is contradictory with the goal of learning cross-lingual representations. In

Models	en	es	de	ar	hi	vi	zh	Avg
MBERT*	77.7 / 65.2	64.3 / 46.6	57.9 / 44.3	45.7 / 29.8	43.8 / 29.7	57.1 / 38.6	57.5 / 37.3	57.7 / 41.6
XML*	74.9 / 62.4	68.0 / 49.8	62.2 / 47.6	54.8 / 36.3	48.8 / 27.3	61.4 / 41.8	61.1 / 39.6	61.6 / 43.5
UNICODER	80.6 / -	68.6 / -	62.7 / -	57.8 / -	62.7 / -	67.5 / -	62.1 / -	66.0 / -
XML-R	77.1 / 64.6	67.4 / 49.6	60.9 / 46.7	54.9 / 36.6	59.4 / 42.9	64.5 / 44.7	61.8 / 39.3	63.7 / 46.3
XML-R (reimpl)	80.2 / 67.0	67.7 / 49.9	62.1 / 47.7	56.1 / 37.2	61.1 / 44.0	67.0 / 46.3	61.4 / 38.5	65.1 / 47.2
INFOXML	81.6 / 68.3	69.8 / 51.6	64.3 / 49.4	60.6 / 40.9	65.2 / 47.1	70.2 / 49.0	64.8 / 41.3	<b>68.1 / 49.6</b>
-XLCO	81.2 / 68.1	69.6 / 51.9	64.0 / 49.3	59.7 / 40.2	64.0 / 46.3	69.3 / 48.0	64.1 / 40.6	67.4 / 49.2
XML-R <sub>LARGE</sub>	80.6 / 67.8	74.1 / 56.0	68.5 / 53.6	63.1 / 43.5	69.2 / 51.6	71.3 / 50.9	68.0 / 45.4	70.7 / 52.7
XML-R <sub>LARGE</sub> (reimpl)	84.0 / 71.1	74.4 / 56.4	70.2 / 55.0	66.5 / 46.3	71.1 / 53.2	74.4 / 53.5	68.6 / 44.6	72.7 / 54.3
INFOXML <sub>LARGE</sub>	84.5 / 71.6	75.1 / 57.3	71.2 / 56.2	67.6 / 47.6	72.5 / 54.2	75.2 / 54.1	69.2 / 45.4	<b>73.6 / 55.2</b>

Table 4: Evaluation results on MLQA cross-lingual question answering. We report the F1 and exact match (EM) scores. Results with “\*” are taken from (Lewis et al., 2020). “(reimpl)” is our reimplementation of fine-tuning, which is the same as INFOXML. Results of INFOXML and XML-R (reimpl) are averaged over five runs. “-XLCO” is the model without cross-lingual contrast.

Models	MLQA	XNLI	Average
MBERT	23.3	16.9	20.1
XML-R	17.6	10.4	14.0
INFOXML	<b>15.8</b>	<b>10.3</b>	<b>13.1</b>
-XLCO	16.1	11.0	13.6

Table 5: Cross-lingual transfer gap scores, i.e., averaged performance drop between English and other languages in zero-shot transfer. Smaller gap indicates better transferability. “-XLCO” is the model without cross-lingual contrast.

contrast, INFOXML still provides high retrieval accuracy at the last few layers, which indicates that INFOXML provides better aligned representations than XML-R. Moreover, we find that the performance is further improved when removing TLM, demonstrating that XLCO is more effective than TLM for aligning cross-lingual representations, although TLM helps to improve zero-shot cross-lingual transfer.

**Effect of Cross-Lingual Pre-training Tasks** To better understand the effect of the cross-lingual pre-training tasks, we perform ablation studies on the pre-training tasks of INFOXML, by removing XLCO, TLM, or both. We present the experimental results in Table 7. Comparing the results of -TLM and -XLCO with the results of -TLM-XLCO, we find that both XLCO and TLM effectively improve cross-lingual transferability of the pre-trained INFOXML model. TLM is more effective for XNLI while XLCO is more effective for MLQA. Moreover, the performance can be further improved by jointly learning XLCO and TLM.

**Effect of Contrast on Universal Layer** We conduct experiments to investigate whether contrast

Model	XLCO Layer	XNLI	MLQA
INFOXML15	8	76.45	67.87 / 49.58
INFOXML15	12	76.12	67.83 / 49.50
INFOXML15-TLM	8	75.58	67.42 / 49.27
INFOXML15-TLM	12	75.85	67.84 / 49.54

Table 6: Contrast on the universal layer v.s. on the last layer. Results are averaged over five runs. “-TLM” is the ablation variant without TLM.

	Model	XNLI	MLQA
[0]	INFOXML15	<b>76.45</b>	<b>67.87 / 49.58</b>
[1]	[0]-XLCO	76.24	67.43 / 49.23
[2]	[0]-TLM	75.85	67.84 / 49.54
[3]	[2]-XLCO	75.33	66.86 / 48.82
[4]	[2]-Mixup	75.43	67.21 / 49.19
[5]	[2]-Momentum	75.32	66.58 / 48.66

Table 7: Ablation results on components of INFOXML. Results are averaged over five runs.

on the universal layer improves cross-lingual pre-training. As shown in Table 6, we compare the evaluation results of four variants of INFOXML, where XLCO is applied on the layer 8 (i.e., universal layer) or on the layer 12 (i.e., the last layer). We find that contrast on the layer 8 provides better results for INFOXML. However, conducting XLCO on layer 12 performs better when the TLM task is excluded. The results show that maximizing context-sequence (TLM) and sequence-level (XLCO) mutual information at the last layer tends to interfere with each other. Thus, we suggest applying XLCO on the universal layer for pre-training INFOXML.

**Effect of Mixup Contrast** We conduct an ablation study on the mixup contrast strategy. We pre-



train a model that directly uses translation pairs for XLCO without mixup contrast (–TLM–Mixup). As shown in Table 7, we present the evaluation results on XNLI and MLQA. We observe that mixup contrast improves the performance of INFOXLM on both datasets.

**Effect of Momentum Contrast** In order to show whether our pre-trained model benefits from momentum contrast, we pretrain a revised version of INFOXLM without momentum contrast. In other words, the parameters of the key encoder are always the same as the query encoder. As shown in Table 7, we report evaluation results (indicated by “–TLM–Momentum”) of removing momentum contrast on XNLI and MLQA. We observe a performance descent after removing the momentum contrast from INFOXLM, which indicates that momentum contrast improves the learned language representations of INFOXLM.

## 5 Conclusion

In this paper, we present a cross-lingual pre-trained model INFOXLM that is trained with both monolingual and parallel corpora. The model is motivated by the unified view of cross-lingual pre-training from an information-theoretic perspective. Specifically, in addition to the masked language modeling and translation language modeling tasks, INFOXLM is jointly pre-trained with a newly introduced cross-lingual contrastive learning task. The cross-lingual contrast leverages bilingual pairs as the two views of the same meaning, and encourages their encoded representations to be more similar than the negative examples. Experimental results on several cross-lingual language understanding tasks show that INFOXLM can considerably improve the performance.

## 6 Ethical Considerations

Currently, most NLP research works and applications are English-centric, which makes non-English users hard to access to NLP-related services. Our work focuses on cross-lingual language model pre-training. With the pre-trained model, we are able to transfer end-task knowledge from high-resource languages to low-resource languages, which helps to build more accessible NLP applications. Additionally, incorporating parallel corpora into the pre-training procedure improves the training efficiency, which potentially reduces the computational cost for building multilingual NLP applications.

## Acknowledgements

We appreciate the helpful discussions with Bo Zheng, Shaohan Huang, Shuming Ma, and Yue Cao. H.H. is the corresponding author. X.M. and H.H. are partly supported by National Key R&D Plan (No. 2018YFB1005100), National Natural Science Foundation of China (No. 61751201, 61602197 and 61772076), Natural Science Fund of Beijing (No. Z181100008918002), and the funds of Beijing Advanced Innovation Center for Language Resources (No. TYZ19005).

## References

- Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. 2019. A theoretical analysis of contrastive unsupervised representation learning. In *International Conference on Machine Learning*.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. [On the cross-lingual transferability of monolingual representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 4623–4637. Association for Computational Linguistics.
- Mikel Artetxe and Holger Schwenk. 2019. [Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond](#). *Transactions of the Association for Computational Linguistics*, 7(0):597–610.
- Philip Bachman, R Devon Hjelm, and William Buchwalter. 2019. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems*, pages 15509–15519.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. [A simple framework for contrastive learning of visual representations](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR.
- Zewen Chi, Li Dong, Furu Wei, Wenhui Wang, Xian-Ling Mao, and Heyan Huang. 2020. [Cross-lingual natural language generation via pre-training](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 7570–7577. AAAI Press.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020a. [Unsupervised cross-lingual representation learning at scale](#). In

- Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems*, pages 7057–7067. Curran Associates, Inc.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Alexis Conneau, Shijie Wu, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2020b. [Emerging cross-lingual structure in pretrained language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6022–6034, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Philipp Dufter and Hinrich Schütze. 2020. [Identifying elements essential for BERT’s multilinguality](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4423–4437, Online. Association for Computational Linguistics.
- Robert M. Fano. 1963. *Transmission of Information: A Statistical Theory of Communications*. M.I.T. Press.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. Language-agnostic BERT sentence embedding. *arXiv preprint arXiv:2007.01852*.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Olivier Henaff. 2020. [Data-efficient image recognition with contrastive predictive coding](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4182–4192. PMLR.
- R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. 2019. [Learning deep representations by mutual information estimation and maximization](#). In *International Conference on Learning Representations*.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421. PMLR.
- Haoyang Huang, Yaobo Liang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, and Ming Zhou. 2019. [Unicoder: A universal language encoder by pre-training with multiple cross-lingual tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 2485–2494, Hong Kong, China. Association for Computational Linguistics.
- Karthikeyan K, Zihan Wang, Stephen Mayhew, and Dan Roth. 2020. [Cross-lingual ability of multilingual bert: An empirical study](#). In *International Conference on Learning Representations*.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations*, San Diego, CA.
- Lingpeng Kong, Cyprien de Masson d’Autume, Lei Yu, Wang Ling, Zihang Dai, and Dani Yogatama. 2020. [A mutual information maximization perspective of language representation learning](#). In *International Conference on Learning Representations*.
- Anoop Kunchukuttan, Pratik Mehta, and Pushpak Bhattacharyya. 2018. [The IIT Bombay English-Hindi parallel corpus](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation*, Miyazaki, Japan. European Language Resources Association.
- Patrick Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2020. [MLQA: Evaluating cross-lingual extractive question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7315–7330, Online. Association for Computational Linguistics.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fengei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. [XGLUE: A new benchmark dataset for cross-lingual pre-training](#).

- understanding and generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, Online. Association for Computational Linguistics.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8(0):726–742.
- Xuanlong Nguyen, Martin J Wainwright, and Michael I Jordan. 2010. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Masoud Jalili Sabet, Philipp Dufter, and Hinrich Schütze. 2020. Simalign: High quality word alignments without parallel training data using static and contextualized embeddings. *arXiv preprint arXiv:2004.08728*.
- Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2019. WikiMatrix: Mining 135M parallel sentences in 1620 language pairs from wikipedia. *arXiv preprint arXiv:1907.05791*.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. [MASS: Masked sequence to sequence pre-training for language generation](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5926–5936. PMLR.
- Yonglong Tian, Dilip Krishnan, and Phillip Isola. 2019. Contrastive multiview coding. *ArXiv*, abs/1906.05849.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association.
- Naftali Tishby and Noga Zaslavsky. 2015. Deep learning and the information bottleneck principle. *2015 IEEE Information Theory Workshop (ITW)*, pages 1–5.
- T. Tsujishita. 1995. [On triple mutual information](#). *Advances in Applied Mathematics*, 16(3):269 – 274.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, pages 5998–6008. Curran Associates, Inc.
- 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*, pages 833–844, Hong Kong, China. Association for Computational Linguistics.
- Jian Yang, Shuming Ma, Dongdong Zhang, Shuangzhi Wu, Zhoujun Li, and Ming Zhou. 2020. Alternating language modeling for cross-lingual pre-training. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*.
- Michał Ziemski, Marcin Junczys-Dowmunt, and Bruno Pouliquen. 2016. The united nations parallel corpus v1. 0. In *LREC*, pages 3530–3534.

## A Pre-Training Data

We reconstruct CCNet<sup>2</sup> and follow (Conneau et al., 2020a) to reproduce the CC-100 corpus for monolingual texts. The resulting corpus contains 94 languages. Table 8 reports the language codes and data size in our work. Notice that several languages can share the same ISO language code, e.g., zh represents both Simplified Chinese and Traditional Chinese. Moreover, Table 9 shows the statistics of the parallel data.

Code	Size (GB)	Code	Size (GB)	Code	Size (GB)
af	0.2	hr	1.4	pa	0.8
am	0.4	hu	9.5	pl	28.6
ar	16.1	hy	0.7	ps	0.4
as	0.1	id	17.2	pt	39.4
az	0.8	is	0.5	ro	11.0
ba	0.2	it	47.2	ru	253.3
be	0.5	ja	86.8	sa	0.2
bg	7.0	ka	1.0	sd	0.2
bn	5.5	kk	0.6	si	1.3
ca	3.0	km	0.2	sk	13.6
ckb	0.6	kn	0.3	sl	6.2
cs	14.9	ko	40.0	sq	3.0
cy	0.4	ky	0.5	sr	7.2
da	6.9	la	0.3	sv	60.4
de	99.0	lo	0.2	sw	0.3
el	13.1	lt	2.3	ta	7.9
en	731.6	lv	1.3	te	2.3
eo	0.5	mk	0.6	tg	0.7
es	85.6	ml	1.3	th	33.0
et	1.4	mn	0.4	tl	1.2
eu	1.0	mr	0.5	tr	56.4
fa	19.0	ms	0.7	tt	0.6
fi	5.9	mt	0.2	ug	0.2
fr	89.9	my	0.4	uk	13.4
ga	0.2	ne	0.6	ur	3.0
gl	1.5	nl	25.9	uz	0.1
gu	0.3	nn	0.4	vi	74.5
he	4.4	no	5.5	yi	0.3
hi	5.0	or	0.3	zh	96.8

Table 8: The statistics of CCNet used corpus for pre-training.

ISO Code	Size (GB)	ISO Code	Size (GB)
en-ar	5.88	en-ru	7.72
en-bg	0.49	en-sw	0.06
en-de	4.21	en-th	0.47
en-el	2.28	en-tr	0.34
en-es	7.09	en-ur	0.39
en-fr	7.63	en-vi	0.86
en-hi	0.62	en-zh	4.02

Table 9: Parallel data used for pre-training.

## B Results of Training From Scratch

We conduct experiments under the setting of training from scratch. The Transformer size and hyperparameters follow BERT-base (Devlin et al., 2019). The parameters are randomly initialized from  $U[-0.02, 0.02]$ . We optimize the models with

<sup>2</sup>[https://github.com/facebookresearch/cc\\_net](https://github.com/facebookresearch/cc_net)

Model	XNLI	MLQA
Metrics	Acc.	F1 / EM
MMLM <sub>SCRATCH</sub>	69.40	55.02 / 37.90
INFOXML <sub>SCRATCH</sub>	<b>70.71</b>	<b>59.71 / 41.46</b>
–XLCO	70.64	57.70 / 40.21
–TLM	69.76	58.22 / 40.78
–MMLM	63.06	52.81 / 35.01

Table 10: Ablation results of the models pre-trained from scratch. Results are averaged over five runs.

Adam using a batch size of 256 for a total of 1M steps. The learning rate is scheduled with a linear decay with 10K warmup steps, where the peak learning rate is set as 0.0001. For cross-lingual contrast, we set the queue length as 16,384. We use a warmup of 200K steps for the key encoder and then enable cross-lingual contrast. We use an inverse square root scheduler to set the momentum coefficient, i.e.,  $m = \min(1 - t^{-0.51}, 0.9995)$ , where  $t$  is training step.

Table 10 shows the results of INFOXML<sub>SCRATCH</sub> and various ablations. INFOXML<sub>SCRATCH</sub> significantly outperforms MMLM<sub>SCRATCH</sub> on both XNLI and MLQA. We also evaluate the pre-training objectives of INFOXML, where we ablate XLCO, TLM, and MMLM, respectively. The findings agree with the results in Table 7.

## C Hyperparameters for Pre-Training

As shown in Table 11, we present the hyperparameters for pre-training INFOXML. We use the same vocabulary with XLM-R (Conneau et al., 2020a).

Hyperparameters	FROM SCRATCH	BASE	LARGE
Layers	12	12	24
Hidden size	768	768	1,024
FFN inner hidden size	3,072	3,072	4,096
Attention heads	12	12	16
Training steps	1M	150K	200K
Batch size	256	2,048	2,048
Adam $\epsilon$	1e-6	1e-6	1e-6
Adam $\beta$	(0.9, 0.999)	(0.9, 0.98)	(0.9, 0.98)
Learning rate	1e-4	2e-4	1e-4
Learning rate schedule	Linear	Linear	Linear
Warmup steps	10,000	10,000	10,000
Gradient clipping	1.0	1.0	1.0
Weight decay	0.01	0.01	0.01
Momentum coefficient	0.9995*	0.9999	0.999
Queue length	16,384	131,072	131,072
Universal layer	8	8	12

Table 11: Hyperparameters used for INFOXML pre-training. \*: the momentum coefficient uses an inverse square root scheduler  $m = \min(1 - t^{-0.51}, 0.9995)$ .

	XNLI	MLQA
Batch size	32	{16, 32}
Learning rate	{5e-6, 7e-6, 1e-5}	{2e-5, 3e-5, 5e-5}
LR schedule	Linear	Linear
Warmup	12,500 steps	10%
Weight decay	0	0
Epochs	10	{2, 3, 4}

Table 12: Hyperparameters used for fine-tuning BASE-size models on XNLI and MLQA.

	XNLI	MLQA
Batch size	32	32
Learning rate	{4e-6, 5e-6, 6e-6}	{2e-5, 3e-5, 5e-5}
LR schedule	Linear	Linear
Warmup	5,000 steps	10%
Weight decay	{0, 0.01}	0
Epochs	10	{2, 3, 4}

Table 13: Hyperparameters used for fine-tuning LARGE-size models on XNLI and MLQA.

## D Hyperparameters for Fine-Tuning

In Table 12 and Table 13, we present the hyperparameters for fine-tuning on XNLI and MLQA. For each task, the hyperparameters are searched on the joint validation set of all languages ( $\#M=1$ ). For XNLI, we evaluate the model every 5,000 steps, and select the model with the best accuracy score on the validation set. For MLQA, we directly use the final learned model. The final scores are averaged over five random seeds.