

VECO: Variable and Flexible Cross-lingual Pre-training for Language Understanding and Generation

Fuli Luo*, Wei Wang*, Jiahao Liu, Yijia Liu, Bin Bi, Songfang Huang, Fei Huang, Luo Si
Alibaba Group

{lf1259702, hebian.ww, glacier.ljh, yanshan.lyj}@alibaba-inc.com
{b.bi, songfang.hsf, f.huang, luo.si}@alibaba-inc.com

Abstract

Existing work in multilingual pretraining has demonstrated the potential of cross-lingual transferability by training a unified Transformer encoder for multiple languages. However, much of this work only relies on the shared vocabulary and bilingual contexts to encourage the correlation across languages, which is loose and implicit for aligning the contextual representations between languages. In this paper, we plug a cross-attention module into the Transformer encoder to explicitly build the interdependence between languages. It can effectively avoid the degeneration of predicting masked words only conditioned on the context in its own language. More importantly, when fine-tuning on downstream tasks, the cross-attention module can be plugged in or out on-demand, thus naturally benefiting a wider range of cross-lingual tasks, from language understanding to generation.

As a result, the proposed cross-lingual model delivers new state-of-the-art results on various cross-lingual understanding tasks of the XTREME benchmark, covering text classification, sequence labeling, question answering, and sentence retrieval. For cross-lingual generation tasks, it also outperforms all existing cross-lingual models and state-of-the-art Transformer variants on WMT14 English-to-German and English-to-French translation datasets, with gains of up to 1~2 BLEU.¹

1 Introduction

Cross-lingual pre-trained models like mBERT (Devlin et al., 2019), XLM (Lample and Conneau, 2019) and XLM-R (Conneau et al., 2019) that target providing contextualized representations for the inputs across languages, have shown large poten-

*Equal contribution.

¹Code and model are available at <https://github.com/alibaba/AliceMind/tree/main/VECO>




Figure 1: The attention scores of XLM and XLM-R with the input of a pair of parallel sentences: *Take a seat and have a rest* in English and its translated Chinese sentence. The darker line denotes a higher score. We can found that there are only a few attention patterns across English and Chinese subwords.

tial on a variety of cross-lingual understanding and generation tasks.

Behind the great success, two major factors play the role of aligning the contextual representations between languages: 1) build the shared vocabulary across languages through subword tokenization, which supports the simple extension of masked language modeling (MLM) from English corpus to multilingual corpus; 2) capture the alignment in parallel data via concatenating two sentences as input, called translation language modeling (TLM). However, both of these two mechanisms rely on the *self-attention* module (query=key/value) of the Transformer encoder to *implicitly* enhance the interdependence between languages, which may lead to few attention patterns across languages. Taking Figure 1 as an example, even though inputting a pair of parallel sentences, both models *only* attend to the English context to build the representation of English tokens, while ignoring the se-

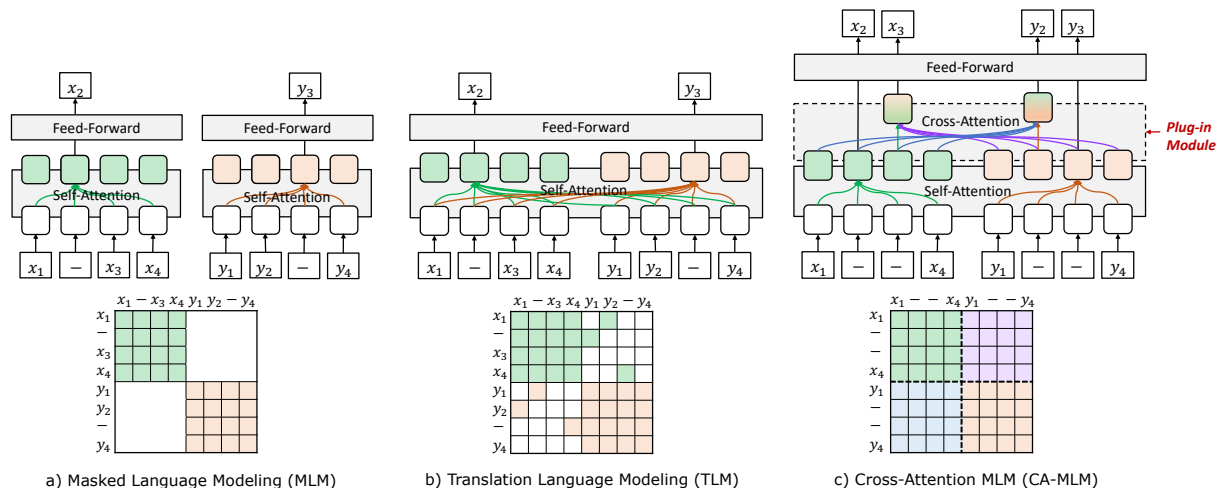


Figure 2: A schematic comparison of cross-lingual pre-training tasks and their attention matrices. When predicting the masked words of different languages: a) MLM can only attend to the context in its own language; b) TLM implicitly attend to a part of words across languages (as shown in Figure 1). However, c) the proposed CA-MLM can: (1) not only attend to the context in its own language to predict words x_2 and y_3 , (2) but also can firstly attend to its own context and then explicitly attend to all words across languages to predict words x_3 and y_2 via a plug-in cross-attention module.

matically related Chinese tokens. That is, the self-attention module captures little communication across languages, which is crucial for learning universal cross-lingual representations.

Based on the above observation, we propose to plug a *cross-attention* module (query!=key/value) into the Transformer encoder and design a cross-attention MLM task to *explicitly* capture the interdependence between languages. As illustrated in Figure 2 (c), the cross-attention module takes the representation of x as query and y as key/value (purple lines) to build the representations of x in the next layer, thus explicitly aligning the representations across languages (purple attention matrices). It can effectively avoid the degeneration of predicting masked words only conditioned on the context in its own language. Moreover, what distinguishes our work from pre-training an encoder-decoder model (Liu et al., 2020b) is that we also keep the good nature (i.e., bidirectional contextual modeling) of the original encoder by *unplugging* the cross-attention from the model to predicting the masked words (e.g., x_2 and y_3).

Furthermore, when fine-tuning on various downstream tasks, we can choose either *plug-in* or *plug-out* the cross-attention module on-demand, thus making it suitable for both cross-lingual language understanding (NLU) and generation tasks (NLG). For cross-lingual NLU tasks, if plugging the cross-attention module out, we can adopt the same fine-

tuning methods as an encoder-only model like XLM. However, we find that plugging the cross-attention module in fine-tuning can better utilize the bilingual context to boost the performance. For cross-lingual NLG like machine translation (MT), the cross attention is already jointly pre-trained with the whole network. Therefore, the parameters of the decoder do not need to be re-adjusted substantially in the following tuning process, thus fundamentally solving the main drawback of utilizing pre-trained encoders like XLM for initializing encoder-decoder models.

We call our approach VECO for “Variable and Flexible Cross-lingual Pre-training”. We validate VECO on a variety of representative cross-lingual understanding and generation benchmarks. Regrading cross-lingual understanding tasks, we conduct experiments on the XTREME benchmark consisting of 9 cross-lingual tasks, including text classification, sequence labeling, question answering, and sentence retrieval. VECO ranks first at the XTREME leaderboard² at the submission deadline. Regrading cross-lingual generation tasks, we validate VECO on the widely used WMT14 English-German and English-French machine translation benchmarks. VECO obtains 44.5 and 31.7 BLEU scores, consistently outperforming existing cross-lingual pre-training approaches and state-of-the-art Transformer variants by around 1~2 BLEU.

²<https://sites.research.google/xtreme>

2 Pre-training of VECO

2.1 Overview of VECO

VECO extends from a multi-layer Transformer encoder and plugs a cross-attention module in each layer. Given a pair of input (\mathbf{x}, \mathbf{y}) and its corrupted version $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ via randomly masking part of its tokens, the model builds two types of contextualized vector representation for each token:

- One suit of contextual representations \mathbf{H} , denoted as green blocks and yellow blocks in Figure 2 (c), are only build on self-attention module (i.e., unplugging the cross-attention module) in each layer.
- Another suit of contextual representations \mathbf{S} , denoted as mixed color blocks in Figure 2 (c), are build on both the self-attention and cross-attention modules³.

The model is trained to predict the masked tokens via two corresponding representations, conditioning on both its own context and paired context, respectively. Take predicting the masked words in sequence \mathbf{x} as an example, the training objective is the cross-entropy of the gold distribution and predicted distribution $P(\mathbf{x}|\hat{\mathbf{x}})$ and $P(\mathbf{x}|\hat{\mathbf{y}}, \hat{\mathbf{x}})$ computed via the above two suits of contextual representations. Thus, the training objective of cross-attention masked language modeling (CA-MLM) can be formulated as

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \mathbf{y}) = & \\ & - \log P(\mathbf{x}|\hat{\mathbf{x}}; \theta_s) - \log P(\mathbf{x}|\hat{\mathbf{y}}, \hat{\mathbf{x}}; \theta_s, \theta_c) \quad (1) \\ & - \log P(\mathbf{y}|\hat{\mathbf{y}}; \theta_s) - \log P(\mathbf{y}|\hat{\mathbf{x}}, \hat{\mathbf{y}}; \theta_s, \theta_c) \end{aligned}$$

where θ_s and θ_c are the parameters of self-attention and cross-attention modules.

2.2 Architecture

The backbone network of VECO is composed of a stack of N Transformer layers. Each layer has three modules: a required self-attention module, a plug-and-play cross-attention module, and a required feed-forward linear module. Both self-attention and cross-attention modules are based on the multi-head attention (Vaswani et al., 2017). An attention function can be described as mapping a query (\mathbf{Q}) and a set of key-value (\mathbf{K} - \mathbf{V}) pairs to an output.

³For simplicity of illustration, we only show the mixed representations \mathbf{S} of \mathbf{x}_3 and \mathbf{y}_2 in Figure 2 (c).

For the self-attention module, all the queries, keys and values are the same representations from the previous layer. Specifically, for the l -th Transformer layer, the output of a self-attention head \mathbf{A}_l^s is computed via:

$$\mathbf{Q} = \mathbf{H}^{l-1} \mathbf{W}_l^Q \quad (2)$$

$$\mathbf{K} = \mathbf{H}^{l-1} \mathbf{W}_l^K \quad (3)$$

$$\mathbf{V} = \mathbf{H}^{l-1} \mathbf{W}_l^V \quad (4)$$

$$\mathbf{A}_l^s = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}}\right)\mathbf{V} \quad (5)$$

where \mathbf{H}^{l-1} are the previous layer's outputs, $\mathbf{W}_l^Q, \mathbf{W}_l^K, \mathbf{W}_l^V$ are the parameter matrices of self-attention modules.

For the cross-attention module, the queries come from the previous layer, and the keys and values come from the last layer's representations of paired input. Specifically, for the l -th layer, the output of a cross-attention head \mathbf{A}_l^c is computed via:

$$\mathbf{Q} = \mathbf{S}^{l-1} \mathbf{U}_l^Q \quad (6)$$

$$\mathbf{K} = \mathbf{H}^L \mathbf{U}_l^K \quad (7)$$

$$\mathbf{V} = \mathbf{H}^L \mathbf{U}_l^V \quad (8)$$

$$\mathbf{A}_l^c = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}}\right)\mathbf{V} \quad (9)$$

where \mathbf{S}^{l-1} are the previous layer's outputs, $\mathbf{U}_l^Q, \mathbf{U}_l^K, \mathbf{U}_l^V$ are the parameter matrices of cross-attention modules.

Finally, the output \mathbf{H}^L of the last layer is used to recover the masked tokens of \mathbf{x} , conditioning on its own context.

$$P(\mathbf{x}|\hat{\mathbf{x}}) = \text{softmax}(f(\mathbf{H}_x^L)) \quad (10)$$

$$P(\mathbf{y}|\hat{\mathbf{y}}) = \text{softmax}(f(\mathbf{H}_y^L)) \quad (11)$$

where f is the feed-forward network that maps the output vectors into the dictionary. \mathbf{H}_x^L and \mathbf{H}_y^L are computed via Eq 2~5 when \mathbf{H}_x^0 and \mathbf{H}_y^0 are the word embeddings of \mathbf{x} and \mathbf{y} , respectively.

Meanwhile, \mathbf{S}^L , conditioning on the context of the paired sequence $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$, is used to predict the masked tokens of \mathbf{y} .

$$P(\mathbf{x}|\hat{\mathbf{y}}, \hat{\mathbf{x}}) = \text{softmax}(f(\mathbf{S}_x^L)) \quad (12)$$

$$P(\mathbf{y}|\hat{\mathbf{x}}, \hat{\mathbf{y}}) = \text{softmax}(f(\mathbf{S}_y^L)) \quad (13)$$

where \mathbf{S}_x^L and \mathbf{S}_y^L are computed via Eq 6~9 with the corresponding word embeddings and \mathbf{H}^L .

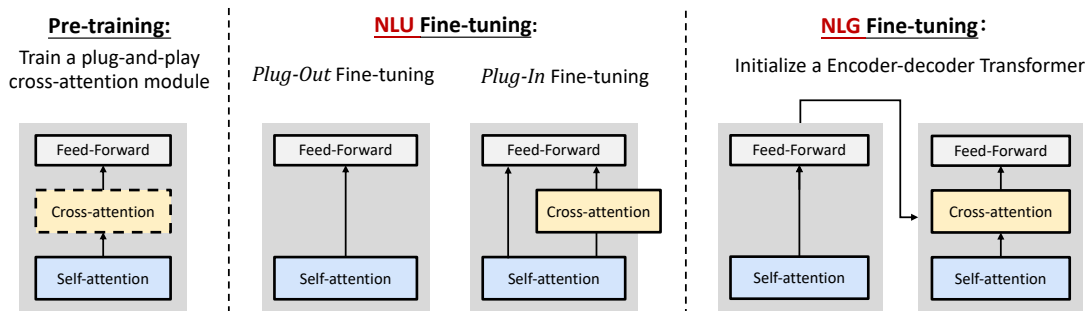


Figure 3: The overview of VECO. During pre-training, a plug-and-play cross-attention module is jointly pre-trained along with the self-attention module. When fine-tuning on natural language understanding (NLU) tasks, the cross-attention module can be either plug-in or plug-out on demand. When fine-tuning on natural language generation (NLG) tasks, VECO can initialize an encoder-decoder module (the mainstream backbone model of generation tasks) since all those necessary modules in the encoder and decoder are already pre-trained.

Note that when optimizing the objectives based on Eq 12 and Eq 13, we apply a stop-gradients operation (Chen and He, 2020) to \mathbf{H}^L (i.e., \mathbf{H}^L is treated as a constant in this term). This operation can largely speed up the training by avoiding the backpropagation on a $2L$ -layer network. Moreover, it even stabilizes the training of deep post-layernorm Transformer, which requires non-trivial efforts regarding carefully designing learning rate schedulers and cutting-edge optimizers (Liu et al., 2020a; Bachlechner et al., 2020).

3 Fine-tuning VECO for Downstream Cross-lingual Understanding and Generation Tasks

As Figure 3 illustrated, when fine-tuning on various downstream tasks, one advantage of VECO is its flexibility for initializing both the encoder-only Transformer for understanding tasks and encoder-decoder Transformer for generation tasks. Beyond it, we also explore a fine-tuning approach combined with the characteristics of VECO .

3.1 VECO for Cross-lingual Understanding

Due to the plug-and-play cross-attention module, we explore two fine-tuning approaches:

- *Plug-Out* fine-tuning is to unplug the cross-attention module from the pre-trained model. In other words, the architecture of the fine-tuned model is almost the same as mBERT or XLM. Specifically, the contextual representations from the last layer \mathbf{H}_x^L is used to predict the label of input x .
- *Plug-In* fine-tuning is to plug the cross-attention module into the fine-tuned model, if

the bilingual or automatically translated training data \mathbf{y} is available in the downstream task. Specifically, we concatenated the two representations $[\mathbf{H}_x^L : \mathbf{S}_x^L]$ to predict the label of x , $[\mathbf{H}_y^L : \mathbf{S}_y^L]$ to predict the label of y .⁴

3.2 VECO for Cross-lingual Generation

For pre-trained encoders like XLM, it is not a trivial problem to incorporate them into the sequence-to-sequence architecture – the mainstream backbone model of generation tasks (Zhu et al., 2020). One of the drawbacks or challenges could be that the encoder-to-decoder attention is not pre-trained. Therefore, the parameters of the decoder need to be re-adjusted along with the encoder in the following fine-tuning process (Ren et al., 2019).

However, under the framework of VECO , the cross-attention is jointly pre-trained along with the whole network, making it easy to provide full initialization for sequence-to-sequence models. Specifically, the self-attention module is used to initialize both the corresponding modules in the encoder and decoder for contextual modeling, while the cross-attention module is used to initialize the encoder-to-decoder attention. It’s okay whether you continue to tie the self-attention parameters during fine-tuning. Directly pre-training a sequence-to-sequence model like mBART (Liu et al., 2020b) could be another solution for NLG tasks, but we found mBART is not so effective in cross-lingual NLU tasks. We refer the reader to the Section 7 for detailed experiments and analysis.

⁴*Plug-In* fine-tuning is not suitable for the zero-shot setting (also called *cross-lingual transfer*) due to the lack of bilingual or translated pair (x, y)

Model	Architecture	#Parameters	Enc Layers	Dec Layers	#Languages	#Vocab	Training Data
mBERT (Devlin et al., 2019)	Encoder-only	110M	12	-	104	110k	Wikipedia
XLM (Lample and Conneau, 2019)	Encoder-only	570M	24	-	100	200k	Wikipedia
XLM-R (Conneau et al., 2019)	Encoder-only	550M	24	-	100	250k	CommonCrawl
mRASP (Lin et al., 2020)	Encoder-decoder	375M	6	6	32	64k	Translation
MMTE (Siddhant et al., 2020)	Encoder-decoder	375M	6	6	103	64k	Translation
mBART (Liu et al., 2020b)	Encoder-decoder	680M	12	12	25	250k	CommonCrawl
VECO	Flexible	662M		24*	50	250k	CommonCrawl + Translation

Table 1: Comparison of large cross-lingual models. * denotes VECO unifies the encoder and decoder.

4 Pre-training Setup

Model Configuration We pre-train a 24-layer model with 1024 embedding/hidden size and 4096 feed-forward size. We do not use language embeddings to allow our model to better deal with downstream tasks of unseen languages. We adopt the same 250K vocabulary that is also used by XLM-R (Conneau et al., 2019). Table 1 shows the other details of baselines and VECO.

Pre-Training Data We collect monolingual and bilingual corpus covering 50 languages. For monolingual training datasets, we reconstruct CommonCrawl Corpus used in XLM-R (Conneau et al., 2019). We extract 1.36TB data in 50 languages, which contains 6.5G sentences and 0.4G documents. We up/down-sample the monolingual text like XLM from each language with a smoothing parameter $\alpha = 0.5$. For bilingual data, we collect from the OPUS website⁵ like previous works (Lample and Conneau, 2019; Chi et al., 2020b). There are 6.4G parallel sentences, covering 879 language pairs across 50 languages. See more statistics of training data in Appendix A.

Optimization Settings For each iteration, we alternately sample a batch of adjacent segments from the monolingual corpus and a batch of parallel sentences from bilingual datasets to conduct a pair of masked input (\hat{x}, \hat{y}) . We adopt the translation language modeling (TLM) when the inputs are parallel bilingual sentences. Thus the overall training objective is the sum of TLM and the proposed CA-MLM objectives. During training, the model parameters except for cross-attention are initialized by XLM-R. We first freeze the parameters of XLM-R and only update the cross-attention parameters for faster convergence. Then, we jointly train the whole model. We pre-train our model with mixed-precision training using 64 Nvidia Tesla V100 32GB GPUs. Appendix A shows additional details.

⁵<http://opus.nlpl.eu/>

5 Experiments on Cross-lingual Understanding Tasks

5.1 Experimental Setup

Downstream Tasks We conduct cross-lingual NLU evaluations on XTREME (Hu et al., 2020), a representative massively multilingual benchmark that consists of 9 understanding tasks over 40 languages. XTREME tasks can be classified into four different categories: (1) sentence-pair classification: XNLI (Conneau et al., 2018), PAWS-X (Yang et al., 2019); (2) structured prediction: POS (Nivre et al., 2018), Wikiann NER (Pan et al., 2017); (3) question answering: XQuAD (Artetxe et al., 2020), MLQA (Lewis et al., 2020), TyDiQA (Clark et al., 2020); (4) sentence retrieval: BUCC 2018 (Zweigenbaum et al., 2017), Tatoeba (Artetxe and Schwenk, 2019). Tasks in the first three categories are provided: 1) golden training corpus in English, 2) translated training corpus in other languages, and 3) dev/test set in all languages. For sentence retrieval tasks, no training datasets are provided. We refer the reader to Hu et al. (2020) for additional details about the datasets.

Fine-tuning Setting Following previous works (Conneau et al., 2019; Hu et al., 2020), we consider two typical fine-tuning settings: (1) *Cross-lingual Transfer* which fine-tunes the pre-trained model using English golden data only and directly performs inference on the test data of different target languages; (2) *Translate-Train-All* fine-tunes a multilingual model on the concatenation of all data (golden training corpus in English and translated training corpus in other languages). Note that for two sequence-labeling tasks (POS, NER), the position of token labels in the translated text generally differs from that in the source text. Following FILTER (Fang et al., 2020), we use the model trained only on the English training dataset as a teacher, to label the translated text. To have a fair comparison with the strong baseline XLM-R (Conneau et al., 2019)

Datasets	XNLI	PAWS-X	POS	NER	XQuAD	MLQA	TyDiQA	BUCC	Tatoeba	
#Languages	15	7	33	40	11	7	9	5	33	
Metrics	Acc	Acc	F1	F1	F1/EM	F1/EM	F1/EM	F1	Acc	Avg.
<i>Cross-lingual Transfer: Fine-tune model on English training set and test on all languages</i>										
MMTE [†]	67.4	81.3	73.5	58.3	64.4/46.2	60.3/41.4	58.1/43.8	59.8	37.9	59.5
mBERT [†]	65.4	81.9	70.3	62.2	64.5/49.4	61.4/44.2	59.7/43.0	56.7	38.7	59.6
XLM [†]	69.1	80.9	70.1	61.2	59.8/44.3	48.5/32.6	43.6/29.1	56.8	32.6	55.5
XLM-R [†]	79.2	86.4	72.6	65.4	76.6/60.8	71.6/53.2	65.1/45.0	66.0	57.3	68.1
VECO _{out}	79.9	88.7	75.1	65.7	77.3/61.8	71.7/53.2	67.6/49.1	85.0	75.1	73.1
<i>Translate-Train-All: Fine-tune model on English training data and translated data of the target language</i>										
XLM-R [‡]	82.6	90.4	-	-	80.2/65.9	72.8/54.3	66.5/47.7	-	-	-
XLM-R*	82.8	90.2	72.6	65.4	80.0/65.8	73.0/54.3	74.5/58.3	80.2	75.2	74.4
FILTER	83.9	91.4	76.2	67.7	82.4/68.0	76.2/57.7	68.3/50.9	84.5	84.5	77.0
VECO _{out}	83.0	91.1	75.1	65.7	79.9/66.3	73.1/54.9	75.0/58.9	89.3	86.9	77.2
VECO _{in}	84.3	92.8	79.8	71.0	83.9/70.9	77.5/59.3	79.4/63.7	92.6	91.1	81.0

Table 2: XTREME results on each dataset (as of ACL submission deadline). Averaged results on the four categories can be found at leaderboard: <https://sites.research.google/xtreme>. “†” and “‡” indicates results from Hu et al. (2020) and Fang et al. (2020), respectively. “*” indicates the results obtained by our implementation. The detailed results for each language are in Appendix D.

under the translate-train-all setting, we also show the results of XLM-R using the same fine-tuning hyperparameters as VECO.

5.2 Experimental Results

The detailed test results of nine tasks on the XTREME benchmark are shown in Table 2. It demonstrates that the proposed VECO outperforms previous cross-lingual models on all datasets. Compared to XLM-R, it averages 5.0 and 6.6 points higher under the cross-lingual transfer and translation-train-all settings, respectively.

In the cross-lingual transfer setting, VECO delivers a large improvement compared to XLM-R, especially on zero-shot sentence retrieval tasks (BUCC, Tatoeba). This phenomenon reflects that our model can better build the interdependence between languages. Thus it can better mine parallel sentences in a multilingual corpus.

Under the translation-train-all setting, it can be observed that VECO with Plug-In fine-tuning (VECO_{in}) is better than Plug-Out fine-tuning (VECO_{out}). We conclude the reasons as two-fold. On the input side, the Plug-Out fine-tuning individually takes multilingual instances as input, while the Plug-In fine-tuning considers the bilingual instances⁶ at each run. On the model side, the Plug-In fine-tuning can encourage correspondence across language via the cross-attention module. Note that the Plug-In fine-tuning method also outperforms FILTER (Fang et al., 2020), an enhanced cross-lingual fine-tuning method that also takes the

bilingual instance as the input of XLM-R. It further demonstrates the effectiveness of VECO and its specialized fine-tuning method.

We conclude the reasons for the above performance improvement as two-fold: 1) the introduction of bilingual data during pre-training, which is a direct way to enhance the cross-lingual ability of the model; 2) Stronger ability to enhance the interdependence and fusion among languages via the proposed CA-MLM pre-training tasks. To analyze which plays a leading role, we conduct a set of more fair experiments in Section 7.

6 Experiments on Cross-lingual Generation Tasks

6.1 Experimental Setup

Datasets We choose the machine translation (MT) task, a typical cross-lingual generation scenario. In order to illustrate the generality of our approach and have a fair comparison with the most recent state-of-the-art Transformer work (Liu et al., 2020a), we choose two most widely used datasets: WMT14 English→German (En-De) and English→French (En-Fr) translation. WMT14 En-De is a medium-resource dataset that provides 4.5M pairs for training and validation. We adopt standard newstest2014 as the test set. WMT14 En-Fr is a high-resource dataset that contains 36M pairs of parallel sentences. We use newstest2012+newstest2013 for validation and newstest2016 for test. We measure case-insensitive tokenized BLEU with `multi-bleu.perl` and de-

⁶English instance with its translated one.

Model	WMT14 En-Fr		WMT14 En-De	
	BLEU	SacreBLEU	BLEU	SacreBLEU
<i>Randomly Initialize</i>				
Baseline	42.9	40.4	28.7	27.8
Liu et al. (2020a)	43.8	41.8	30.1	29.5
<i>Randomly Initialize + More Bilingual Data*</i>				
Baseline*	-	-	30.6	29.5
<i>Cross-lingual Model Initialize</i>				
mBART	43.2	41.0	30.0	29.1
mRASP	44.3	41.7	30.3	-
XLM-R	43.8	41.2	30.9	29.9
VECO	44.5	42.0	31.7	30.6




Table 3: (left) Results on machine translation. (right) Learning curves of different initialization methods.

tokenized SacreBLEU⁷ to avoid the influence of different tokenization and normalization between models (Post, 2018).

Fine-tuning Setting We fine-tune our model using fairseq⁸ toolkit and adopt comparable training settings with baselines. We run WMT 14 En-De and En-Fr MT experiments on 16 and 32 V100 GPUs, respectively. The batch size is 64k for En-De and 256k for En-Fr. The total training updates are set to 100k. The learning rate is $1e-4/2e-4$, with linear warm-up over the first 16k steps and linear decay. We average the last 10 checkpoints and use beam search with a beam size of 5.

Baselines We consider two types of Transformer baselines: randomly initialized and cross-lingual models initialized. For random initialization, we reproduce a Transformer baseline that adopts the same architecture and fine-tuning hyperparameters as VECO but with random initialization. Besides, we compare to the state-of-the-art Deep Transformer (Liu et al., 2020a). For cross-lingual encoder-decoder models, we include mBART (Liu et al., 2020b) and mRASP (Lin et al., 2020), which show impressive results on MT. Note that since we tied the self-attention weights of each encoder layer with each decoder layer, the whole parameters of mBART and VECO are comparable. We also conduct the WMT experiments for XLM-R, following the totally same fine-tuning settings as VECO, but leaving the encoder-to-decoder attention un-initialized.

⁷Hash: BLEU+case.mixed+lang.en-{de,fr}+numrefs.1+smooth.exp+test.wmt14/full+tok.13a+version.1.4.9

⁸<https://github.com/pytorch/fairseq>

6.2 Experimental Results

Table 3 (left) shows the results on the machine translation. We can observe that VECO can largely outperform the randomly initialized same-sized Transformer baseline by 2.3 BLEU points. Moreover, it even beats the (randomly initialized) state-of-the-art Deep-Transformer (Liu et al., 2020a), which is three times deep as VECO. Among the cross-lingual models, VECO can consistently outperform the best models, averaged on two datasets, by 0.8 BLEU points.

Table 3 (right) displays the BLEU scores of same-sized models during training. We find that VECO initialized model can get a surprising more than 28 SacreBLEU score just after 10 epochs, which is better than the final score of the randomly initialized model at 35 epochs. It reveals that VECO can provide a fairly good initialization for the machine translation model, which can converge quickly and further boost the results.

One might suspect that the main reason for the performance improvement is leveraging parallel corpus during pre-training. To figure it out, we conduct a more comparable experiment. We first train an out-of-domain Transformer model using the whole En-De parallel data ($\sim 68M$) used in VECO pre-training, and then continue to train the model on the in-domain WMT14 En-De training dataset. Results are shown in Table 3 (left) marked with *. Under this set of a totally fair comparison, VECO still maintains a lead of 1.1 BLEU score. This directly confirms that the improvement in MT is not only due to the use of bilingual data. More importantly, CA-MLM ensures *better* use of bilingual and large-scale unlabeled multilingual corpus.

Method	#Layers	WMT14 En-De	
		BLEU	SacreBLEU
<i>Randomly Initialize</i>	3	28.5	27.6
	6	28.6	27.7
<i>VECO Initialize</i>	First-3	30.8	29.8
	Last-3	31.2	30.3
	First-6	31.1	30.1
	Last-6	31.5	30.5
	Full-24	31.7	30.6

Table 4: Results of utilizing VECO to initialize deep encoder and shallow decoder (3/6-layer) Transformers.

6.3 Potential of Initializing Shallow Decoder

Online translation applications usually have a restriction of inference time. The most direct way is to reduce the decoder layers since previous MT works (Liu et al., 2020a) have shown that deeper encoders are more worthwhile than deeper decoders. Based on this, we also explore the potential of the VECO to initialize deep encoder and shallow decoder Transformers, which is a blank in the cross-lingual pre-training works.

Table 4 contrasts two ways of initializing a Transformer with n decoder layers ($n < 24$) via selecting: (1) the first n layers; (2) the last n layers from a 24-layer pre-trained VECO model. We consider $n = \{3, 6\}$ to conduct experiments. We find that selecting the last n layers exhibits better performance than selecting the first n layers. It reveals that the last several layers play a more important role in making predictions over the whole vocabulary. Moreover, we can find that there is 0.2~0.3 BLEU gain when increasing the decoder layers from 3 to 6. However, we observe that only marginal improvement can be gained when further increasing the decoder layers from 6 to 24, which is also in line with the findings in Liu et al. (2020a). Regardless of the initialization method, the VECO initialized model can gain consistent 1~2 BLEU improvement over the randomly initialized model.

7 Analysis and Ablation Study

We perform an ablation study to investigate where the improvement in cross-lingual NLU and NLG tasks mainly comes from. Specifically, there are three main aspects we have studied:

1. How much performance improvement comes from the parallel translation corpus used in pre-training?
2. How effective of the CA-MLM pre-training

Data	Models	Tasks	XNLI	IWSLT
Mono.	XLM	MLM	59.8	33.7
	mBART	MLM	57.3	32.9
	VECO	CA-MLM	60.6	34.0
Bili.	XLM	MLM+TLM	64.5	33.9
	mBART	MLM+TLM	60.8	34.5
	VECO	CA-MLM+TLM	67.7	36.0

Table 5: Ablation study of small-sized models on XNLI and IWSLT14 De-En translation dataset.

task, especially compared to the MLM and TLM pre-training tasks?

3. How about pre-training a sequence-to-sequence model like mBART for NLU and NLG tasks?

To figure out these questions, we train XLM, mBART and VECO model from scratch using the same datasets and parameter settings (see Appendix A for more details). All of them is pre-trained via MLM and TLM tasks. Note that the MLM task generally refers to predict the masked words of source language, while the TLM task generally refers to predict the words of the target language. Specifically for mBART that is under the framework of encoder-decoder, the input of encoder is masked sequence \hat{x} , and the target of decoder is the masked words of source input x (for MLM task), or the parallel sentence y (for TLM task).

Table 5 shows the results of two representative datasets of cross-lingual NLU and NLG. We can observe that, when using monolingual corpus only, VECO can outperform XLM by 0.8 points on the XNLI dataset and 0.3 BLEU scores on the IWSLT14 De-En translation dataset. It suggests that the CA-MLM can still benefit from adjacent sentences in monolingual corpus⁹, to be equipped with a stronger ability of contextual modeling. Moreover, when pre-training both on the monolingual and bilingual corpus, VECO can even achieve a larger improvement compared to XLM, with 3.2 and 2.1 points improvement on two datasets, respectively. It reveals that CA-MLM objective of VECO can better utilize the bilingual corpus, compared to only optimized by TLM and MLM of XLM.

Moreover, we find that pre-training a sequence-to-sequence model like mBART (Liu et al., 2020b)

⁹As noted in Section 4, we take two adjacent sentences in the monolingual corpus as (x, y) .

performs worst on NLU tasks like XNLI¹⁰, almost 6 points worse than VECO and near 2 points worse than XLM. One possible explanation could be that the unidirectional language modeling in the decoder might be sub-optimal for NLU tasks. And even on the machine translation task, mBART still performs worse than VECO when pre-training on the same bilingual datasets. We conclude that it is because that VECO can do better in the contextual modeling of source input x via a explicit masked language modeling objective in Eq 10 applied to x_2 in Figure 2 (c).

8 Related Work

mBERT (Devlin et al., 2019) is a key step towards building a unified contextual language representation over multiple languages. It simply shares all languages’ vocabulary and trains a bidirectional Transformer encoder, achieving promising results in various cross-lingual NLU tasks. There have been several extensions that follow the same encoder-only backbone as mBERT. The main difference is the introduction of more training corpus (e.g., bilingual data) and pre-training tasks. XLM (Lample and Conneau, 2019) utilizes both monolingual and bilingual corpus to perform the masked language modeling. XLM-R (Conneau et al., 2019) extends to be built on RoBERTa (Liu et al., 2019) using larger monolingual training data. Other works (Huang et al., 2019; Yang et al., 2020; Chi et al., 2020b) propose new pre-training tasks to utilize the bilingual data better. However, there are two main drawbacks of these works. First, they mainly rely on the self-attention module in the Transformer encoder to implicitly build the interdependence between languages, leading to few attention patterns across languages due to the “lazy” network. Second, even though they show impressive performance improvement on cross-lingual understanding tasks like XNLI, only marginal improvement has been gained on cross-lingual generation tasks like machine translation, especially on high-resource languages.

A feasible solution for cross-language generation is to pre-train a denoising auto-encoder like mBART (Liu et al., 2020b). It extends BART (Lewis et al., 2019) to the multilingual setting, demonstrating significant gains in low/medium-resource machine translation, but

¹⁰We follow BART (Lewis et al., 2019) by utilizing the final representation from the decoder for classification tasks.

with a decrease in high resource languages. Unlike mBART, Chi et al. (2020a) first trains an encoder via MLM and then frozen the encoder to train the decoder only via two generative tasks. A similar approach is also proposed in Liang et al. (2020) and Lin et al. (2020), with the main difference in the joint training of encoder-decoder with code-switch tricks. However, all these cross-lingual models emphasize training a dedicated model for NLG. Thus they may hurt the NLU capabilities of the model. The ablation study in Section 7 also validates that it is sub-optimal to train an encoder-encoder network for NLU tasks.

This paper endeavors to build a unified cross-lingual model for NLU and NLG tasks via a plug-and-play cross-attention module. More importantly, the cross-attention module plays a role in the explicit alignment of encoded representations of different languages, thus largely contributing to building a unified cross-lingual model.

9 Conclusion

We present VECO, a variable and flexible cross-lingual pre-training model, targets at explicitly capturing the interdependence between languages via a plug-and-play cross-attention module. Based on the flexible characteristics, VECO can initialize both NLU preferred encoder-only and NLG specialized encoder-decoder Transformer. Moreover, we also introduce a Plug-In fine-tuning approach to encourage the fusion between languages, combining the feature of VECO and cross-language downstream tasks.

Taken together, VECO achieves consistent improvements on various language understanding and generation tasks, broadening the way of thinking about pre-trained backbone architecture and fine-tuning methods under the cross-lingual scenario.

References

- 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](#). *Trans. Assoc. Comput. Linguistics*, 7:597–610.

- Thomas Bachlechner, Bodhisattwa Prasad Majumder, Huanru Henry Mao, Garrison W Cottrell, and Julian McAuley. 2020. [Rezero is all you need: Fast convergence at large depth.](#) *arXiv preprint arXiv:2003.04887*.
- Xinlei Chen and Kaiming He. 2020. [Exploring simple siamese representation learning.](#) *CoRR*, abs/2011.10566.
- Zewen Chi, Li Dong, Furu Wei, Wenhui Wang, Xian-Ling Mao, and Heyan Huang. 2020a. [Cross-lingual natural language generation via pre-training.](#) In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2020b. [InfoXLM: An information-theoretic framework for cross-lingual language model pre-training.](#) *arXiv preprint arXiv:2007.07834*.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. [TyDi QA: A Benchmark for Information-Seeking Question Answering in Typologically Diverse Languages.](#) In *Transactions of the Association of Computational Linguistics*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Unsupervised cross-lingual representation learning at scale.](#) *arXiv preprint arXiv:1911.02116*.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations.](#) In *Proceedings of EMNLP 2018*, pages 2475–2485.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding.](#) In *NAACL-HLT*.
- Yuwei Fang, Shuohang Wang, Zhe Gan, Siqi Sun, and Jingjing Liu. 2020. [FILTER: An enhanced fusion method for cross-lingual language understanding.](#) *arXiv preprint arXiv:2009.05166*.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization.](#) *arXiv preprint arXiv:2003.11080*.
- Haoyang Huang, Yaobo Liang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, and Ming Zhou. 2019. [Unicoder: A universal language encoder by pre-training with multiple cross-lingual tasks.](#) *arXiv preprint arXiv:1909.00964*.
- Taku Kudo and John Richardson. 2018. [Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing.](#) *arXiv preprint arXiv:1808.06226*.
- Guillaume Lample and Alexis Conneau. 2019. [Cross-lingual language model pretraining.](#) *arXiv preprint arXiv:1901.07291*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension.](#) *arXiv preprint arXiv:1910.13461*.
- Patrick Lewis, Barlas Oğuz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2020. [MLQA: Evaluating Cross-lingual Extractive Question Answering.](#) In *Proceedings of ACL 2020*.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, et al. 2020. [XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation.](#) *arXiv preprint arXiv:2004.01401*.
- Zehui Lin, Xiao Pan, Mingxuan Wang, Xipeng Qiu, Jiangtao Feng, Hao Zhou, and Lei Li. 2020. [Pre-training multilingual neural machine translation by leveraging alignment information.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2649–2663. Association for Computational Linguistics.
- Xiaodong Liu, Kevin Duh, Liyuan Liu, and Jianfeng Gao. 2020a. [Very deep transformers for neural machine translation.](#) *arXiv preprint arXiv:2008.07772*.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020b. [Multilingual denoising pre-training for neural machine translation.](#) *arXiv preprint arXiv:2001.08210*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A robustly optimized BERT pretraining approach.](#) *arXiv preprint arXiv:1907.11692*.
- Joakim Nivre, Mitchell Abrams, Željko Agić, Lars Ahrenberg, Lene Antonsen, Maria Jesus Aranzabe, Gashaw Arutie, Masayuki Asahara, Luma Ateyah, Mohammed Attia, et al. 2018. [Universal dependencies 2.2.](#)
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. [Cross-lingual name tagging and linking for 282 languages.](#) In *Proceedings of ACL 2017*, pages 1946–1958.

Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation*.

Shuo Ren, Yu Wu, Shujie Liu, Ming Zhou, and Shuai Ma. 2019. [Explicit cross-lingual pre-training for unsupervised machine translation](#). volume abs/1909.00180.

Aditya Siddhant, Melvin Johnson, Henry Tsai, Naveen Ari, Jason Riesa, Ankur Bapna, Orhan Firat, and Karthik Raman. 2020. [Evaluating the cross-lingual effectiveness of massively multilingual neural machine translation](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pages 8854–8861.

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*.

Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzman, Armand Joulin, and Edouard Grave. 2019. [Ccnnet: Extracting high quality monolingual datasets from web crawl data](#). *arXiv preprint arXiv:1911.00359*.

Jian Yang, Shuming Ma, Dongdong Zhang, Shuangzhi Wu, Zhoujun Li, and Ming Zhou. 2020. [Alternating language modeling for cross-lingual pre-training](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Yinfei Yang, Yuan Zhang, Chris Tar, and Jason Baldridge. 2019. [PAWS-X: A cross-lingual adversarial dataset for paraphrase identification](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3685–3690. Association for Computational Linguistics.

Jinhua Zhu, Yingce Xia, Lijun Wu, Di He, Tao Qin, Wengang Zhou, Houqiang Li, and Tie-Yan Liu. 2020. [Incorporating BERT into neural machine translation](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Pierre Zweigenbaum, Serge Sharoff, and Reinhard Rapp. 2017. [Overview of the second BUCC shared task: Spotting parallel sentences in comparable corpora](#). In *Proceedings of the 10th Workshop on Building and Using Comparable Corpora, BUCC@ACL 2017, Vancouver, Canada, August 3, 2017*, pages 60–67. Association for Computational Linguistics.

A Pre-Training Details

For monolingual data, following XLM-R (Conneau et al., 2019), we build a clean CommonCrawl Corpus using an open-source tool CCNet (Wenzek

et al., 2019). There are 1.36TB monolingual data in 50 languages before up/down-sampling. Table 6 reports the language codes and statistics of pre-training data. We collect bilingual corpus in 50 languages from the OPUS website¹¹, including MultiUN, UNPC, Bombay, EU-bookshop, OpenSubtitles2018, Tanzil, GlobalVoices, ParaCrawl, Multi-ParaCrawl, DGT, Tilde, Europarl, Wikipedia, ECB, TED2013, News-Commentary, Ubuntu, Books, UN, infopankki-v1, EUconst, and Bianet. In total, there are 1TB bilingual training data before pre-processing, covering 879 language pairs. Table 7 lists the statistics for each language pair. We then apply subword tokenization directly on raw text data using Sentence Piece Model (Kudo and Richardson, 2018) without any additional preprocessing.

We use the whole corpus to train VECO and a subset ($\sim 1/4$) that contains 33 languages to train small-sized XLM, mBART and VECO. The full set of pre-training hyperparameters for small-sized and large-sized VECO (default) are listed in Table 8.

B More details about Illustrated Attention

The models illustrated with attention patterns in Figure 1 of main paper (not appendix), are the base-sized XLM¹² and XLM-R¹³. We show the attention scores averaged on all heads in the middle layer.

C Fine-Tuning Details on XTERME

We select the model with the best average result over all the languages on the dev sets, by searching the learning rate over [5e-6, 8e-6, 1e-5, 2e-5, 3e-5] for the *Cross-lingual Transfer* setting and [5e-6, 6e-6, 7e-6, 8e-6, 9e-6] for *Translate-Train-All* setting, training epoch over [3, 5, 10], and batch size over [16, 32, 64].

D Detailed Results on XTREME

The detailed results of each XTREME task under the cross-lingual transfer and translate-train-all settings on all languages are listed in the following tables.

¹¹<http://opus.nlpl.eu/>

¹²<https://huggingface.co/xlm-mlm-tlm-xnli15-1024>

¹³<https://huggingface.co/xlm-roberta-base>

Language	#Document(M)	#Sentence(M)	Size(GB)
af	0.023	0.522	0.107
ar	2.823	42.659	11.786
bg	0.919	14.743	5.217
bn	0.750	9.217	4.264
cs	3.980	55.754	9.668
de	21.410	310.942	66.333
el	1.740	24.334	9.737
en	130.087	2,215.534	479.099
es	17.569	267.764	58.774
et	0.347	5.252	0.877
eu	0.342	5.216	0.613
fr	15.819	267.888	58.023
fa	2.506	43.570	13.831
fi	1.530	23.790	3.940
fy	0.027	0.537	0.054
gu	0.039	0.519	0.228
gd	0.009	0.126	0.020
he	0.755	12.338	3.073
hi	0.536	7.303	3.762
hu	1.816	29.962	6.421
id	3.417	60.908	11.528
it	9.336	133.006	30.854
ja	27.967	588.926	71.785
ju	0.002	0.138	0.030
ka	0.141	1.756	0.766
kk	0.061	1.545	0.448
ko	11.609	227.396	27.837
lt	0.552	7.996	1.480
lv	0.281	4.159	0.798
ms	0.334	3.762	0.455
ml	0.162	2.615	1.025
my	0.045	0.893	0.306
mr	0.059	0.708	0.365
pl	6.642	93.760	19.082
pt	8.623	128.107	25.612
ne	0.080	0.829	0.429
nl	6.513	85.997	16.648
ru	35.887	580.291	203.105
ro	1.944	31.929	7.056
si	0.132	2.927	0.902
sw	0.057	0.945	0.179
ta	0.876	20.376	6.422
te	0.288	4.995	1.721
tr	18.547	291.081	40.321
th	6.278	117.826	27.941
tl	0.166	5.611	0.679
vi	12.183	234.071	37.919
ur	0.460	7.509	2.003
yo	0.0002	0.003	0.0005
zh	27.067	497.408	87.005
Total	382.735	6,475.444	1,360.526

Table 6: The statistics of monolingual pre-training corpus.

Pre-training Hyperparameters	Large	Small
Number of layers	24	6
Hidden Size	1024	768
FFN inner hidden size	4096	3072
Attention heads	16	12
Attention head size	64	64
Embedding Size	1024	768
Mask percent (monolingual/ bilingual)	15%/25%	15%/25%
Learning Rate Decay	Linear	Linear
Warmup steps	12k	12k
Learning Rate	2e-4	3e-4
Adam ϵ	1e-6	1e-6
Adam β_1	0.9	0.9
Adam β_2	0.98	0.999
Attention Dropout	0.1	0.1
Dropout	0.1	0.1
Weight Decay	0.01	0.01
Max Sequence Length (monolingual/bilingual)	512/128	512/128
Batch Size (monolingual/bilingual)	1024/4096	1024/4096
Train Steps	240k	240k
Total Parameters	662M	247M

Table 8: The pre-training hyperparameters.

Model	en	ar	bg	de	el	es	fr	hi	ru	sw	th	tr	ur	vi	zh	Avg.
Cross-lingual Transfer																
XLm-R	88.7	77.2	83.0	82.5	80.8	83.7	82.2	75.6	79.1	71.2	77.4	78.0	71.7	79.3	78.2	79.2
VECO _{out}	88.2	79.2	83.1	82.9	81.2	84.2	82.8	76.2	80.3	74.3	77.0	78.4	71.3	80.4	79.1	79.9
Translate-Train-All																
XLm-R	88.6	82.2	85.2	84.5	84.5	85.7	84.2	80.8	81.8	77.0	80.2	82.1	77.7	82.6	82.7	82.6
VECO _{out}	88.9	82.4	86.0	84.7	85.3	86.2	85.8	80.1	83.0	77.2	80.9	82.8	75.3	83.1	83.0	83.0
VECO _{in}	89.3	83.7	87.0	85.9	85.8	87.3	86.7	81.8	83.6	79.9	82.5	84.3	77.7	84.4	84.0	84.3

Table 9: XNLI accuracy scores for each language.

Model	en	de	es	fr	ja	ko	zh	Avg.	Model	de	fr	ru	zh	Avg.
Cross-lingual Transfer									Cross-lingual Transfer					
XLm-R	94.7	89.7	90.1	90.4	78.7	79.0	82.3	86.4	XLm-R	67.5	66.5	73.5	56.7	66.0
VECO _{out}	96.2	91.3	91.4	92.0	81.8	82.9	85.1	88.7	VECO _{out}	89.6	84.6	87.4	78.5	85.0
Translate-Train-All									Translate-Train-All					
VECO _{out}	96.4	93.0	93.0	93.5	87.2	86.8	87.9	91.1	VECO _{out}	93.0	88.7	89.9	85.7	89.3
VECO _{in}	96.5	94.4	94.3	94.0	89.0	90.3	91.0	92.8	VECO _{in}	95.4	91.9	93.1	89.9	92.6

Table 10: PAWS-X accuracy scores.

Table 11: BUCC F1 results.

Model	af	ar	bg	de	el	en	es	et	eu	fa	fi	fr	he	hi	hu	id	it
Cross-lingual Transfer																	
XLm-R	89.8	67.5	88.1	88.5	86.3	96.1	88.3	86.5	72.5	70.6	85.8	87.2	68.3	76.4	82.6	72.4	89.4
VECO _{out}	88.3	67.4	87.4	88.5	86.7	95.9	89.0	87.8	75.1	70.9	86.2	88.9	67.5	76.2	82.9	72.9	89.9
Translate-Train-All																	
VECO _{in}	92.5	73.7	93.4	91.8	90.4	95.2	91.3	90.6	79.1	79.8	89.5	91.4	79.1	80.6	88.4	74.8	91.8
	ja	kk	ko	mr	nl	pt	ru	ta	te	th	tl	tr	ur	vi	yo	zh	Avg.
Cross-lingual Transfer																	
XLm-R	15.9	78.1	53.9	80.8	89.5	87.6	89.5	65.2	86.6	47.2	92.2	76.3	70.3	56.8	24.6	25.7	73.8
VECO _{out}	31.4	79.3	53.1	84.3	89.8	88.3	90.2	64.3	85.8	48.0	93.7	77.2	69.2	58.1	26.2	39.4	75.1
Translate-Train-All																	
VECO _{in}	45.1	78.0	63.7	84.5	92.7	90.1	92.6	72.6	88.5	55.2	88.8	76.8	75.0	70.5	24.3	63.0	79.8

Table 12: POS results (Accuracy) for each language.

Model	en	af	ar	bg	bn	de	el	es	et	eu	fa	fi	fr	he	hi	hu	id	it	ja	jv
Cross-lingual Transfer																				
XML-R	84.7	78.9	53.0	81.4	78.8	78.8	79.5	79.6	79.1	60.9	61.9	79.2	80.5	56.8	73.0	79.8	53.0	81.3	23.2	62.5
VECO _{out}	83.8	77.5	48.2	83.9	77.2	79.4	79.3	75.4	80.4	68.3	68.2	80.6	80.1	55.0	71.0	80.9	52.9	81.7	19.4	63.2
Translate-Train-All																				
VECO _{in}	80.7	82.5	66.4	84.1	78.4	82.2	82.4	79.7	84.7	78.2	68.8	84.9	79.1	69.7	76.6	85.1	77.3	83.8	21.3	70.3
	ka	kk	ko	ml	mr	ms	my	nl	pt	ru	sw	ta	te	th	tl	tr	ur	vi	yo	zh
Cross-lingual Transfer																				
XML-R	71.6	56.2	60.0	67.8	68.1	57.1	54.3	84.0	81.9	69.1	70.5	59.5	55.8	1.3	73.2	76.1	56.4	79.4	33.6	33.1
VECO _{out}	67.1	51.2	59.9	63.4	65.0	70.0	56.1	83.4	83.1	71.3	70.5	60.5	56.2	1.4	71.3	80.4	69.3	76.0	37.4	29.1
Translate-Train-All																				
VECO _{in}	77.0	67.2	71.0	73.3	74.1	71.8	63.8	85.5	80.8	72.8	77.0	69.1	67.5	2.6	74.0	85.2	71.5	76.4	32.8	31.0

Table 13: NER results (F1) for each language.

Model	en	ar	de	el	es	hi	ru	th	tr	vi	zh	Avg.
Cross-lingual Transfer												
XML-R	86.5 / 75.7	68.6 / 49.0	80.4 / 63.4	79.8 / 61.7	82.0 / 63.9	76.7 / 59.7	80.1 / 64.3	74.2 / 62.8	75.9 / 59.3	79.1 / 59.0	59.3 / 50.0	76.6 / 60.8
VECO _{out}	87.6 / 76.5	73.6 / 56.1	79.8 / 62.2	79.6 / 61.6	81.2 / 61.6	74.7 / 57.6	78.7 / 62.1	72.8 / 60.6	75.1 / 58.3	79.0 / 59.8	69.2 / 59.2	77.3 / 61.8
Translate-Train-All												
VECO _{out}	88.3 / 77.9	76.9 / 61.1	80.5 / 64.6	81.5 / 64.1	84.2 / 66.8	78.8 / 62.5	80.2 / 66.1	77.0 / 70.4	77.8 / 62.2	82.5 / 63.7	71.6 / 69.4	79.9 / 66.3
VECO _{in}	90.2 / 79.5	81.8 / 66.4	85.4 / 69.8	85.3 / 69.0	87.2 / 70.8	83.7 / 67.9	85.6 / 71.6	80.0 / 74.7	82.4 / 68.6	85.8 / 68.3	74.9 / 73.1	83.9 / 70.9

Table 14: XQuAD results (F1 / EM) for each language.

Model	en	ar	de	es	hi	vi	zh	Avg.
Cross-lingual Transfer								
XML-R	83.5 / 70.6	66.6 / 47.1	70.1 / 54.9	74.1 / 56.6	70.6 / 53.1	74.0 / 52.9	62.1 / 37.0	71.6 / 53.2
VECO _{out}	83.6 / 70.5	65.0 / 44.6	69.8 / 54.6	73.8 / 55.6	69.1 / 51.4	73.1 / 51.8	67.3 / 43.6	71.7 / 53.2
Translate-Train-All								
VECO _{out}	84.1 / 71.3	67.8 / 47.1	70.7 / 55.8	74.6 / 56.6	71.1 / 53.4	74.8 / 54.4	68.8 / 45.8	73.1 / 54.9
VECO _{in}	87.5 / 75.5	72.3 / 52.1	75.7 / 61.1	78.8 / 61.6	76.6 / 58.6	79.3 / 59.1	72.1 / 46.8	77.5 / 59.3

Table 15: MLQA results (F1 / EM) for each language.

Model	en	ar	bn	fi	id	ko	ru	sw	te	Avg.
Cross-lingual Transfer										
XML-R	71.5 / 56.8	67.6 / 40.4	64.0 / 47.8	70.5 / 53.2	77.4 / 61.9	31.9 / 10.9	67.0 / 42.1	66.1 / 48.1	70.1 / 43.6	65.1 / 45.0
VECO _{out}	71.3 / 58.2	73.1 / 52.8	58.9 / 42.5	70.9 / 55.1	77.2 / 60.0	54.2 / 39.9	66.1 / 37.6	65.8 / 45.7	70.6 / 50.7	67.6 / 49.1
Translate-Train-All										
VECO _{out}	77.2 / 64.8	77.0 / 57.5	72.2 / 56.6	76.6 / 59.3	80.0 / 64.4	63.4 / 52.2	72.8 / 50.5	79.4 / 67.1	76.0 / 58.0	75.0 / 58.9
VECO _{in}	79.4 / 65.2	80.1 / 60.9	80.8 / 68.1	81.6 / 65.5	84.3 / 69.7	65.4 / 50.4	77.8 / 55.8	83.7 / 74.1	81.0 / 63.4	79.4 / 63.7

Table 16: TyDiQA-GolP results (F1 / EM) for each language.

Model	af	ar	bg	bn	de	el	es	et	eu	fa	fi	fr	he	hi	hu	id	it	ja
Cross-lingual Transfer																		
XML-R	58.2	47.5	71.6	43	88.8	61.8	75.7	52.2	35.8	70.5	71.6	73.7	66.4	72.2	65.4	77	68.3	60.6
VECO _{out}	48.2	70.9	86.7	57.7	97.5	81.5	94.8	89.7	62.9	82.1	87.9	88.8	74.7	80.7	87.6	89.6	89.2	83.2
Translate-Train-All																		
VECO _{out}	80.9	85.1	91.3	78.1	98.5	89.5	97.4	94.8	79.8	93.1	95.4	93.7	85.8	94.2	93.8	93.0	92.2	92.8
VECO _{in}	88.5	88.7	91.5	84.2	98.9	91.5	97.9	96.4	85.8	95.3	95.9	95.6	89.6	97.0	95.1	94.2	94.1	94.0
	jv	ka	kk	ko	ml	mr	nl	pt	ru	sw	ta	te	th	tl	tr	ur	vi	zh
Cross-lingual Transfer																		
XML-R	14.1	52.1	48.5	61.4	65.4	56.8	80.8	82.2	74.1	20.3	26.4	35.9	29.4	36.7	65.7	24.3	74.7	68.3
VECO _{out}	17.6	58.5	53.9	75.3	80.1	64.2	94.4	92.8	88.6	37.4	61.9	65.8	84.5	52.5	89.3	64.3	85.8	82.7
Translate-Train-All																		
VECO _{out}	35.1	83.0	74.1	88.7	94.8	82.5	95.9	94.6	92.2	69.7	82.4	91.0	94.7	73.0	95.2	63.8	95.1	93.9
VECO _{in}	49.3	86.6	83.7	91.2	97.1	87.9	97.6	96.1	93.8	82.6	88.9	95.3	95.1	79.8	97.6	91.4	97.2	95.2

Table 17: Tatoeba results (Accuracy) for each language