

BERT, mBERT, or BiBERT? A Study on Contextualized Embeddings for Neural Machine Translation

Haoran Xu, Benjamin Van Durme, Kenton Murray

Johns Hopkins University
{h xu64, vandurme, kent on}@jhu.edu

Abstract

The success of bidirectional encoders using masked language models, such as BERT, on numerous natural language processing tasks has prompted researchers to attempt to incorporate these pre-trained models into neural machine translation (NMT) systems. However, proposed methods for incorporating pre-trained models are non-trivial and mainly focus on BERT, which lacks a comparison of the impact that other pre-trained models may have on translation performance. In this paper, we demonstrate that simply using the output (contextualized embeddings) of a tailored and suitable bilingual pre-trained language model (dubbed BiBERT) as the input of the NMT encoder achieves state-of-the-art translation performance. Moreover, we also propose a *stochastic layer selection* approach and a concept of *dual-directional translation model* to ensure the sufficient utilization of contextualized embeddings. In the case of without using back translation, our best models achieve BLEU scores of **30.45** for En→De and **38.61** for De→En on the IWSLT’14 dataset, and **31.26** for En→De and **34.94** for De→En on the WMT’14 dataset, which exceeds all published numbers¹².

1 Introduction

Pre-trained language models (LMs), trained on a large-scale unlabeled data to capture rich representations of the input, such as ELMO (Peters et al., 2018), BERT (Devlin et al., 2019), XLNET (Yang et al., 2019) and XLM (Conneau and Lample, 2019) have increasingly attracted attention in various NLP tasks. Either utilizing context-aware representations of input tokens (Peters et al., 2018) or fine-tuning the pre-trained parameters (Devlin

et al., 2019) both lead to significant improvement for downstream tasks.




Figure 1: The overview of methods: a series of additive improvements to the use of contextualized embeddings on IWSLT’14 dataset. Experimenting over various pre-trained language models, we show that our BiBERT, a bilingual English-German language model, vastly outperforms all other methods (Section 2). Adding *stochastic layer selection* to BiBERT improves performance (Section 3). Finally, innovative *dual-directional training* and fine-tuning with the previous two methods yield around 2 BLEU point gains over the previous state-of-the-art result (Wu et al., 2021) (Section 4).

Inspired by the superior performance of BERT on many other tasks, researchers have investigated leveraging using this pre-trained masked language model to enhance translation models, e.g., initializing the parameters of the model’s encoder with BERT parameters (Rothe et al., 2020), and incorporating the output of BERT to each layer of the encoder (Zhu et al., 2020; Weng et al., 2020). In this paper, we demonstrate simply using the output of a pre-trained language model as the input of NMT systems can achieve state-of-the-art results on IWSLT’14 (Cettolo et al., 2014) and WMT’14 (Bojar et al., 2014) English↔German (En↔De) translation tasks in the case of without using back translation (Sennrich et al., 2016; Edunov et al., 2018)³. After conducting a thor-

¹Code is available at: <https://github.com/felixxu/BiBERT>.

²Our BiBERT is released at: <https://huggingface.co/jhu-clsp/bibert-ende>.

³Though we use a language model that has been trained

ough evaluation of numerous pre-trained language models, we demonstrate that specialized bilingual models perform the best. We then introduce two further refinements, *stochastic layer selection* and *dual-directional training* that yield further improvements. The overview of methods are shown in Figure 1. Overall, our best systems beat published state-of-the-art BLEU scores by around 2 points.

Our main contributions are listed as follows:

- We release our English-German bilingual pre-trained language model, BiBERT, and demonstrate that it outperforms both monolingual and multi-lingual language models for machine translation (Section 2).
- Expanding upon our bilingual language model results, we introduce **stochastic layer selection** which incorporates information from more layers in the pre-trained language model to improve machine translation (Section 3).
- We introduce **dual-directional translation models** which leverages the inherent bilingual nature of BiBERT with mixed domain training and fine-tuning. When combined with stochastic layer selection, it achieves state-of-the-art performance, i.e., **30.45** for En→De and **38.61** for De→En on the IWSLT’14 dataset, and **31.26** for En→De on the WMT’14 dataset (Section 4).

2 Contextualized Embeddings for NMT

2.1 Method

In this section, we focus on investigating the effectiveness of using the output (contextualized embeddings) of the last layer of pre-trained language models on building NMT models. Our basic NMT models are six-layer transformer translation models, though it is model agnostic assuming there are encoder embeddings (Vaswani et al., 2017). Specifically, our method relies on extracting contextualized embeddings of source sentences from the final layer of a frozen pre-trained language model and feeding them to the embedding layer of the NMT encoder. Rather than randomly initializing the source embedding layer, we use the output of these pre-trained models and do not allow these parameters to update during training. To allow for a deep analysis, we concentrate on one language pair, English↔German (En↔De). In the

with additional monolingual data, we only use the provided bitexts during machine translation training.

following subsections, we first explore how much translation performance can be improved by simply using contextualized embeddings, and then explore the internal factors of various pre-trained language models that may affect NMT models. We then introduce our bilingual pre-trained language model and demonstrate that using its contextualized embeddings achieves state-of-the-art results.

2.2 Existing Pre-Trained Models

We first describe four influential pre-trained models that we incorporate into NMT – two monolingual and two multilingual models.

ROBERTA An optimized version of BERT which is trained on a larger dataset, with a dynamic masked language model training regiment that also removes the next sentence prediction (Liu et al., 2019). This model matches or exceeds the performance of BERT on multiple NLP tasks.

GottBERT A state-of-the-art pure German Roberta model (Scheible et al., 2020) trained on 145G German text data portion of OSCAR (Ortiz Suárez et al., 2020), a huge multilingual corpus extracted from Common Crawl. This has been shown to outperform the other two existing German monolingual models (i.e., German BERT⁴ from deepset and dbmz BERT⁵) on NER and text classification tasks.

MBERT (cased) A multilingual BERT (Devlin et al., 2019) pre-trained on 104 highest-resource languages in Wikipedia.

XLM-R (base) A transformer-based (Vaswani et al., 2017) masked language model trained on 100 languages, using more than two terabytes of filtered CommonCrawl data, which outperforms MBERT on a variety of cross-lingual benchmarks (Conneau et al., 2020).

2.3 How Do Pre-Trained LMs Affect NMT?

First we investigate how contextualized embeddings of aforementioned pre-trained language models help NMT models, and explore possible positive and negative factors that may affect NMT models.

Dataset We initially consider a low-resource scenario and then show further experiments in a high-resource scenario in Section 5. We conduct experi-

⁴<https://deepset.ai/german-bert>

⁵<https://github.com/dbmdz/berts>

ments on the IWSLT’14 English-German dataset, which has 160K parallel bilingual sentence pairs.

Settings Our model configuration is `transformer_iwslt_de_en`, a six-layer transformer architecture (Vaswani et al., 2017), with FFN dimension size 1024 and 4 attention heads. We use an embedding dimension of 768 to match the dimension of pre-trained language models. For a consistent comparison with previous works, the evaluation metric is the commonly used tokenized BLEU (Papineni et al., 2002) score calculated with the `multi-bleu.perl` script. More training details are described in Appendix A.

Methods		En→De	De→En
ROBERTA	<i>random</i>	27.3	-
	<i>pre-trained</i>	28.74(+1.44)	-
GOTTBERT	<i>random</i>	-	33.56
	<i>pre-trained</i>	-	36.32(+2.76)
MBERT	<i>random</i>	27.80	34.01
	<i>pre-trained</i>	27.37(-0.43)	34.26(+0.25)
XLM-R	<i>random</i>	27.87	33.67
	<i>pre-trained</i>	27.85(-0.02)	35.38 (+1.71)
BiBERT	<i>random</i>	27.53	33.52
	<i>pre-trained</i>	29.65(+2.12)	37.58(+4.06)

Table 1: IWSLT’14 En↔De BLEU scores utilizing contextualized embeddings from various pre-trained language models. *random* represents the embedding layer of the NMT encoder that is randomly initialized but uses the same vocabulary of the assigned pre-trained language model. *pre-trained* means the embedding layer of the NMT encoder use the output of the assigned frozen pre-trained language model during MT training. Numbers in the bracket show the increment/deduction compared with the corresponding model compared to randomly initialized embeddings.

Observations The main IWSLT’14 results are shown in Table 1. We first conduct experiments with randomly initialized embeddings to obtain baselines. Feeding the output of a pre-trained language model into an NMT model necessitates that the vocabulary of the encoder should be the same as the one used for the language model. To ensure that improvements are not the result of choosing a better vocabulary, we train randomly initialized baseline systems using identical vocabularies for each encoder. For these experiments, the decoder’s vocabulary size is fixed to 8K in order to make fair comparisons. We investigate decoder vocabulary size selection in more detail in Section 2.5. When the embedding layer of the MT encoder is randomly initialized, as opposed to using the pre-trained language model, we ob-

serve similar BLEU scores for all baselines from English-to-German (around 27.6) and German-to-English (around 33.7). By replacing the embedding layer with contextualized embeddings, GOTTBERT boosts the BLEU scores of De→En from 33.56 to 36.32, and ROBERTA strengthens the En→De translation from 27.3 to 28.74. However, the MBERT and XLM-R only provide modest improvement in De→En translation and even degenerate the performance of En→De translation.

Curse of Multilinguality We first note the deterioration caused by MBERT and XLM-R on En→De over the randomly initialized baselines, as well as the comparatively small gains versus the monolingual models of De→En. We hypothesize that contextualized embeddings from MBERT and XLM-R are hurt by the *curse of multilinguality* (Conneau et al., 2020), i.e., low-resource language performance can be improved by adding higher-resource languages during pre-training, but unfortunately high-resource performance suffers and degrades. MBERT and XLM-R are trained on 100 and 104 languages respectively and the *curse of multilinguality* may lead to model capacity issues that degenerate the contextualized embeddings of high-resource languages such as English and German. We attribute the slightly higher improvements of XLM-R over MBERT to the larger amounts of data used in pre-training. The large monolingual models, ROBERTA and GOTTBERT significantly beat a randomized baseline, but also significantly beat the multilingual models. Note that even though XLM-R has 55.6B English tokens used for pre-training, it still helps less than ROBERTA using around 28B English tokens, which is possibly due to interference and constrained capacity (Arivazhagan et al., 2019; Johnson et al., 2017; Tan et al., 2019). Therefore, a suitable pre-trained language model for NMT intuitively should be trained on a large amount of data, but with special care to avoid using too many languages during pre-training.

2.4 Customized Pre-Trained LM

Pre-trained monolingual language models can improve performance of machine translation systems, yet machine translation is inherently a bilingual task. We hypothesize that a pre-trained language model can further improve the translation performance if **its training data is composed of a mixture of texts in both source and target lan-**


guages. In other words, we expect the source and target language data to enrich the contextualized information for each other to better facilitate translation for both directions (En↔De). Therefore, we propose our bilingual pre-trained language models, dubbed **BIBERT**.

Our $\text{BIBERT}_{\text{EN-DE}}$ is based on the RoBERTa architecture (Liu et al., 2019) and implemented using the fairseq framework (Ott et al., 2019). In order to make a direct comparison, $\text{BIBERT}_{\text{EN-DE}}$ is trained on the same German texts as GOTTBERT – just with an additional 146GB of English texts. These are a subset of the English portion in OSCAR – the same dataset the German texts come from. We combine English and German data and shuffle them before training. We train the model using the same number of update steps on German texts as GOTTBERT⁶. We train a unified 52K vocabulary using the WordPiece tokenizer (Wu et al., 2016), with 67GB English and 67GB German texts which are randomly sampled from the training set. $\text{BIBERT}_{\text{EN-DE}}$ is trained on TPU v3-8 for four weeks. More details about optimization for $\text{BIBERT}_{\text{EN-DE}}$ are described in Appendix B.

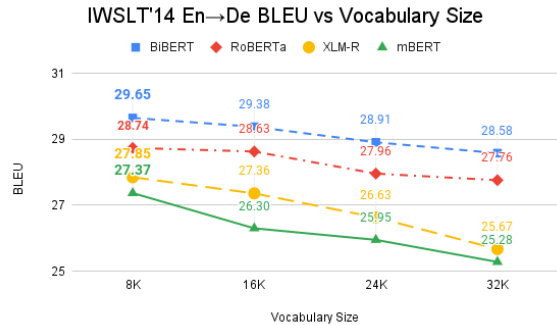
2.5 Vocabulary Size Selection

The vocabulary is fixed for the encoder but still indeterminate for the decoder. In a low-resource machine translation setting, performance is highly sensitive to decoder vocabulary size selection. Gowda and May (2020) demonstrated that a decoder vocabulary using 8K BPE operations performed best across a large grid search. To ensure that 8K vocabulary size is also a suitable choice for the IWSLT’14 (160K parallel sentences) dataset when combined with our method, we search over four candidate decoder vocabulary sizes (8K, 16K, 24K, and 32K) for all aforementioned pre-trained language models. As shown in Figure 2, 8K yields the highest BLEU score for all of our NMT models for De↔En. Thus we select 8K as the vocabulary size of the decoder and use this for all subsequent experiments on IWSLT’14 unless otherwise noted. Interestingly, we also notice that the performance of the translation model with $\text{BIBERT}_{\text{EN-DE}}$ is robust for De→En, and basically unaffected by the vocabulary size.

⁶Note that GOTTBERT uses 100K update steps, but our training data is roughly double that due to the extra English data, so we adopt 200K update steps.



(a) De→En BLEU on IWSLT’14 test set



(b) En→De BLEU on IWSLT’14 test set

Figure 2: En↔De BLEU as a function of vocabulary size with various pre-trained language models on IWSLT’14 test set. Models obtain highest scores with 8K vocabulary size.

2.6 BIBERT Performance

$\text{BIBERT}_{\text{EN-DE}}$ results As indicated in the last row of Table 1, our bilingual model help the transformer model achieve **29.65** score for En→De — a gain of **2.12** over the baseline. For De→En, our model achieves **37.58** score with a gain of **4.06**. Recall that $\text{BIBERT}_{\text{EN-DE}}$ uses the same settings as GOTTBERT — the only difference is the addition of extra English training data from the OSCAR corpus — yet $\text{BIBERT}_{\text{EN-DE}}$ yields an additional 1.30 BLEU point improvement over GOTTBERT.

Analysis Based on the superior performance of $\text{BIBERT}_{\text{EN-DE}}$, we hypothesize that contextualized embeddings output from $\text{BIBERT}_{\text{EN-DE}}$ contain richer German information than GOTTBERT and better assist the model in translation by learning extra English data. Furthermore, we theorize training on German texts also enhances the quality of English contextualized embeddings — note that even though ROBERTA and $\text{BIBERT}_{\text{EN-DE}}$ are not directly comparable due to different English pre-training data, $\text{BIBERT}_{\text{EN-DE}}$ still had a 0.68 BLEU point improvement over ROBERTA

even while using less English training data. Some other explanations for the superior performance of BiBERT_{EN-DE} are 1) it learns the aligned embeddings for the tokens with similar meanings across two languages. Hence, the source embeddings can offer the encoder a hint of aligned target embeddings to help translation. 2) Embeddings of overlapping En-De sub-word units⁷ fed to NMT encoders may facilitate translation by bilingual information.

Algorithms	De → En
Adversarial MLE (Wang et al., 2019)	35.18
DynamicConv(Wu et al., 2019)	35.20
Macaron Net (Lu* et al., 2020)	35.40
BERT-Fuse (Zhu et al., 2020)	36.11
MAT (Fan et al., 2020)	36.22
Mixed Representations (Wu et al., 2020)	36.41
UniDrop (Wu et al., 2021)	36.88
Ours, GOTTBERT	36.32
Ours, BiBERT	37.58

Table 2: Comparison of our work and most recent existing methods on IWSLT’14 De→En.

2.7 Comparison with Existing Work

Table 2 shows a comparison of our work with the recent literature on IWSLT’14 German to English translation. These works propose improvements to transformer models in different aspects, e.g., incorporating BERT into every layer of encoders and decoders with additional multi-head attentions (Zhu et al., 2020), multi-branch encoders (Fan et al., 2020), mixed representations from different tokenizers (Wu et al., 2020) and uniting different dropout techniques into NMT models (Wu et al., 2021). Our straightforward method of simply using the final layer of BiBERT_{EN-DE} outperforms all of them. Furthermore, even the model that only uses the monolingual GOTTBERT achieves a competitive result (36.32) compared with the previous state-of-the-art approach (36.88). Our method is easy to implement, so it can be used in conjunction with other methods in the literature.

2.8 Time Costs

Leveraging an external pre-trained language model leads to higher computational complexity. Our approach takes approximately 20% additional time during training and 13% extra time during inference. Considering the significant BLEU gains, we argue that they justify the higher time costs.

⁷Such as ##n, which uses shared En-De information.

3 Layer Selection

Jawahar et al. (2019) demonstrates that different layers of BERT capture differing linguistic information in a rich, hierarchical structure that mimics classical, compositional tree-like structures. Information in the lower layer (e.g., phrase-level information) gets gradually diluted in higher layers. Thus, to potentially leverage more information encapsulated in the pre-trained language models, we are also interested in exploring how other layers of contextualized embeddings can improve NMT models — rather than simply using the last layer.

We denote \mathcal{X} as the collection of source language sentences. For each source sentence $x \in \mathcal{X}$, let $H_B^i(x)$ denote the contextualized embeddings of x obtained from the i th layer of the pre-trained language model. In our settings, we consider top K layers of the pre-trained language model, i.e., we consider $H_B^i(x) \forall i \in [M - K + 1, M]$, where K is a hyperparameter, and M is the total number of layers of the pre-trained language model.

3.1 Stochastic Layer Selection

During training of deep neural networks, various methods of stochastically freezing groups of parameters in a model for individual training examples have been shown to improve performance. For instance, dropout (Srivastava et al., 2014) samples parameters from a Bernoulli distribution to not update, and drop-net (Zhu et al., 2020) and drop-branch (Fan et al., 2020) randomly active a candidate net and freeze the others in a uniform distribution. We propose *stochastic layer selection*, a novel approach to encapsulate more features and information from more layers of the pre-trained language models. Specifically, for each batch, we randomly pick the output from one layer rather than all of them as the input for the NMT encoder (Figure 3). We denote the input embeddings of sentence x to the NMT encoder as $H_E(x)$, which is defined in the following way during training:

$$H_E(x) = \sum_{i=1}^K \mathbb{1}\left(\frac{i-1}{K} < p \leq \frac{i}{K}\right) H_B^{M-i+1}(x) \quad (1)$$

where $\mathbb{1}(\cdot)$ is the indicator function and p is a random variable which is uniformly sampled from $[0,1]$. In the inference step, the output is the expectation of outputs of all layers used for training, i.e., $\mathbb{E}_{p \sim \text{uniform}[0,1]}[H_E(x)]$, which leads to the modifi-

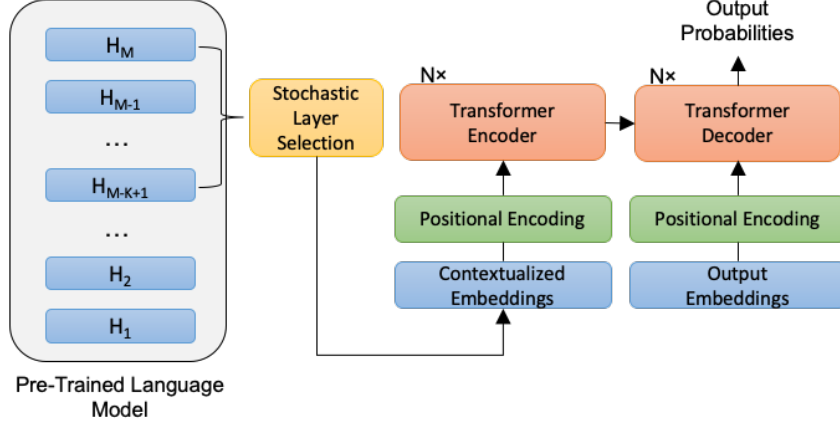


Figure 3: The overall framework of *stochastic layer selection* method. Top K layers of the pre-trained language model are considered and fed to the NMT encoder.

cation of Equation 1:

$$H_E(x) = \frac{1}{K} \sum_{i=1}^K H_B^{M-i+1}(x). \quad (2)$$

3.2 Experiments and Results

Based on the results of Table 1, we select the pre-trained model performing best for NMT, BiBERT_{EN-DE}, and use it as the basis for all subsequent experiments. To be consistent with the results in Section 2, we once again use the IWSLT’14 dataset. Figure 4 illustrates the impact of stochastic layer selection. We conduct experiments for En↔De with the number of layers K ranging from 2 to M ($M = 12$ for BiBERT_{EN-DE}). Note that setting $K = 1$ reduces to the case of only selecting the last layer as in Section 2.

In all cases, the stochastic layer selection obtains substantial gains compared with our previous best scores in En→De (29.65) and De→En (37.58) in Section 2. In both situations of En→De and De→En, the translation model gets the highest score (37.94 for De→En and 30.04 for En→De) when stochastic layer selection uses 8 layers.

4 One Model, Dual-Directional Translations

In this section, different from ordinary one-way translation models, we introduce our dual-directional translation models, i.e., a model can translate both En→De and De→En. The model architecture is the same as the one in Section 3.

One of the biggest advantages of the shared English-German vocabulary of BiBERT_{EN-DE} is that our encoder has the capability of receiving




Figure 4: IWSLT’14 En→De (left, blue bars) and De→En (right, red bars) BLEU as a function of number of layers K considered in the *stochastic layer selection* module for NMT models. Note that when $K = 1$, it reduces to the case of selecting the last layer. However, any value of $K > 1$ selected for stochastic layer selection beats this very strong baseline with $K = 8$ obtaining the highest BLEU scores in both directions.

contextualized embeddings of both source and target tokens. During the training step, we feed source sentences to the model and expect the generation of a target translation, yet also, inversely, feed target sentences and expect translations in the source language. The motivation behind the dual-directional translation model is that we expect the contextualized representations of source and target sentences could enhance each other to build a better encoder for the translation model. From the aspect of data augmentation, the target sentences play a role in augmented data in the task of translating from the source language to the target language, vice versa. With the method of swapping source and target sentences once as an additional dataset, our experiments show superior performance for both direc-

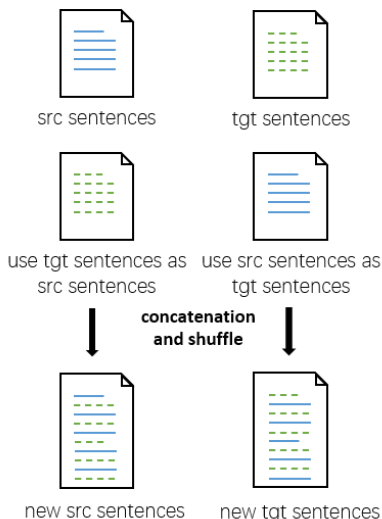


Figure 5: Workflow of data preprocessing. We swap source and target sentences, and concatenate swapped sentence pairs and original sentence pairs. Finally, we shuffle the concatenated data for dual-directional translation model training.

tional translations. Two advantages of this method are 1) obtaining improvement without extra bitexts, and 2) only slight modification for data preprocessing and no changes for the model architecture.

4.1 Dataset Preprocessing

For consistent comparisons, the dataset is still IWSLT’14 En→De. The details of data preprocessing for the dual-directional translation model are illustrated in Figure 5. Using only the same exact parallel sentences in our bitext for training, we simply leverage the dataset in reverse, by swapping our original target sentences to use as new source sentences and original source sentences as new target sentences. We then concatenate and jointly shuffle the original and new data to acquire our mixed training data. We use a joint English-German vocabulary of size 12k for the decoder.

4.2 Fine-tuning

Inspired by the findings of Xu et al. (2021), where training on a mix of in- and out-of-domain of data initially, and then gradually fine-tuning until only in-domain data is used, substantially improved model performance, we treat our concatenated sentences as mixed domain data, and the source and target languages are separate language domains. Each language data can be the out-of-domain data for the other language. Following this perspective, we first train our dual-directional model on mixed

data, and then fine-tune it on the source or target data to obtain one-way translation models⁸.

4.3 Experiments and Results

We additionally conduct one-way translation models with 12K bilingual vocabulary to have a fair baseline for dual-directional models. Overall results are shown in Table 3. We first discuss the models trained without stochastic layer selection. The dual-directional model substantially outperforms the one-way model by obtaining a gain of 0.52 in En→De and 0.72 in De→En. Moreover, fine-tuning on the in-domain data further improves BLEU from 29.89 to 30.33 in En→De and from 37.97 to 38.12 in De→En. Both positive results indicated by the dual-directional model and fine-tuning approach show their effectiveness in helping translation. A similar discussion holds for the models with the stochastic layer selection method. Compared with our previous models in Section 3 (30.04 En→De and 37.94 in De→En), our best model achieves new state-of-the-art results both in En→De and De→En, which respectively obtain **30.45** and **38.61** BLEU.

Methods	En→De	De→En
<i>No Stochastic Layer Selection:</i>		
One-Way (vocab size=12K)	29.37	37.25
Dual-Directional Training + Fine-Tuning	29.89	37.97
	30.33	38.12
<i>Stochastic Layer Selection, K = 8:</i>		
One-Way (vocab size=12K)	30.00	37.69
Dual-Directional Training + Fine-Tuning	30.30	38.37
	30.45	38.61

Table 3: Comparison of dual-directional and ordinary (one-way) translation models, with and without stochastic layer selection, on IWSLT’14 En↔De.

5 High-Resource Scenario

5.1 Dataset and Training Details

For the high-resource scenario, we evaluate models on the WMT’14 English-German dataset, which contains 4.5M parallel sentence pairs. We combine newstest2012 and newstest2013 as the validation set and use newstest2014 (3003 sentence pairs) as the test set. Our model configuration is `transformer_vaswani_wmt_en_de_big`, a ‘big’ transformer with 4096 FFN dimension and

⁸In the view of the same amount of training data for the source and target language, we only apply one-stage fine-tuning, which is slightly different from the multiple stages of fine-tuning in Xu et al. (2021).

Methods	En → De	De → En
Transformer + Large Batch (Ott et al., 2018)	29.3	-
Evolved Transformer (So et al., 2019)	29.8	-
BERT Initialization (12 layers) (Rothe et al., 2020)	30.6	33.6
BERT-Fuse (Zhu et al., 2020)	30.75	-
BiBERT Contextualized Embeddings + Stochastic Layer Selection	30.91	34.94
+ Dual-Directional Training	30.31	34.54
+ Fine-Tuning	31.26	34.68

Table 4: WMT’14 En↔De results on newstest2014 test set.

16 attention heads. We replace hidden size 1024 with 768 to keep the dimensions consistent with BiBERT_{EN-DE}. The evaluation strategy is the same as IWSLT’14 tasks.

Following the findings that En↔De translation has similar results for vocabularies ranging from 32K to 64K in high-resource scenarios (4.5M training samples) (Gowda and May, 2020), we use a bilingual vocabulary with 52K size for the decoder, which is larger than the ones (8K and 12K) used in IWSLT experiments.

5.2 Results

We compare our methods with prior existing works that achieve highest scores by only using provided bi-texts in Table 4. With BiBERT_{EN-DE} contextualized embeddings and stochastic layer selection, our model achieves state-of-the-art BLEU both on En→De (30.91) and De→En (34.94). Interestingly, dual-directional translation training does not show the same strong effectiveness as it did in the low-resource scenario. One possible reason is that model capacity is not large enough to handle mixed domain data (Arivazhagan et al., 2019). However, it still additively improves En→De to **31.26** BLEU. It is worth mentioning that our NMT model achieves better performance with less training parameters — the hidden size of our NMT model is 768 but 1024 for the prior existing works.

6 Related Work

6.1 Pre-Trained Embeddings

Traditional pre-trained embeddings are investigated in type level, e.g., word2vec (Mikolov et al., 2013), glove (Pennington et al., 2014) and fastText (Bojanowski et al., 2017). Peters et al. (2018) moved further from this line and proposed context-aware embeddings output from pre-trained bidirectional LSTM (ELMo). Following the attention-based transformer module (Vaswani et al., 2017), the architectures of GPT models (Radford et al., 2018, 2019; Brown et al., 2020) and BERT (Devlin et al.,

2019) respectively are based on stacking deep transformer decoders and encoders and significantly boost downstream tasks. Beyond pure English models, pre-trained language models for other languages have also showed up, e.g., CAMEMBERT for French (Martin et al., 2020) and ARABERT for Arabic (Baly et al., 2020). Multilingual representations, e.g. mBERT and XLMs (Conneau and Lample, 2019) have been shown to be effective to facilitate cross-lingual learning. XLM-R (Conneau et al., 2020), a model learning cross-lingual representation at scale achieved state-of-the-art results on multiple cross-lingual benchmarks. Recently, an English-Arabic bilingual BERT (Lan et al., 2020) outperformed ARABERT, mBERT and XLM-R on supervised and zero-shot transfer settings.

6.2 MT with Context-Aware Representations

Imamura and Sumita (2019) removed the NMT encoder part and directly fed the output of BERT to the attention mechanism in the decoder. They train the model with two optimization stages, i.e., only training the decoder and fine-tuning BERT. Similarly, Clinchant et al. (2019) have incorporated BERT into NMT models by replacing the embedding layer with BERT parameters and initializing encoder with BERT, but they still notice that NMT model with BERT is not as robust as expected. Rothe et al. (2020) also leveraged pre-trained checkpoints (e.g., BERT and GPT) to initialize 12-layer NMT encoder and decoder and achieved state-of-the-art results. Interestingly, they showed that the models with decoder initialized by GPT fail to improve the translation performance and are even worse than the one whose decoder is randomly initialized. Similarly, Ma et al. (2020) initialize both transformer encoder and decoder by XLM-R but fine-tune it on multiple bilingual corpora to obtain a multilingual translation model. The preliminary experiments from Zhu et al. (2020) indicate that NMT models simply fed by the output of BERT outperform the models initialized

by BERT or XLM. However, only limited experiments and little analysis on this method has been done in their work. They mainly focused on the *BERT-fuse* approach, i.e., the output of BERT is fed to each layer of NMT encoder and decoder with extra multi-head attentions. Instead of only using the last layer of BERT, [Weng et al. \(2020\)](#) introduced layer-aware attention mechanism to capture compound contextual information from BERT. Moreover, they also proposed the knowledge distillation paradigm to learn pre-trained representation in the training process. On an English-Arabic translation task, [Yarmohammadi et al. \(2021\)](#) use a precursor of this method though it lacks all of the refinements described here. However, it was shown to further help in downstream cross-lingual information extraction tasks.

7 Conclusion

We have shown that our BiBERT trained on a large amount of mixed texts of the source and target languages can better help NMT models improve translation performance compared with other existing pre-trained language models and achieve state-of-the-art results by simply using the output of the last layer. Moreover, we introduce the stochastic layer selection method and demonstrated its effectiveness in improving translation performance. Finally, experiments on the dual-directional translation model illustrate that source and target data can augment each other to further boost performance.

Acknowledgements

We thank anonymous reviewers for their valuable comments. We thank Kelly Marchisio and Kevin Duh for their helpful suggestions. We also thank the Google TFRC program for providing free TPU access. This work was supported in part by IARPA BETTER (#2019-19051600005). The views and conclusions contained in this work are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, or endorsements of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George Foster, Colin Cherry, et al. 2019. Massively multilingual neural machine translation in the wild: Findings and challenges. *arXiv preprint arXiv:1907.05019*.
- Fady Baly, Hazem Hajj, et al. 2020. Arabert: Transformer-based model for arabic language understanding. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 9–15.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, et al. 2014. Findings of the 2014 workshop on statistical machine translation. In *Proceedings of the ninth workshop on statistical machine translation*, pages 12–58.
- Nicolas Boulanger-Lewandowski, Yoshua Bengio, and Pascal Vincent. 2013. Audio chord recognition with recurrent neural networks. In *ISMIR*, pages 335–340. Citeseer.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Mauro Cettolo, J Niehues, S Stüker, Luisa Bentivogli, and Marcello Federico. 2014. Report on the 11th iwslt evaluation campaign, iwslt 2014. In *IWSLT-International Workshop on Spoken Language Processing*, pages 2–17. Marcello Federico, Sebastian Stüker, François Yvon.
- Stephane Clinchant, Kweon Woo Jung, and Vassilina Nikoulina. 2019. [On the use of BERT for neural machine translation](#). In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 108–117, Hong Kong. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco

- Guzmán, Édouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- 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.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. Understanding back-translation at scale. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500.
- Yang Fan, Shufang Xie, Yingce Xia, Lijun Wu, Tao Qin, Xiang-Yang Li, and Tie-Yan Liu. 2020. Multi-branch attentive transformer. *arXiv preprint arXiv:2006.10270*.
- Thamme Gowda and Jonathan May. 2020. [Finding the optimal vocabulary size for neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3955–3964, Online. Association for Computational Linguistics.
- Kenji Imamura and Eiichiro Sumita. 2019. [Recycling a pre-trained BERT encoder for neural machine translation](#). In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 23–31, Hong Kong. Association for Computational Linguistics.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. [What does BERT learn about the structure of language?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. 2017. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- Wuwei Lan, Yang Chen, Wei Xu, and Alan Ritter. 2020. [An empirical study of pre-trained transformers for Arabic information extraction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4727–4734, Online. Association for Computational Linguistics.
- 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*.
- Yiping Lu*, Zhuohan Li*, Di He, Zhiqing Sun, Bin Dong, Tao Qin, Liwei Wang, and Tie yan Liu. 2020. [Understanding and improving transformer from a multi-particle dynamic system point of view](#). In *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations*.
- Shuming Ma, Jian Yang, Haoyang Huang, Zewen Chi, Li Dong, Dongdong Zhang, Hany Hassan Awadalla, Alexandre Muzio, Akiko Eriguchi, Saksham Singhal, et al. 2020. Xlm-t: Scaling up multilingual machine translation with pretrained cross-lingual transformer encoders. *arXiv preprint arXiv:2012.15547*.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric Villenote de la Clergerie, Djamé Seddah, and Benoît Sagot. 2020. Camembert: a tasty french language model. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Pedro Javier Ortiz Suárez, Laurent Romary, and Benoît Sagot. 2020. [A monolingual approach to contextualized word embeddings for mid-resource languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1703–1714, Online. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. [Scaling neural machine translation](#). In *Proceedings of the Third Conference on*

- Machine Translation: Research Papers*, pages 1–9, Brussels, Belgium. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Sascha Rothe, Shashi Narayan, and Aliaksei Severyn. 2020. [Leveraging pre-trained checkpoints for sequence generation tasks](#). *Transactions of the Association for Computational Linguistics*, 8:264–280.
- Raphael Scheible, Fabian Thomczyk, Patric Tippmann, Victor Jaravine, and Martin Boeker. 2020. Götbert: a pure german language model. *arXiv preprint arXiv:2012.02110*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96.
- David So, Quoc Le, and Chen Liang. 2019. The evolved transformer. In *International Conference on Machine Learning*, pages 5877–5886. PMLR.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.
- Xu Tan, Yi Ren, Di He, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2019. Multilingual neural machine translation with knowledge distillation. *arXiv preprint arXiv:1902.10461*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Dilin Wang, ChengYue Gong, and Qiang Liu. 2019. [Improving neural language modeling via adversarial training](#). In *ICML*, pages 6555–6565.
- Rongxiang Weng, Heng Yu, Shujian Huang, Shanbo Cheng, and Weihua Luo. 2020. Acquiring knowledge from pre-trained model to neural machine translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9266–9273.
- Felix Wu, Angela Fan, Alexei Baevski, Yann Dauphin, and Michael Auli. 2019. [Pay less attention with lightweight and dynamic convolutions](#). In *International Conference on Learning Representations*.
- Lijun Wu, Shufang Xie, Yingce Xia, Yang Fan, Jian-Huang Lai, Tao Qin, and Tieyan Liu. 2020. Sequence generation with mixed representations. In *International Conference on Machine Learning*, pages 10388–10398. PMLR.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Zhen Wu, Lijun Wu, Qi Meng, Yingce Xia, Shufang Xie, Tao Qin, Xinyu Dai, and Tie-Yan Liu. 2021. [UniDrop: A simple yet effective technique to improve transformer without extra cost](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3865–3878, Online. Association for Computational Linguistics.
- Haoran Xu, Seth Ebner, Mahsa Yarmohammadi, Aaron Steven White, Benjamin Van Durme, and Kenton Murray. 2021. [Gradual fine-tuning for low-resource domain adaptation](#). In *Proceedings of the Second Workshop on Domain Adaptation for NLP*, pages 214–221, Kyiv, Ukraine. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Mahsa Yarmohammadi, Shijie Wu, Marc Marone, Haoran Xu, Seth Ebner, Guanghui Qin, Yunmo Chen, Jialiang Guo, Craig Harman, Kenton Murray, Aaron Steven White, Mark Dredze, and Benjamin

Van Durme. 2021. Everything is all it takes: A multipronged strategy for zero-shot cross-lingual information extraction. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.

Jinhua Zhu, Yingce Xia, Lijun Wu, Di He, Tao Qin, Wengang Zhou, Houqiang Li, and Tiejun Liu. 2020. [Incorporating bert into neural machine translation](#). In *International Conference on Learning Representations*.

A Training Details

A.1 IWSLT'14 Training Details

We use 4 NVIDIA 2080 Ti GPUs with 2048 tokens per GPU and accumulate the gradient 4 times. The learning rate is 0.0004. The optimizer is Adam (Kingma and Ba, 2014) with `inverse_sqrt` learning rate scheduler. At inference time, we use beam search with width 4 and use a length penalty of 0.6 (Boulanger-Lewandowski et al., 2013; Wu et al., 2016; Koehn and Knowles, 2017).

A.2 WMT'14 Training Details

We use 4 NVIDIA V100 GPUs with a batch size of 4096 tokens per GPU. Following the recommendation of the training settings from Ott et al. (2018), we accumulate the gradient 32 times to simulate 128 GPU training settings. We set the initial learning rate as 0.001.

B Optimization for BiBERT

We use `fairseq` to pre-train our bilingual language models on an 8-core TPU v3-8. We train the models in 200K update steps using a batch size of 8192. We use Adam optimizer (Kingma and Ba, 2014) with a learning rate of $4e-4$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, L2 weight decay of 0.01. The learning rate is warmed up over the first 20K steps to a peak value of $4e-4$, from which the learning rate polynomially decayed. We apply a dropout rate of 0.1 to all layers.