

Evaluating the Cross-Lingual Effectiveness of Massively Multilingual Neural Machine Translation

Aditya Siddhant, Melvin Johnson, Henry Tsai, Naveen Ari,
Jason Riesa, Ankur Bapna, Orhan Firat, Karthik Raman

Google Research

{adisid, melvinp, henrytsai, navari, reisa, ankurbpn, orhanf, karthikraman}@google.com

Abstract

The recently proposed massively multilingual neural machine translation (NMT) system has been shown to be capable of translating over 100 languages to and from English within a single model (Aharoni, Johnson, and Firat 2019). Its improved translation performance on low resource languages hints at potential cross-lingual transfer capability for downstream tasks. In this paper, we evaluate the cross-lingual effectiveness of representations from the encoder of a massively multilingual NMT model on 5 downstream classification and sequence labeling tasks covering a diverse set of over 50 languages. We compare against a strong baseline, multilingual BERT (mBERT) (Devlin et al. 2018), in different cross-lingual transfer learning scenarios and show gains in zero-shot transfer in 4 out of these 5 tasks.

1 Introduction

English has an abundance of labeled data that can be used for various Natural Language Processing (NLP) tasks, such as part-of-speech tagging (POS), named entity recognition (NER), and natural language inference (NLI). This richness of labeled data manifests itself as a boost in accuracy in the current era of data-hungry deep learning algorithms. However, the same is not true for many other languages where task specific data is scarce and expensive to acquire. This motivates the need for cross-lingual transfer learning – the ability to leverage the knowledge from task specific data available in one or more languages to solve that task in languages with little or no task-specific data.

Recent progress in NMT has enabled one to train multilingual systems that support translation from multiple source languages into multiple target languages within a single model (Firat, Cho, and Bengio 2016; Johnson et al. 2017; Aharoni, Johnson, and Firat 2019). Such multilingual NMT (mNMT) systems often demonstrate large improvements in translation quality on low resource languages. This positive transfer originates from the model’s ability to learn representations which are transferable across languages. Previous work has shown that these representations can then be used for cross-lingual transfer in other downstream NLP tasks -

albeit on only a pair of language pairs (Eriguchi et al. 2018), or by limiting the decoder to use a pooled vector representation of the entire sentence from the encoder (Artetxe and Schwenk 2018).

In this paper we scale up the number of translation directions used in the NMT model to include 102 languages to and from English. Unlike Artetxe and Schwenk (2018), we do not apply any restricting operations such as pooling while training mNMT which allows us to obtain token level representations making it possible to transfer them to sequence tagging tasks as well. We find that mNMT models trained using plain translation losses can out of the box emerge as competitive alternatives to other methods at the forefront of cross-lingual transfer learning (Devlin et al. 2018; Artetxe and Schwenk 2018)

Our contributions in this paper are threefold:

- We use representations from a Massively Multilingual Translation Encoder (MMTE) that can handle 103 languages to achieve cross-lingual transfer on 5 classification and sequence tagging tasks spanning more than 50 languages.
- We compare MMTE to mBERT in different cross-lingual transfer scenarios including zero-shot, few-shot, fine-tuning, and feature extraction scenarios.
- We outperform the state-of-the-art on zero-shot cross-lingual POS tagging [Universal Dependencies 2.3 dataset (Nivre et al. 2018)], intent classification (Schuster et al. 2018), and achieve results comparable to state-of-the-art on document classification [ML-Doc dataset (Schwenk and Li 2018)].

The remainder of this paper is organized as follows. Section 2 describes our MMTE model in detail and points out its differences from mBERT. All experimental details, results and analysis are given in Sections 3 and 4. This is followed by a discussion of related work. In Section 6, we summarize our findings and present directions for future research. We emphasize that the primary motivation of the paper is not to challenge the state-of-the-art but instead to investigate the effectiveness of representations learned from an mNMT model in various transfer-learning settings.

2 Massively Multilingual Neural Machine Translation Model

In this section, we describe our massively multilingual NMT system. Similar to BERT, our transfer learning setup has two distinct steps: *pre-training* and *fine-tuning*. During pre-training, the NMT model is trained on large amounts of parallel data to perform translation. During fine-tuning, we initialize our downstream model with the pre-trained parameters from the encoder of the NMT system, and then all of the parameters are fine-tuned using labeled data from the downstream tasks.

2.1 Model Architecture

We train our Massively Multilingual NMT system using the Transformer architecture (Vaswani et al. 2017) in the open-source implementation under the Lingvo framework (Shen et al. 2019). We use a larger version of Transformer Big containing 375M parameters (6 layers, 16 heads, 8192 hidden dimension) (Chen et al. 2018), and a shared source-target sentence-piece model (SPM)¹ (Kudo and Richardson 2018) vocabulary with 64k individual tokens. All our models are trained with Adafactor (Shazeer and Stern 2018) with momentum factorization, a learning rate schedule of $(3.0, 40k)^2$ and a per-parameter norm clipping threshold of 1.0. The encoder of this NMT model comprises approximately 190M parameters and is subsequently used for fine-tuning.

2.2 Pre-training

Objective We train a massively multilingual NMT system which is capable of translating between a large number of language pairs at the same time by optimizing the translation objective between language pairs. To train such a multilingual system within a single model, we use the strategy proposed in (Johnson et al. 2017) which suggests prepending a target language token to every source sequence to be translated. This simple and effective strategy enables us to share the encoder, decoder, and attention mechanisms across all language pairs.

Data We train our multilingual NMT system on a massive scale, using an in-house corpus generated by crawling and extracting parallel sentences from the web (Uszkoreit et al. 2010). This corpus contains parallel documents for 102 languages, to and from English, comprising a total of 25 billion sentence pairs. The number of parallel sentences per language in our corpus ranges from around 35 thousand to almost 2 billion. Figure 1 illustrates the data distribution for all 204 language pairs used to train the NMT model. Language ids for all the languages are also provided in supplementary material.

¹<https://github.com/google/sentencepiece>

²The shorthand form $(3.0, 40k)$ corresponds to a learning rate of 3.0, with 40k warm-up steps for the schedule, which is decayed with the inverse square root of the number of training steps after warm-up.




Figure 1: Per language pair data distribution of the dataset used for our multilingual NMT model. The y-axis depicts the number of training examples available per language pair on a logarithmic scale. Dataset sizes range from 35k for the lowest resource language pairs to 2 billion for the largest.

Data sampling policy Given the wide distribution of data across language pairs, we used a temperature based data balancing strategy. For a given language pair, l , let D_l be the size of the available parallel corpus. Then if we adopt a naive strategy and sample from the union of the datasets, the probability of the sample being from language pair l will be $p_l = \frac{D_l}{\sum_l D_l}$. However, this strategy would starve low resource language pairs. To control for the ratio of samples from different language pairs, we sample a fixed number of sentences from the training data, with the probability of a sentence belonging to language pair l being proportional to $p_l^{\frac{1}{T}}$, where T is the sampling temperature. As a result, $T = 1$ would correspond to a true data distribution, and, $T = 100$ yields an (almost) equal number of samples for each language pair (close to a uniform distribution with over-sampling for low-resource language-pairs). We set $T = 5$ for a balanced sampling strategy. To control the contribution of each language pair when constructing the vocabulary, we use the same temperature based sampling strategy with $T = 5$. Our SPM vocabulary has a character coverage of 0.999995.

Model quality We use BLEU score (Papineni et al. 2002) to evaluate the quality of our translation model(s). Our mNMT model performs worse than the bilingual baseline on high resource language pairs but improves upon it on low resource language pairs. The average drop in BLEU score on 204 language pairs as compared to bilingual baselines is just 0.25 BLEU. This is impressive considering we are comparing one multilingual model to 204 different bilingual models. Table 1 compares the BLEU scores achieved by mNMT to that of the bilingual baselines on 10 representative language pairs.³ These scores are obtained on an internal evaluation set which contains around 5k examples per language

³We chose a diverse set of language pairs with varying language families, scripts, and dataset sizes.

	baseline*	mNMT		baseline*	mNMT
ur-en	27.84	27.24	en-bg	31.31	29.36
mr-en	27.81	28.61	en-es	35.23	34.35
be-en	24.23	24.66	en-sw	18.73	19.79
ru-en	26.54	26.56	en-pt	37.19	37.41
fr-en	37.49	34.02	en-hi	16.46	16.63

Table 1: BLEU scores on ten language pairs with mNMT. *baseline refers to a transformer based bilingual NMT model trained on only one language pair.

pair.

2.3 Fine-tuning mNMT Encoder

Fine-tuning involves taking the encoder of our mNMT model, named Massively Multilingual Translation Encoder (MMTE), and adapting it to the downstream task. For tasks which involve single input, the text is directly fed into the encoder. For tasks such as entailment which involve input pairs, we concatenate the two inputs using a separator token and pass this through the encoder. For each downstream task, the inputs and outputs are passed through the encoder and we fine-tune all the parameters end-to-end. The encoder encodes the input through the stack of Transformer layers and produces representations for each token at the output. For sequence tagging tasks, these token level representations are individually fed into a task-specific output layer. For classification or entailment tasks, we apply max-pooling on the token level representations and feed this into the task-specific output layer.

It should be noted that fine-tuning is relatively inexpensive and fast. All of the results can be obtained within a few thousand gradient steps. The individual task-specific modeling details are described in detail in section 3. It is also important to note that while the encoder, the attention mechanism, and the decoder of the model are trained in the pre-training phase, only the encoder is used during fine-tuning.

2.4 Differences with mBERT

We point out some of the major difference between mBERT and MMTE are:

- mBERT uses two unsupervised pre-training objectives called masked language modeling (MLM) and next sentence prediction (NSP) which are both trained on monolingual data in 104 languages. MMTE on the other hand uses parallel data in 103 languages (102 languages to and from English) for supervised training with negative log-likelihood as the loss. It should be noted that mBERT uses clean Wikipedia data while MMTE is pre-trained on noisy parallel data from the web.
- mBERT uses 12 transformer layers, 12 attention heads, 768 hidden dimensions and has 178M parameters while MMTE uses 6 transformer layers, 16 attention heads, and 8196 hidden dimensions with 190M parameters. Note that, the effective capacity of these two models cannot easily be compared by simply counting number of pa-

rameters, due to the added characteristic complexity with depth and width.

- MMTE uses SPM to tokenize input with 64k vocabulary size while mBERT uses a Wordpiece model (Wu et al. 2016) with 110k vocabulary size.

3 Experiments and Results

As stated earlier, we use MMTE to perform downstream cross-lingual transfer on 5 NLP tasks. These include 3 classification tasks: NLI (XNLI dataset), document classification (MLDoc dataset) and intent classification, and 2 sequence tagging tasks: POS tagging and NER. We detail all of the experiments in this section.

3.1 XNLI: Cross-lingual NLI

XNLI is a popularly used corpus for evaluating cross-lingual sentence classification. It contains data in 15 languages (Conneau et al. 2018). Evaluation is based on classification accuracy for pairs of sentences as one of entailment, neutral, or contradiction. We feed the text pair separated by a special token into MMTE and add a small network on top of it to build a classifier. This small network consists of a pre-pool feed-forward layer with 64 units, a max-pool layer which pools word level representations to get the sentence representation, and a post-pool feed-forward layer with 64 units. The optimizer used is Adafactor with a learning rate schedule of (0.2, 90k). The classifier is trained on English only and evaluated on all the 15 languages. Results are reported in Table 2. Please refer to Appendix Table 1 for language names associated with the codes.

	mBERT*	MMTE	SOTA [†]
en	82.1	79.6	85.0
ar	64.9	64.9	73.1
bg	68.9	70.4	77.4
de	71.1	68.2	77.8
el	66.4	67.3	76.6
es	74.3	71.6	78.9
fr	73.8	69.5	78.7
hi	60.0	63.5	69.6
ru	69.0	66.2	75.3
sw	50.4	61.9	68.4
th	55.8	66.2	73.2
tr	61.6	63.6	72.5
ur	58.0	60.0	67.3
vi	69.5	69.7	76.1
zh	69.3	69.2	76.5
avg	66.3	67.5	75.1

Table 2: Accuracies on the test set of the XNLI dataset. *mBERT numbers have been taken from Wu and Dredze (2019). [†] SOTA in the last column refers to the MLM + translation language modeling (TLM) results reported in Lample and Conneau (2019).

MMTE outperforms mBERT on 9 out of 15 languages

and by 1.2 points on average. BERT achieves excellent results on English, outperforming our system by 2.5 points but its zero-shot cross-lingual transfer performance is weaker than MMTE. We see most gains in low resource languages such as ar, hi, ur, and sw. MMTE however falls short of the current state-of-the-art (SOTA) on XNLI (Lample and Conneau 2019). We hypothesize this might be because of 2 reasons: (1) They use only the 15 languages associated with the XNLI task for pre-training their model, and (2) They use both monolingual and parallel data for pre-training while we just use parallel data. We confirm our first hypothesis later in Section 4 where we see that decreasing the number of languages in mNMT improves the performance on XNLI.

3.2 MLDoc: Document Classification

MLDoc is a balanced subset of the Reuters corpus covering 8 languages for document classification (Schwenk and Li 2018). This is a 4-way classification task of identifying topics between CCAT (Corporate/Industrial), ECAT (Economics), GCAT (Government/Social), and MCAT (Markets). Performance is evaluated based on classification accuracy. We split the document using the sentence-piece model and feed the first 200 tokens into the encoder for classification. The task-specific network and the optimizer used is same as the one used for XNLI. Learning rate schedule is (0.2,5k). We perform both in-language and zero-shot evaluation. The in-language setting has training, development and test sets from the language. In the zero-shot setting, the train and dev sets contain only English examples but we test on all the languages. The results of both the experiments are reported in Table 3.

	In language		Zero-shot		
	mBERT*	MMTE	mBERT*	MMTE	SOTA†
en	94.2	94.7	94.2	94.7	89.9
de	93.3	93.4	80.2	77.4	84.8
zh	89.3	90.0	76.9	73.4	71.9
es	95.7	95.6	72.6	73.0	77.3
fr	93.4	92.7	72.6	77.2	78.0
it	88.0	87.6	68.9	64.2	69.4
ja	88.4	88.1	56.5	69.0	60.3
ru	87.5	87.4	73.7	68.9	67.8
avg	91.2	91.2	74.5	74.7	74.9

Table 3: Accuracies on the test set of the MLDoc dataset. *mBERT numbers have been taken from Wu and Dredze (2019). mBERT is also the state-of-the-art for in-language training. † For zero-shot SOTA refers to Artetxe and Schwenk (2018).

MMTE performance is on par with mBERT for in-language training on all the languages. It slightly edges over mBERT on zero-shot transfer while lagging behind SOTA by 0.2 points. Interestingly, MMTE beats SOTA on Japanese by more than 8 points. This may be due to the different nature and amount of data used for pre-training by these methods.

3.3 Cross-lingual Intent Classification

Schuster et al. (2018) recently presented a dataset for multi-lingual task oriented dialog. This dataset contains 57k annotated utterances in English (43k), Spanish (8.6k), and Thai (5k) with 12 different intents across the domains weather, alarm, and reminder. The evaluation metric used is classification accuracy. We use this data for both in-language training and zero-shot transfer. The task-specific network and the optimizer used is the same as the one used for the above two tasks. The learning rate schedule is (0.1,100k). Results are reported in Table 4. MMTE outperforms both mBERT and previous SOTA in both in-language and zero-shot setting on all 3 languages and establishes a new SOTA for this dataset.

	In language			Zero-shot		
	mBERT*	MMTE	SOTA†	mBERT*	MMTE	SOTA†
en	99.3	99.4	99.1	99.3	99.4	99.1
es	98.4	98.8	98.6	69.2	93.6	85.4
th	97.0	97.6	97.4	43.4	89.6	95.9
avg	98.2	98.6	98.4	70.6	94.2	93.5

Table 4: Accuracies on the test set of the intent classification dataset. *mBERT numbers are from our own implementation using the publicly available mBERT checkpoint. † SOTA refers to the numbers reported in Schuster et al. (2018).

3.4 POS Tagging

We use universal dependencies POS tagging data from the Universal Dependency v2.3 (Nivre et al. 2018; Zeman et al. 2018). Gold segmentation is used for training, tuning and testing. The POS tagging task has 17 labels for all languages. We consider 48 different languages. These languages are chosen based on intersection of languages for which POS labels are available in the universal dependencies dataset and the languages supported by our mNMT model. The task-specific network consists of a one layer feed-forward neural network with 784 units. Since MMTE operates on the subword-level, we only consider the representation of the first subword token of each word. The optimizer used is Adafactor with learning rate schedule (0.1,40k). The evaluation metric used is F1-score, which is same as accuracy in our case since we use gold-segmented data. Results of both in-language and zero-shot setting are reported in Table 5.

While mBERT outperforms MMTE on in-language training by a small margin of 0.16 points, MMTE beats mBERT by nearly 0.6 points in the zero-shot setting. Similar to results in XNLI, we see MMTE outperform mBERT on low resource languages. Since mBERT is SOTA for zero-shot cross-lingual transfer on POS tagging task (Wu and Dredze 2019), we also establish state-of-the-art on this dataset by beating mBERT in this setting.

3.5 Named Entity Recognition

For NER, we use the dataset from the CoNLL 2002 and 2003 NER shared tasks, which when combined have 4 languages (Sang 2002; Sang and De Meulder 2003). The label-

	In language		Zero-shot	
	mBERT*	MMTE	mBERT*	MMTE
ar	97.25	96.72	58.67	66.87
bg	98.86	98.93	86.65	87.61
de	96.21	95.92	91.63	89.22
en	96.11	96.34	96.16	96.34
es	97.70	97.79	87.56	85.11
fr	97.87	97.83	89.05	86.22
hi	97.39	97.53	72.00	74.48
mr	84.15	85.10	60.23	68.40
nl	97.82	97.27	90.30	88.40
pt	98.04	97.74	90.33	88.60
ro	97.52	97.43	79.44	78.80
sv	97.35	96.92	91.60	89.95
te	93.66	93.73	81.12	83.40
ur	94.43	94.42	65.88	69.36
zh	96.25	95.87	67.38	67.51
48L	95.65	95.49	74.87	75.43

Table 5: F1 scores on the dev set of the POS tagging dataset. *mBERT numbers are from our own implementation using the publicly available mBERT checkpoint. For readability, we only report results on 15 select languages and the 48 Language average (48L).

ing scheme is IOB with 4 types of named entities. The task-specific network, optimizer, and the learning rate schedule is the same as in the setup for POS tagging. The evaluation metric is span-based F1. Table 6 reports the results of both in-language and zero-shot settings.

	In language		Zero-shot	
	mBERT*	MMTE	mBERT*	MMTE
en	92.0	87.8	92.0	87.8
es	87.4	78.6	75.0	52.0
de	82.8	78.7	69.6	51.8
nl	90.9	81.7	77.6	50.8
avg	88.3	81.7	78.5	60.6

Table 6: F1 scores on the test set of the NER task. *mBERT numbers have been taken from Wu and Dredze (2019). mBERT is also the state-of-the-art for zero-shot setting.

MMTE performs significantly worse than mBERT on the NER task in all languages. On average, mBERT beats MMTE by 7 F1 points in the in-language setting and by more than 18 points in the zero-shot setting. We hypothesize that this might be because of two reasons: (1) mBERT is trained on clean Wikipedia data which is entity-rich while MMTE is trained on noisy web data with fewer entities, and (2) the translation task just copies the entities from the source to the target and therefore might not be able to accurately recognize them. This result points to the importance of the type of pre-training data and objective on down-stream

task performance. We plan to investigate this further in future work.

4 Analysis

In this section, we consider some additional settings for comparing mBERT and MMTE. We also investigate the impact of the number of languages and the target language token on MMTE performance.

Feature-based Approach In this setting, instead of fine-tuning the entire network of mBERT or MMTE, we only fine-tune the task-specific network which only has a small percentage of the total number of parameters. The rest of the model parameters are frozen. We perform this experiment on POS tagging task by fine-tuning a single layer feed-forward neural network stacked on top of mBERT and MMTE. We report the results in Table 7. While the scores of the feature-based approach are significantly lower than those obtained via full fine-tuning (5), we see that MMTE still outperforms mBERT on both in-language and zero-shot settings by an even bigger margin. This is particularly interesting as the feature-based approach has its own advantages: 1) it is applicable to downstream tasks which require significant task-specific parameters on top of a transformer encoder, 2) it is computationally cheaper to train and tune the downstream model, and 3) it is compact and scalable since we only need a small number of task-specific parameters.

	In language		Zero-shot	
	mBERT	MMTE	mBERT	MMTE
ar	87.65	91.50	61.08	60.35
bg	95.82	96.72	88.08	83.48
de	90.68	93.34	85.10	89.52
en	92.73	92.15	92.73	92.33
es	93.03	92.45	76.95	83.66
fr	92.20	93.41	77.64	83.58
hi	89.86	92.15	64.21	73.29
mr	83.50	84.49	52.27	66.36
nl	92.87	94.00	86.38	87.32
pt	93.14	93.84	85.95	86.86
ro	90.76	89.68	71.92	77.77
sv	93.67	94.35	85.98	87.82
te	90.57	92.75	70.09	81.12
ur	90.02	90.32	58.98	67.90
zh	86.29	84.34	60.82	56.50
avg	90.85	91.70	74.55	78.52

Table 7: F1 scores on dev set of POS tagging dataset using both mBERT and MMTE as feature extractor.

Few Shot Transfer While zero-shot transfer is a good measure of a model’s natural cross-lingual effectiveness, the more practical setting is the few-shot transfer scenario as we almost always have access to, or can cheaply acquire, a small

amount of data in the target language. We report the few-shot transfer results of mBERT and MMTE on the POS tagging dataset in 8. To simulate the few-shot setting, in addition to using English data, we use 10 examples from each language (upsampled to 1000). MMTE outperforms mBERT in few-shot setting by 0.6 points averaged over 48 languages. Once again, we see that the gains are more pronounced in low resource languages.

	mBERT	MMTE		mBERT	MMTE
ar	85.20	85.53	nl	90.01	90.26
bg	93.27	92.13	pt	92.18	92.29
de	92.24	91.57	ro	84.58	84.68
en	96.07	96.29	sv	91.61	92.11
es	92.52	92.12	te	83.84	84.74
fr	92.54	91.12	ur	81.38	82.94
hi	86.29	87.32	zh	81.14	76.32
mr	66.59	79.32	48L	83.97	84.58

Table 8: F1 scores on dev set of POS tagging dataset in few-shot setting using 10 examples from each language in addition to English data.

One Model for all Languages Another setting of importance is the in-language training where instead of training one model for each language, we concatenate all the data and train one model jointly on all languages. We perform this experiment on the POS tagging dataset with 48 languages and report results in Table 9. We observe that MMTE performance is on par with mBERT. We also find that the 48 language average improves by 0.2 points as compared to the one model per language setting in Table 5.

	mBERT	MMTE		mBERT	MMTE
ar	97.09	96.92	nl	97.60	97.59
bg	98.83	98.94	pt	98.42	98.02
de	95.94	95.68	ro	97.48	97.47
en	96.04	96.28	sv	97.27	97.22
es	97.71	97.77	te	94.86	94.81
fr	97.74	97.86	ur	94.39	94.77
hi	97.34	97.62	zh	96.01	95.84
mr	84.70	86.13	48L	95.67	95.69

Table 9: F1 scores on dev set of POS tagging dataset in one model for all language setting.

Number of Languages in mNMT We perform an ablation where we vary the number of languages used in the pre-training step. Apart from the 103 language setting, we consider 2 additional settings: 1) where we train mNMT on 4 languages to and from English, and 2) where we use 25 languages. The results are presented in Table 10. We see that as we scale up the languages the zero-shot performance goes

down on both POS tagging and XNLI tasks. These losses align with the relative BLEU scores of these models suggesting that the regressions are due to interference arising from the large number of languages attenuating the capacity of the NMT model. Scaling up the mNMT model to include more languages without diminishing cross-lingual effectiveness is a direction for future work.

	POS			XNLI		
	4x4	25x25	102x102	4x4	25x25	102x102
en	96.30	96.25	96.28	80.1	80.0	79.6
es	86.61	—	85.13	74.2	—	71.6
fr	86.62	85.81	85.23	72.9	72.1	69.5
de	90.52	89.84	88.57	71.4	71.1	68.2
zh	—	67.83	67.51	—	72.7	69.2

Table 10: Effect of number of languages used for mNMT training on downstream zero-shot performance. XNLI numbers are accuracies. POS numbers are F1 scores. — indicate that the language was not present in the subset that mNMT was trained on.

Effect of the Target Language Token During the pre-training step, when we perform the translation task using the mNMT system, we prepend a $\langle 2xx \rangle$ token to the source sentence, where xx indicates the target language. The encoder therefore has always seen a $\langle 2en \rangle$ token in front of non-English sentences and variety of different tokens depending on the target language in front of English sentence. However, when fine-tuning on downstream tasks, we do not use this token. We believe this creates a mismatch between the pre-training and fine-tuning steps. To investigate this further, we perform a small scale study where we train an mNMT model on 4 languages to and from English in two different settings: 1) where we prepend the $\langle 2xx \rangle$ token, and 2) where we don't prepend the $\langle 2xx \rangle$ token but instead encode it separately. The decoder jointly attends over both the source sentence encoder and the $\langle 2xx \rangle$ token encoding. The BLEU scores on the translation tasks are comparable using both these approaches.

	POS		XNLI	
	with $\langle 2xx \rangle$	without	with $\langle 2xx \rangle$	without
en	96.3	96.4	80.1	80.5
es	85.1	86.6	74.2	74.5
fr	86.6	87.8	72.9	73.7
de	90.5	91.1	71.4	72.9
avg	89.6	90.5	74.6	75.4

Table 11: Effect of $\langle 2xx \rangle$ token on zero-shot cross-lingual performance. XNLI numbers are accuracies. POS numbers are F1 scores.

The results on cross-lingual zero-shot transfer in both settings are provided in Table 11. Removing the $\langle 2xx \rangle$ to-

ken from the source sentence during mNMT training improves cross-lingual effectiveness on both POS tagging and XNLI task. Training a massively multilingual NMT model that supports translation of 102 languages to and from English without using the $\langle 2xx \rangle$ token in the encoder is another direction for future work.

5 Related Work

We briefly review widely used approaches in cross-lingual transfer learning and some of the recent work in learning contextual word representations (CWR).

Multilingual Word Embeddings For cross-lingual transfer, the most widely studied approach is to use multilingual word embeddings as features in neural network models. Several recent efforts have explored methods that align vector spaces for words in different languages (Faruqui and Dyer 2014; Upadhyay et al. 2016; Ruder, Vulić, and Søgaard 2017).

Unsupervised CWR More recent work has shown that CWRs obtained using unsupervised generative pre-training techniques such as language modeling or cloze task (Taylor 1953) have led to state-of-the-art results beyond what was achieved with traditional word type representations on many monolingual NLP tasks (Peters et al. 2018; Devlin et al. 2018; Howard and Ruder 2018; Radford et al.) such as sentence classification, sequence tagging, and question answering. Subsequently, these contextual methods have been extended to produce multilingual representations by training a single model on text from multiple languages which have proven to be very effective for cross lingual transfer (Wu and Dredze 2019; Mulcaire, Kasai, and Smith 2019; Pires, Schlinger, and Garette 2019). Lample and Conneau (2019) show that adding a translation language modeling (TLM) objective to mBERT’s MLM objective utilizes both monolingual and parallel data to further improve the cross-lingual effectiveness.

Representations from NMT The encoder from an NMT model has been used as yet another effective way to contextualize word vectors (McCann et al. 2017). Additionally, recent progress in NMT has enabled one to train multilingual NMT systems that support translation from multiple source languages into multiple target languages within a single model (Johnson et al. 2017). Our work is more closely related to two very recent works which explore the encoder from multilingual NMT model for cross-lingual transfer learning (Eriguchi et al. 2018; Artetxe and Schwenk 2018). While Eriguchi et al. (2018) also consider multilingual systems, they do so on a much smaller scale, training it on only 2 languages. Artetxe and Schwenk (2018) uses a large scale model comparable to ours with 93 languages but they constrain the model by pooling encoder representations and therefore only obtain a single vector per sequence. Neither of these approaches have been used on token level sequence tagging tasks. Further, neither concern themselves with the performance of the actual translation task whereas we our

mNMT model performs comparable to bilingual baselines in terms of translation quality.

6 Conclusion and Future Work

We train a massively multilingual NMT system using parallel data from 103 languages and exploit representations extracted from the encoder for cross-lingual transfer on various classification and sequence tagging tasks spanning over 50 languages. We find that the positive language transfer visible in improved translation quality for low resource languages is also reflected in the cross-lingual transferability of the extracted representations. The gains observed on various tasks over mBERT suggest that the translation objective is competitive with specialized approaches to learn cross-lingual embeddings.

We find that there is a trade off between the number of languages in the multilingual model and efficiency of the learned representations due to the limited capacity. Scaling up the model to include more languages without diminishing transfer learning capability is a direction for future work. Finally, one could also consider integrating mBERT’s objective with the translation objective to pre-train the mNMT system.

References

- Aharoni, R.; Johnson, M.; and Firat, O. 2019. Massively multilingual neural machine translation. *arXiv preprint arXiv:1903.00089*.
- Artetxe, M., and Schwenk, H. 2018. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *arXiv preprint arXiv:1812.10464*.
- Chen, M. X.; Firat, O.; Bapna, A.; Johnson, M.; Macherey, W.; Foster, G.; Jones, L.; Schuster, M.; Shazeer, N.; Parmar, N.; Vaswani, A.; Uszkoreit, J.; Kaiser, L.; Chen, Z.; Wu, Y.; and Hughes, M. 2018. The best of both worlds: Combining recent advances in neural machine translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 76–86. Melbourne, Australia: Association for Computational Linguistics.
- Conneau, A.; Rinott, R.; Lample, G.; Williams, A.; Bowman, S. R.; Schwenk, H.; and Stoyanov, V. 2018. Xnli: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Eriguchi, A.; Johnson, M.; Firat, O.; Kazawa, H.; and Macherey, W. 2018. Zero-shot cross-lingual classification using multilingual neural machine translation. *arXiv preprint arXiv:1809.04686*.
- Faruqui, M., and Dyer, C. 2014. Improving vector space word representations using multilingual correlation. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, 462–471.

- Firat, O.; Cho, K.; and Bengio, Y. 2016. Multi-way, multilingual neural machine translation with a shared attention mechanism. *arXiv preprint arXiv:1601.01073*.
- Howard, J., and Ruder, S. 2018. Universal language model fine-tuning for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 328–339.
- Johnson, M.; Schuster, M.; Le, Q. V.; Krikun, M.; Wu, Y.; Chen, Z.; Thorat, N.; Viégas, F.; Wattenberg, M.; Corrado, G.; et al. 2017. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association of Computational Linguistics* 5(1):339–351.
- Kudo, T., and Richardson, J. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 66–71.
- Lample, G., and Conneau, A. 2019. Cross-lingual language model pretraining. *arXiv preprint arXiv:1901.07291*.
- McCann, B.; Bradbury, J.; Xiong, C.; and Socher, R. 2017. Learned in translation: Contextualized word vectors. In *Advances in Neural Information Processing Systems*, 6294–6305.
- Mulcaire, P.; Kasai, J.; and Smith, N. 2019. Polyglot contextual representations improve crosslingual transfer. *arXiv preprint arXiv:1902.09697*.
- Nivre, J.; Abrams, M.; Agić, Ž.; Ahrenberg, L.; Antonsen, L.; Aranzabe, M. J.; Arutie, G.; Asahara, M.; Ateyah, L.; Attia, M.; et al. 2018. Universal dependencies 2.2.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, 311–318. Association for Computational Linguistics.
- Peters, M.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 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)*, 2227–2237.
- Pires, T.; Schlinger, E.; and Garette, D. 2019. How multilingual is multilingual bert? In *Proceedings of Association of Computational Linguistics (ACL)*.
- Radford, A.; Narasimhan, K.; Salimans, T.; and Sutskever, I. Improving language understanding by generative pre-training.
- Ruder, S.; Vulić, I.; and Søgaard, A. 2017. A survey of cross-lingual word embedding models. *arXiv preprint arXiv:1706.04902*.
- Sang, E. F. T. K., and De Meulder, F. 2003. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*.
- Sang, E. F. T. K. 2002. Introduction to the conll-2002 shared task: Language-independent named entity recognition. *CoRR* cs.CL/0209010.
- Schuster, S.; Gupta, S.; Shah, R.; and Lewis, M. 2018. Cross-lingual transfer learning for multilingual task oriented dialog. *arXiv preprint arXiv:1810.13327*.
- Schwenk, H., and Li, X. 2018. A corpus for multilingual document classification in eight languages. In chair), N. C. C.; Choukri, K.; Cieri, C.; Declerck, T.; Goggi, S.; Hasida, K.; Isahara, H.; Maegaard, B.; Mariani, J.; Mazo, H.; Moreno, A.; Odijk, J.; Piperidis, S.; and Tokunaga, T., eds., *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. Paris, France: European Language Resources Association (ELRA).
- Shazeer, N., and Stern, M. 2018. Adafactor: Adaptive learning rates with sublinear memory cost. *arXiv preprint arXiv:1804.04235*.
- Shen, J.; Nguyen, P.; Wu, Y.; and et al., Z. C. 2019. Lingvo: a modular and scalable framework for sequence-to-sequence modeling. *arXiv preprint arXiv:1902.08295*.
- Taylor, W. L. 1953. “cloze procedure”: A new tool for measuring readability. *Journalism Bulletin* 30(4):415–433.
- Upadhyay, S.; Faruqui, M.; Dyer, C.; and Roth, D. 2016. Cross-lingual models of word embeddings: An empirical comparison. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 1661–1670.
- Uszkoreit, J.; Ponte, J. M.; Papat, A. C.; and Dubiner, M. 2010. Large scale parallel document mining for machine translation. In *Proceedings of the 23rd International Conference on Computational Linguistics*, 1101–1109. Association for Computational Linguistics.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, 5998–6008.
- Wu, S., and Dredze, M. 2019. Beto, bentz, becas: The surprising cross-lingual effectiveness of bert. *arXiv preprint arXiv:1904.09077*.
- Wu, Y.; Schuster, M.; Chen, Z.; Le, Q. V.; Norouzi, M.; Macherey, W.; Krikun, M.; Cao, Y.; Gao, Q.; Macherey, K.; et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Zeman, D.; Hajič, J.; Popel, M.; Potthast, M.; Straka, M.; Ginter, F.; Nivre, J.; and Petrov, S. 2018. Conll 2018 shared task: Multilingual parsing from raw text to universal dependencies. In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, 1–21.