

# Ensembling of Distilled Models from Multi-task Teachers for Constrained Resource Language Pairs

Amr Hendy<sup>1</sup>, Esraa A. Gad<sup>1</sup>, Mohamed Abdelghaffar<sup>1</sup>, Jailan S. ElMosalami<sup>1</sup>,  
Mohamed Afify<sup>1</sup>, Ahmed Y. Tawfik<sup>1</sup> and Hany Hassan Awadalla<sup>2</sup>

<sup>1</sup> Microsoft Egypt Development Center, Cairo, Egypt

<sup>2</sup> Microsoft Corporation, Redmond, WA, USA

{amrhendy, v-egad, mohamed.abdelghaar, v-jailanel}@microsoft.com

{mafify, atawfik, hanyh}@microsoft.com

## Abstract

This paper describes our submission to the constrained track of WMT21 shared news translation task. We focus on the three relatively low resource language pairs Bengali  $\leftrightarrow$  Hindi, English  $\leftrightarrow$  Hausa and Xhosa  $\leftrightarrow$  Zulu. To overcome the limitation of relatively low parallel data we train a multilingual model using a multitask objective employing both parallel and monolingual data. In addition, we augment the data using back translation. We also train a bilingual model incorporating back translation and knowledge distillation then combine the two models using sequence-to-sequence mapping. We see around 70% relative gain in BLEU point for  $En \leftrightarrow Ha$  and around 25% relative improvements for  $Bn \leftrightarrow Hi$  and  $Xh \leftrightarrow Zu$  compared to bilingual baselines.

## 1 Introduction

Neural machine translation (NMT) witnessed a lot of success in the past few years especially for high resource languages (Vaswani et al., 2017). Improving the quality of low resource languages is still challenging. Some of the popular techniques are adding high resource helper languages as in multilingual neural machine translation (MNMT) (Dong et al., 2015; Firat et al., 2016; Ha et al., 2016; Johnson et al., 2017; Arivazhagan et al., 2019), using monolingual data including pre-training (Liu et al., 2020), multi-task learning (Wang et al., 2020), back translation (Sennrich et al., 2016) or any combination of these methods (Barrault et al., 2020) and system combination of multiple systems (Liu et al., 2018).

This paper describes the Microsoft Egypt Development Center (EgDC) submission to the WMT21 shared news translation task for three low resource language pairs (six directions), Bengali  $\leftrightarrow$  Hindi ( $Bn \leftrightarrow Hi$ ), English  $\leftrightarrow$  Hausa ( $En \leftrightarrow Ha$ ) and Xhosa  $\leftrightarrow$  Zulu ( $Xh \leftrightarrow Zu$ ). We focus on the constrained track because it is easier to compare

different systems and it is always possible to improve performance by adding more data. The main features of our approach are as follows:

- Using a recently proposed multitask and multilingual learning framework to benefit from monolingual data in both the source and target languages (Wang et al., 2020).
- Using knowledge distillation (Freitag et al., 2017) to create bilingual baselines from the original multilingual model and combining it with the multilingual model.

The paper is organized as follows. Section 2 gives an overview of the data used in the constrained scenario, followed by section 3 that gives a detailed description of our approach. Section 4 presents our experimental evaluation. Finally, our findings are summarized in Section 5.

## 2 Data

Following the constrained track, we use bitext data provided in WMT21 for the following pairs: Bengali  $\leftrightarrow$  Hindi, English  $\leftrightarrow$  Hausa, Xhosa  $\leftrightarrow$  Zulu and English  $\leftrightarrow$  German. Statistics of the parallel data used for the three pairs in addition to the German helper are shown in Table 1. We also use monolingual data for all previously mentioned languages provided in WMT21 for techniques such as multi-task training and back-translation. Statistics of the monolingual data used for the 6 languages in addition to the German helper are shown in Table 2. For very low resource languages, Hausa, Xhosa and Zulu, we use all the available monolingual data, e.g. NewsCrawl + CommonCrawl + Extended CommonCrawl for Hausa, and Extended CommonCrawl for both Xhosa and Zulu. For relatively high resource languages, Bengali, Hindi, English and German, we only use a subset of the provided data mostly from NewsCrawl due to its high-quality. In addition to the NewsCrawl monolingual subset, we add a sampled subset from CommonCrawl to

Language pair	# of sentences
Bengali ↔ Hindi	3.36M
English ↔ Hausa	750K
Xhosa ↔ Zulu	94K
English ↔ German	84.8M

Table 1: Bitext data used for bilingual and multilingual systems. For each language pair, we use all available sources released in WMT21

Language	# of sentences	
	Raw	Cleaned
Bengali	53.8M	53.3M
English	75M	73.5M
German	111.2M	109.9M
Hausa	10.8M	6.2M
Hindi	60.2M	59.8M
Xhosa	1.6M	950K
Zulu	2M	1.4M

Table 2: Monolingual data used for multi-task training and back-translation

avoid biasing into the news domain especially for Bengali ↔ Hindi and Xhosa ↔ Zulu whose target evaluation domain come from Wikipedia content.

## 2.1 Data Filtering

For Bengali, English, Hindi and German, we apply fastText<sup>1</sup> language identification on the monolingual data to remove sentences which are not predicted as the expected language. We do the same for Hausa, Xhosa and Zulu using Polyglot<sup>2</sup> because fastText does not cover these three languages. The resulting size of the monolingual data of each language is shown in Table 2.

## 3 System Architecture

The final MT system in each direction is an ensemble of two NMT models comprising a bilingual model (one for each of the six primary directions) and a multilingual model trained to provide translations for 8 directions (the six primary directions plus English ↔ German). The multilingual system uses a recently proposed multitask framework for training (Wang et al., 2020). We describe the individual systems in Subsection 3.1. This is followed by presenting our system combination techniques in Subsection 3.2. Finally we present the architecture of the submitted system highlighting our

design decisions in Subsection 3.3.

### 3.1 Individual Systems

This subsection describes the individual systems and their training leading to the proposed system combination strategy in the following subsection. We first build bilingual models for the six primary directions using the data shown in Table 1 except the English ↔ German. These serve as baselines to compare to the developed systems. The models use a transformer base architecture comprising 6 encoder and 6 decoder layers and a 24K joint vocabulary built for Bengali ↔ Hindi, a 8K joint vocabulary built for English ↔ Hausa and a 4K joint vocabulary built for Xhosa ↔ Zulu using sentencepiece (Kudo and Richardson, 2018) to learn these subword units to tokenize the sentences. In addition to the baseline bilingual models, we use knowledge distilled (KD) data and back-translated (BT) data generated from a multilingual model to build another set of bilingual models for each of the six primary directions. This multilingual model is described below. The purpose of these models is to participate in the ensemble along with the multilingual models. The latter bilingual models follow the same transformer base architecture and joint vocabulary used in the baseline bilingual models.

The multilingual model combines the 8 translation directions shown in Table 1. These are the six primary directions plus English ↔ German as a helper. The latter is mainly used to improve generation on the English centric directions. The model uses a 64K joint vocabulary constructed using sentencepiece (Kudo and Richardson, 2018) from a subset of the monolingual data of each language as described in Section 2. The transformer model has 12 encoder and 6 decoder layers. In addition, a multitask objective is used during training to make use of monolingual data. The objective comprises the usual parallel data likelihood referred to as MT, a masked language model (MLM) at the encoder and a denoising auto-encoder (DAE) (similar to mBART (Liu et al., 2020)) at the decoder side. The latter two objectives help leverage monolingual data for both the encoder and the decoder sides. The three objectives are combined using different proportions according to a schedule during the training. Please refer to (Wang et al., 2020) for details.

<sup>1</sup><https://fasttext.cc/docs/en/language-identification.html>

<sup>2</sup><https://github.com/aboSamoor/polyglot>

To summarize we build the following models:

- Bilingual models trained using parallel data in Table 1 for the 6 primary directions. These are mainly used as baselines.
- Multilingual models trained using a multitask objective using parallel and monolingual data and comprising 8 directions.
- Bilingual models trained using KD and BT data generated using our best multilingual model. These are combined with the best multilingual model as described in 3.2.

### 3.2 System Combination

System combination or ensembling is known to improve the performance over individual systems. There are many ways to create an ensemble (Liu et al., 2018; Dabre et al., 2019). For example, individual models obtained from different checkpoints during the same training or by training models sharing the same vocab and architecture using different data or simply different random seeds can be combined using model averaging techniques. Here, we opt to combine different models since it generally leads to better performance because different models tend to be more complementary. To this end, we propose a simple and effective method to combine completely different architectures. The proposed method could be also used in conjunction with checkpoint and model averaging for further gains, but we haven’t tried this in our experiments due to time limitations.

The basic idea of our combination is very simple. Assume we have the translation pair  $x \rightarrow y$  where  $y$  is the reference translation. The output of model 1 is the pair  $x \rightarrow y1$  and the output of model 2 is the pair  $x \rightarrow y2$ . This can be generalized to multiple systems but we limited our combination to only two models. We train a new model that takes the set of hypotheses (possibly augmented by the source sentence) from the two models to generate the target sentence. Thus this model combines the outputs of two models in the ensemble to produce a translation closer to the original target sentence i.e.  $\langle HYP \rangle y1 \langle HYP \rangle y2 \rightarrow y$ . We also experimented with adding the source to the input i.e.  $\langle SRC \rangle x \langle HYP \rangle y1 \langle HYP \rangle y2 \rightarrow y$  which led to around 0.3 BLEU improvement for  $Ha \rightarrow En$ , but we haven’t tried on other pairs due to time limitation. All combination models use 6

layers encoder and decoder and a 64K vocabulary similar to the multilingual system. These combination models use the full bitext and dev data provided in WMT21 as shown in Table 1. The system combination is outlined in Figure 1. This ensembling technique can be thought of as providing both system combination and post-editing capabilities.

### 3.3 Overall System

Our overall system is depicted in Figure 2. The first module shows the data input where language identification (LID) is used to filter the monolingual data. As mentioned in Section 2.1 we use fastText and polyglot for LID depending on the language. We first build bilingual baselines which are not shown in the figure. Then as shown in the second module, we build 4 multilingual systems using different task objectives as follows:  $MT$ ,  $MT + MLM$ ,  $MT + DAE$  and  $MT + MLM + DAE$  trained on the 8 directions shown in Table 1 following the temperature-based strategy in (Arivazhagan et al., 2019) to balance the training data in different resource languages using  $T = 5$ . We pick the best system and use it to back translate the selected monolingual data. For most pairs, as detailed in Section 4, we find that  $MT + DAE$  and  $MT + MLM + DAE$  are quite close. Therefore, we use the  $MT + DAE$  to do back translation for all submitted 6 pairs. We use beam search with beam size = 5 when generating the synthetic back-translated data. Once we get the back-translated data (called  $BT_1$ ) we add it to our parallel and monolingual data and build a new multilingual model called  $MT + DAE + BT_1$ . We tag the back-translated data with  $\langle BT \rangle$  tag at beginning of each source sentence so the model can differentiate between the genuine parallel and back-translated data quality. The resulting model is used to regenerate the back-translated data (called  $BT_2$ ) and to knowledge distill the bitext (called  $KD$ ). The latter two data sets are augmented and used to build a bilingual system (called  $MT + KD + BT_2$ ). We upsample the  $KD$  data set and the upsampling ratio is selected based on parameter sweeping and validating the resulting improvement on the validation set. Finally, the latter bilingual model is combined with our final multilingual model using the method in Section 3.2 to create our submission.

## 4 Experimental Results

In this section, we describe the results of our intermediate and final systems. We report Sacre-




Figure 1: The system combination component used for our experiments.

BLEU (Post, 2018) on the validation set released in WMT21, and both SacreBLEU and COMET (Rei et al., 2020) using the available implementation<sup>3</sup> on the official test set released in WMT21. The results for the six submitted language pairs are in Tables 3-5. The first row in each table shows the bilingual baseline which performs relatively poor due to the limited amount of parallel data for each pair. This is followed by the four multilingual systems with different objectives. It is clear that adding a monolingual objective brings nice improvements for all language pairs. The *MT + DAE* and *MT + MLM + DAE* perform closely for all language pairs indicating that target monolingual data is most important. The next two rows show the results of adding back-translated data to the multilingual model and a bilingual baseline using back-translated and knowledge distilled data generated from the best multilingual model. As expected adding back translation brings significant improvement to all language pairs. Also using the multilingual model to create data for a bilingual model shows excellent results that outperform the multilingual model. Finally, the ensemble, as expected, performs better than the individual models. The significant difference between reported improvements in *Ha* ↔ *En* and other directions shows the effectiveness of adding *De* ↔ *En* parallel and monolingual data that helps English centric directions more than other directions. We evaluated the final submitted systems on the official test set released in WMT21 as shown in Table 6.

<sup>3</sup><https://github.com/Unbabel/COMET>

System	Ha-En	En-Ha
bilingual baseline	14.10	13.78
multi. MT	14.32	13.16
+ MLM	16.18	13.94
+ DAE	18.05	14.91
+ MLM + DAE	17.35	15.03
multi. MT + DAE + BT <sub>1</sub>	21.11	20.24
bilingual MT + KD + BT <sub>2</sub>	24.43	20.68
ensemble	24.90	21.00

Table 3: Results of Ha-En and En-Ha systems. We report SacreBLEU scores on the validation set provided in WMT21

System	Bn-Hi	Hi-Bn
bilingual baseline	18.60	10.90
multi. MT	18.21	10.02
+ MLM	18.82	10.67
+ DAE	18.64	10.40
+ MLM + DAE	19.20	11.27
multi. MT + DAE + BT <sub>1</sub>	20.18	12.29
bilingual MT + KD + BT <sub>2</sub>	21.03	12.90
ensemble	21.20	13.30

Table 4: Results of Bn-Hi and Hi-Bn systems. We report SacreBLEU scores on the validation set provided in WMT21

## 5 Summary

This paper describes our submission to the constrained track of WMT21. We focus on the three relatively low resource language pairs *Bn* ↔ *Hi*, *En* ↔ *Ha* and *Xh* ↔ *Zu*. To overcome the limitation of relatively low parallel data we train a multilingual model using a multitask objective recently proposed in (Wang et al., 2020). In addition,




Figure 2: The overall system flow used for our experiments

System	Xh-Zu	Zu-Xh
bilingual baseline	8.00	7.60
multi. MT	7.53	7.47
+ MLM	7.23	7.02
+ DAE	8.53	8.24
+ MLM + DAE	8.20	7.80
multi. MT + DAE + BT <sub>1</sub>	9.06	8.86
bilingual MT + KD + BT <sub>2</sub>	9.80	9.17
ensemble	10.00	9.30

Table 5: Results of Xh-Zu and Zu-Xh systems. We report SacreBLEU scores on the validation set provided in WMT21

Translation direction	BLEU	COMET
$Ha \rightarrow En$	17.13	0.149
$En \rightarrow Ha$	16.13	0.086
$Bn \rightarrow Hi$	21.08	0.532
$Hi \rightarrow Bn$	10.93	0.411
$Xh \rightarrow Zu$	9.94	0.180
$Zu \rightarrow Xh$	9.25	0.299

Table 6: Results of the submitted systems. We report SacreBLEU and COMET scores on the official test set provided in WMT21. For COMET, we use the recommended model “wmt20-comet-da”.

we augment the data using back translation. We also use the resulting multilingual model to create a

bilingual model incorporating back translation and knowledge distillation. Finally, we combine the two models, using a flexible sequence-to-sequence approach, to yield our submitted systems. We see large gains up to 8-10 BLEU points for  $En \leftrightarrow Ha$  and nice improvements of up to 2-3 BLEU points for  $Bn \leftrightarrow Hi$  and  $Xh \leftrightarrow Zu$ .

## References

- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019. [Massively multilingual neural machine translation in the wild: Findings and challenges](#).
- Loïc Barrault, Magdalena Biesialska, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Matthias Huck, Eric Joanis, Tom Kocmi, Philipp Koehn, Chi-kiu Lo, Nikola Ljubešić, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Santanu Pal, Matt Post, and Marcos Zampieri. 2020. [Findings of the 2020 conference on machine translation \(WMT20\)](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1–55, Online. Association for Computational Linguistics.

- Raj Dabre, Fabien Cromieres, and Sadao Kurohashi. 2019. [Enabling multi-source neural machine trans-](#)

- lation by concatenating source sentences in multiple languages.
- Daxiang Dong, Hua Wu, W. He, Dianhai Yu, and Haifeng Wang. 2015. Multi-task learning for multiple language translation. In *ACL*.
- Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016. Multi-way, multilingual neural machine translation with a shared attention mechanism. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 866–875, San Diego, California. Association for Computational Linguistics.
- Markus Freitag, Yaser Al-Onaizan, and Baskaran Sankaran. 2017. Ensemble distillation for neural machine translation.
- Thanh-Le Ha, Jan Niehues, and Alexander Waibel. 2016. Toward multilingual neural machine translation with universal encoder and decoder.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. Google’s multilingual neural machine translation system: Enabling zero-shot translation.
- Taku Kudo and John Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. Multilingual denoising pre-training for neural machine translation.
- Yuchen Liu, Long Zhou, Yining Wang, Yang Zhao, Jiajun Zhang, and Chengqing Zong. 2018. A comparable study on model averaging, ensembling and reranking in nmt. In *Natural Language Processing and Chinese Computing*, pages 299–308, Cham. Springer International Publishing.
- Matt Post. 2018. A call for clarity in reporting bleu scores.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. COMET: A neural framework for MT evaluation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- 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, Berlin, Germany. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need.
- Yiren Wang, ChengXiang Zhai, and Hany Hassan. 2020. Multi-task learning for multilingual neural machine translation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1022–1034, Online. Association for Computational Linguistics.