

The USYD-JD Speech Translation System for IWSLT2021

Liang Ding

The University of Sydney
ldin3097@sydney.edu.au

Di Wu*

Peking University
inbath@163.com

Dacheng Tao

JD Explore Academy, JD.com
dacheng.tao@gmail.com

Abstract

This paper describes the University of Sydney & JD’s joint submission of the IWSLT 2021 low resource speech translation task. We participated in the Swahili→English direction and got the best scoreBLEU (25.3) score among all the participants. Our constrained system is based on a pipeline framework, i.e. ASR and NMT. We trained our models with the officially provided ASR and MT datasets. The ASR system is based on the open-sourced tool Kaldi and this work mainly explores how to make the most of the NMT models. To reduce the punctuation errors generated by ASR model, we employ our previous work SlotRefine to train a punctuation correction model. To achieve better translation performance, we explored the most recent effective strategies, including back translation, knowledge distillation, multi-feature reranking and transductive finetuning. For model structure, we tried autoregressive and non-autoregressive models, respectively. In addition, we proposed two novel pre-train approaches, i.e. *de-noising training* and *bidirectional training* to fully exploit the data. Extensive experiments show that adding the above techniques consistently improves the BLEU scores, and the final submission system outperforms the baseline (Transformer ensemble model trained with the original parallel data) by approximately 10.8 BLEU score, achieving the SOTA performance.

1 Introduction

Recent years have seen a surge of interest in speech translation (ST, Ney 1999) task, that translates the source-side speech to the target-side text directly. The ST task contains two major components, Automatic Speech Recognition (ASR, Jelinek 1997) and Machine Translation (MT, Koehn 2009). In this year’s IWSLT low-resource speech translation

task, our USYD-JD translation team participated in the Swahili to English track. We break the speech translation task into “ASR→NMT” pipeline, and mainly focus on the NMT component.

For model frameworks, we tried autoregressive neural machine translation, including Transformer-BASE and -BIG (Vaswani et al., 2017), and non-autoregressive translation models (Gu et al., 2018). Also, we employ our previous work SlotRefine (Wu et al., 2020a) to tackle the case and punctuation problems after ASR. To make the most of the parallel and monolingual data, we proposed two pre-train strategies, i.e. BIDIRECTIONAL PRETRAINING §2.2 and DENOISING PRETRAINING §2.3, and employed two data augmentation strategies, i.e. BIDIRECTIONAL SELF-TRAINING §2.5 and TAGGED BACK TRANSLATION §2.7. Where the data used for tagged back translation are carefully selected with our proposed multi-feature in-domain selection approach in §2.6. For post finetune/ process, we employed TRANSDUCTIVE FINE-TUNE §2.8 and a simple postprocessing approach §2.10.

This paper is structured as follows: Section 2 describes the major approaches we used. We present the data descriptions in Section 3. The experiments settings and main results are shown in Section 4. Finally, we conclude our work in Section 5.

2 Approaches

2.1 Autoregressive Translation

Given a source sentence \mathbf{x} , an NMT model generates each target word \mathbf{y}_t conditioned on previously generated ones $\mathbf{y}_{<t}$. Accordingly, the probability of generating \mathbf{y} is computed as:

$$p(\mathbf{y}|\mathbf{x}) = \prod_{t=1}^T p(\mathbf{y}_t|\mathbf{x}, \mathbf{y}_{<t}; \theta) \quad (1)$$

Work was done when Di Wu was visiting at JD.

where T is the length of the target sequence and the parameters θ are trained to maximize the likelihood of a set of training examples according to $\mathcal{L}(\theta) = \arg \max_{\theta} \log p(\mathbf{y}|\mathbf{x}; \theta)$. Typically, we choose Transformer (Vaswani et al., 2017) as its SOTA performance. The training examples can be formally defined as follows:

$$\vec{\mathbf{B}} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N \quad (2)$$

where N is the total number of sentence pairs in the training data. Note that in standard MT training, the \mathbf{x} is feed to the encoder and $\mathbf{y}_{<t}$ to the decoder to finish the conditional estimation for \mathbf{y}_t , thus the utilization of $\vec{\mathbf{B}}$ is directional, i.e. $\mathbf{x}_i \rightarrow \mathbf{y}_i$. In the preliminary experiments, we utilized autoregressive translation (AT) model for *translation*, *case correction* and *punctuation generation* tasks as its powerful modelling ability and generation accuracy.

2.2 Bidirectional Pretraining

Motivation The motivation is when human learn foreign languages with translation examples, e.g. \mathbf{x}_i and \mathbf{y}_i . Both directions of this example, i.e. $\mathbf{x}_i \rightarrow \mathbf{y}_i$ and $\mathbf{y}_i \rightarrow \mathbf{x}_i$, may help human easily master the bilingual knowledge. Motivated by this, Levinboim et al. (2015); Liang et al. (2007) propose to modelling the invertibility between bilingual languages. Cohn et al. (2016) introduce extra bidirectional prior regularization to achieve symmetric training from the point view of training objective. He et al. (2018); Zheng et al. (2019) enhance the coordination of bidirectional corpus with model level modifications. Different from the above methods, we model both directions of a given training example by a simple data manipulation strategy.

Our Implementation Many studies have shown that pretraining could transfer the knowledge and data distribution, hence improving the generalization (Hendrycks et al., 2019; Mathis et al., 2021). Here we want to transfer the bidirectional knowledge among the corpus. Specifically, we propose to first pretrain MT models on bidirectional corpus, which can be defined as follows:

$$\overleftrightarrow{\mathbf{B}} = \{(\mathbf{x}_i, \mathbf{y}_i) \cup (\mathbf{y}_i, \mathbf{x}_i)\}_{i=1}^N \quad (3)$$

such that the θ in Equation 1 can be updated by both directions, then the bidirectional pretraining

(BiPT) objective can be formulated as:

$$\mathcal{L}_{\text{BiPT}}(\theta) = \overbrace{\arg \max_{\theta} \log p(\mathbf{y}|\mathbf{x}; \theta)}^{\text{Forward: } \vec{\mathcal{L}}_{\theta}} \quad (4)$$

$$+ \underbrace{\arg \max_{\theta} \log p(\mathbf{x}|\mathbf{y}; \theta)}_{\text{Backward: } \overleftarrow{\mathcal{L}}_{\theta}} \quad (5)$$

where the forward $\vec{\mathcal{L}}_{\theta}$ and backward $\overleftarrow{\mathcal{L}}_{\theta}$ are optimized iteratively. From data perspective, we achieve the bidirectional updating as follows: 1) swapping the source and target sentences of a parallel corpus, and 2) appending the swapped version to the original. Then the training data was doubled to make better and full use of the costly bilingual corpus. The pretraining can acquire general knowledge from bidirectional data, which may help *better* and *faster* learning further tasks. Thus, we early stop bidirectional training at 1/3 of the total steps. To ensure the proper training direction, we further train the pretrained model on required direction $\vec{\mathbf{B}}$ with the rest of 2/3 training steps. Considering the effectiveness of pretraining (Mathis et al., 2021) and clean finetuning (Wu et al., 2019), we introduce a combined pipeline: $\overleftrightarrow{\mathbf{B}} \rightarrow \vec{\mathbf{B}}$ as our best training strategy.

2.3 Denoising Pretraining

Motivation The motivation is when human learn one language, one of the best practices for language acquisition is to correct the sentence errors, e.g. $\text{noised}(\mathbf{x}_i) \rightarrow \mathbf{x}_i$ and $\text{noised}(\mathbf{y}_i) \rightarrow \mathbf{y}_i$. Motivated by this, Lewis et al. (2020) propose several noise adding approaches and denoise them with end-to-end pretraining. Liu et al. (2020b) introduce this idea to the multilingual scenarios. Different from above monolingual denoising pretraining approaches, we proposed a simpler noise function and apply them to each side of the parallel data.

Our Implementation Here we want the model to understand the source- and target-side languages well. For noise function $\text{noised}(\cdot)$, we apply the common noise-injection practice, i.e. removing, replacing, or nearby swapping one time for a random word with a uniform distribution in a sentence (Edunov et al., 2018; Ding et al., 2020a). Then the size of the original parallel data doubled as follows:

$$\mathbf{S}_{\text{src}} = \{\text{noised}(\mathbf{x}_i), \mathbf{x}_i\}_{i=1}^N \quad (6)$$

$$\mathbf{S}_{\text{tgt}} = \{\text{noised}(\mathbf{y}_i), \mathbf{y}_i\}_{i=1}^N \quad (7)$$

where S_{src} and S_{tgt} can be combined to update the end-to-end model to achieve denoising pretraining, such that the θ in Equation 1 can be updated by denoising both the source and target data, then the denoising pretraining (DPT) objective can be formulated as:

$$\mathcal{L}_{\text{DPT}}(\theta) = \underbrace{\arg \max_{\theta} \log p(\mathbf{x} | \text{noised}(\mathbf{x}); \theta)}_{\text{Source Denoising: } \mathcal{L}_{\theta}^S} + \underbrace{\arg \max_{\theta} \log p(\mathbf{y} | \text{noised}(\mathbf{y}); \theta)}_{\text{Target Denoising: } \mathcal{L}_{\theta}^T} \quad (8)$$

$$(9)$$

where the Source Denoising : \mathcal{L}_{θ}^S and Target Denoising : \mathcal{L}_{θ}^T are optimized iteratively. The pretraining can store knowledge of the source and target languages into the shared model parameters, which may help *better* and *faster* learning further tasks. Similar to bidirectional pretraining in §2.2, we early stop denoising training at 1/3 of the total steps, and tune the model normally with the rest of 2/3 training steps. This process can be formally denoted as such pipeline: $S_{\text{src}} + S_{\text{tgt}} \rightarrow \vec{B}$.

Note that Bidirectional Pretraining (BiPT) and Denoising Pretraining (DPT) can be combined and further enhance the model performance (The effect of their complementary can be found in Table 7). In particular, the combination order of BiPT and DPT are empirically inspired by human learning behavior, where a good interpreter will first master at least one language (usually the mother tongue), and then learn other languages and achieve bilingual translation. Thus, the combined pretraining process follows $\text{DPT} \rightarrow \text{BiPT}$. In combined pretraining setting, we will train longer until the model converges completely.

2.4 Nonautoregressive Translation

Different from autoregressive translation (Bahdanau et al., 2015; Vaswani et al., 2017, AT) models that generate each target word conditioned on previously generated ones, non-autoregressive translation (Gu et al., 2018, NAT) models break the autoregressive factorization and produce the target words in parallel. Given a source sentence \mathbf{x} , the probability of generating its target sentence \mathbf{y} with length T is defined by NAT as:

$$p(\mathbf{y} | \mathbf{x}) = p_L(T | \mathbf{x}; \theta) \prod_{t=1}^T p(\mathbf{y}_t | \mathbf{x}; \theta) \quad (10)$$

where $p_L(\cdot)$ is a separate conditional distribution to predict the length of target sequence. Typically, most NAT models are implemented upon the framework of Transformer (Vaswani et al., 2017). In the preliminary experiments, we utilized NAT for *translation*, *case correction* and *punctuation generation* tasks as NAT can well avoid the error accumulation and exposure bias problems during generation. Also, we employ several advanced structure (Gu et al., 2019; Ding et al., 2020b) (*Levenshtein* with source local context modelling) and our proposed training strategies (Ding et al., 2021a,b,c) as default settings.

2.5 Bidirectional Self-Training

Besides improving NMT at model level, many researchers turn to data perspective, including exploiting the parallel and monolingual data. The most representative approaches include: a) Back Translation (BT, Sennrich et al. 2016) combines the synthetic data generated with target-side monolingual data and parallel data; b) Knowledge Distillation (KD, Kim and Rush 2016) trains the model with sequence-level distilled parallel data; c) data diversification (DD, Nguyen et al. 2020) diversifies the data by applying KD and BT on parallel data. Clearly, self-training is at the core of above approaches, that is, they generate the synthetic data either from source to target or reversely, with either monolingual or bilingual data.

To this end, we propose a bidirectional self-training approach for both parallel and monolingual data (including source and target, respectively). Specifically, the base teacher models are trained with original parallel data in the first iteration (Round 1 in Table 6), and based on these forward- and backward-teachers, all available Swahili & English sentences can be used to generate the corresponding synthetic English & Swahili sentences. After balanced-sampling between synthetic and authentic data, the concatenated data can be used to train the second iteration teachers (Round 2 in Table 6).

To reveal why our approach works, we show the results in Table 8 from the point view of data complexity (Zhou et al., 2020). Self-training reduces the data complexity, thus increasing the model deterministic and in turn enhancing the model performance.

	Features
LM Features	BERT LM (Devlin et al., 2019)
	Transformer LM (Bei et al.)
	N-gram LM (Stolcke, 2002)
In-domain features	Moore-Lewis (Moore and Lewis, 2010)
Rule-based features	Illegal characters (Bei et al.)
Count Features	Word count

Table 1: Features for back translation data selection.

2.6 Data Selection Features for Back Translation

Inspired by Ding and Tao (2019), where their cycle-translation strategy (generating high quality in-domain data) for back translation obtain substantial gains, we carefully design criteria for choosing monolingual in-domain corpus. First, we employ rule-based features, language model features. The feature types are described in Table 1. Our BERT language model used here is trained from scratch by the open-source tool¹ with target side data. The Moore-Lewis in-domain scoring strategy (Moore and Lewis, 2010) is used where the language model scores are trained with Transformer (Vaswani et al., 2017). We score all sentences in non-autoregressive fashion² to utilize contextualized information.

According to our observations, by using above multiple data selection filters, issues like illegal characters, unfluent and domain unmatched sentences could be significantly reduced. The data statistics for back translation monolingual data can be found in Table 5.

2.7 Tagged Back Translation

Back-translation (Sennrich et al., 2016; Bojar et al., 2018), translating the large scale monolingual corpus to generate synthetic parallel data by Target-to-Source pretrained model, has been widely utilized to improve the translation quality. However, recent studies find that back translation increase the target-original test set performance rather than source-original ones from the perspective of translationese³ (Zhang and Toral, 2019; Graham et al., 2020). To eliminate such concerns, we leverage tagged back translation (Caswell et al., 2019) to im-

¹<https://github.com/huggingface/pytorch-pretrained-BERT>

²<https://github.com/alphadl/EasyScore>

³Source-Original denotes the testing data originating in the source language, while target-original denotes the data translating from the target language.

src	<i>Msimu uliopita wa Siltala kwenye ligi ilikuwa 2006-07</i>
pred	<i>Siltala’s previous season in the league was 2006 at 07</i>
+post	<i>Siltala’s previous season in the league was 2006-07</i>

Table 2: Example of the effectiveness of post-processing in handling inconsistent number translation.

prove the source-original testing performance. The implementation is straightforward, that is, adding a simple tag on the beginning of each source-side synthetic sentence. The detailed reason why this trick works can be found in Marie et al., 2020.

To ensure tagged back translation works well for our task, we carefully selected the target side in-domain monolingual data (§2.6). Final results in Table 7 show the effectiveness of tagged back translation #9 against competitive model #8 (+1.9 BLEU scores).

2.8 Transductive Fine-Tuning

The key idea of transductive finetune is that source input sentences from the validation and test sets are firstly translated to the target language space with the best well-performed NMT model, which results in a pretranslated synthetic dataset. Then models are finetuned on the generated synthetic dataset. We borrow this concept from previous systems (Wu et al., 2020b; Wang et al.). We empirically show that transductive finetune (#10 – 11 in Table 7) indeed improves the official validation performance but harms the performance of our sampled valid& test set that co-distributed with the training set. Note that we randomly sampled 5K/5K sentences from the training set as valid and test sets, respectively, to avoid the sub-optimal problem caused by the distribution gap. Experimental details can be found in §3 and 4.

2.9 Reranking N-best Hypotheses

As the NMT decoding being generally from left to right, this leads to label bias problem (Lafferty et al., 2001). To alleviate this problem, besides using NAT (§2.4), we rerank the n-best hypotheses through training a k -best batch MIRA ranker (Cherry and Foster, 2012) with multiple features on validation set. The feature pool we integrated include R2L (right-to-left) translation

model, T2S (target-to-source) translation model, language model and IBM model 2 alignment score. After multi-feature reranking, the best hypothesis was retained.

Right-to-Left NMT Model The R2L NMT model using the same training data but with inverted target sentences (*i.e.*, reverse target side characters “a b c d”→“d c b a”). Then, inverting the hypothesis in the n -best list can obtain perplexity score by R2L model.

Target-to-Source NMT Model The T2S model was initially trained for back-translation, we can employ this model to assess the translation adequacy as well by adding the T2S feature to reranking feature pool.

Language Model Besides above features, we employ language models as an auxiliary feature to give the fluent sentences better scores such that the results are easier to understand by human.

2.10 Post Processing

Besides general post-processing (*i.e.*, de-BPE, de-tokenization and de-truecase⁴), we also used a post-processing algorithm (Wang et al., 2018) for inconsistent number, date translation, for example, “2006-07” might be segmented as “2006 -@@ 07” by BPE, resulting in the wrong translation “2006 at 07”. Our post-processing algorithm will search for the best matching number string from the source sentence to replace these types of errors, see Table 2.

3 Data Preparation

For ASR task, we downloaded all available Swahili speech-to-text data⁵, such as openslr⁶ and IARPA Babel⁷ etc., as training corpus and employ all default settings in Kaldi⁸ to preprocess and train them. To simplify the ASR task, we lowercased all Swahili sentences and removed punctuation. To rejuvenate these case and punctuation information, we design two pipeline tasks after ASR: *case correction* task and *punctuation generation*. Also, it is worth noting that we design some rules to perform the “voice activity detection” process for the

⁴<https://github.com/moses-smt/mosesdecoder/tree/master/scripts>

⁵<https://iwslt.org/2021/low-resource>

⁶<https://www.openslr.org/25/>

⁷<https://catalog.ldc.upenn.edu/LDC2017S05>

⁸<https://github.com/kaldi-asr/kaldi>

Available Parallel Corpus	#Sent.
CCAligned	2,044,993
Tanzil	138,253
ParaCrawl	132,517
WikiMatrix	51,387
GlobalVoices	32,307
TED2020	9,754
Gamayun	5,000
WikiMedia	771
Total	2,414,982

Table 3: Statistics of parallel data.

Sampled Mono. Corpus	#Sent.
commoncrawl English	4,366,344
commoncrawl Swahili	38,928
+ upsampling (14×)	544,992

Table 4: Statistics of monolingual data.

official speech testset. Take a piece of speech in Figure 1 for example, partial of speech in the red box will be keep as the valid input.

For NMT task, the parallel datasets we utilized are described at Table 3, including CCAligned (El-Kishky et al., 2020), Tanzil (Tiedemann, 2012), ParaCrawl⁹, WikiMatrix (Schwenk et al., 2019), GlobalVoices (Tiedemann, 2012), TED2020 (Reimers and Gurevych, 2020), WikiMedia (Tiedemann, 2012) and Gamayun¹⁰. The monolingual data we utilized are described in Table 4 and Table 5, where the monolingual data in Table 4 are used to train the system #1 – 8 in Table 7, and data in Table 5 are used to train the system #9 – 11 in Table 7, respectively. Table 6 denotes how the data used and generated by iterative bidirectional self-training (§2.5). The total data size after two round of bidirectional self-training is 50.4M, and after tagged back translation, the final data volume is 60.4M.

To avoid the sub-optimal problem caused by the distribution gap between official validation and training data, we randomly sampled 5K/ 5K sentences from the training set as valid and test sets, respectively. The randomly sampled valid sentences are used to optimize the hype-parameters.

⁹<https://www.paracrawl.eu/index.php>

¹⁰<https://gamayun.translatorswb.org/data/>




Figure 1: An example of how our rule-based voice activity detection model works on the waveform. Note that only the part in the red box will be retained as a valid fragment.

Mono. Corpus for Tagged BT	#Sent.
Totally collected corpus	
commoncrawl English	30,513,498
Cleaned corpus with criteria in §2.6	
in domain English	10,000,000

Table 5: Statistics of monolingual data for Tagged Back-Translation.

4 Experiments

Settings For *case correction* and *punctuation generation* tasks mentioned in §3, we tried Autoregressive Transformer-BASE (AT, §2.1), Non-Autoregressive model (NAT, §2.4) and our previously designed SLOTREFINE (Wu et al., 2020a). In our preliminary experiments, NAT and SlotRefine work better on *case correction* and *punctuation generation* tasks, respectively, thus leaving as the default components in our final speech translation pipeline.

For NMT task, we tried Autoregressive Transformer-BIG (AT, §2.1) and Non-Autoregressive model (NAT, §2.4) in preliminary experiments, and found that AT performs robust on all settings. Thus we employ Transformer-BIG for all MT systems. Inspired by He et al. (2019), we empirically adopt large batch strategy (Edunov et al., 2018) (i.e. 458K tokens/batch) to optimize the performance. The learning rate warms up to 1×10^{-7} for 10K steps, and then decays for 30K (data volumes range from 2M to 10M) / 50K (data volumes large than 10M) steps with the cosine schedule. For regularization, we tune the dropout rate from [0.1, 0.2, 0.3] based on validation performance, and apply weight decay with 0.01 and label smoothing with $\epsilon = 0.1$. We use Adam optimizer (Kingma and Ba, 2015) to train models. We evaluate the performance on an ensemble of last 10 checkpoints to avoid stochasticity.

For fair comparison, the metric we employed is sacreBLEU (Post, 2018). Training set, validation set and test set are processed consistently. Both Swahili and English sentences are performed tok-

#	Data Statistics	#Sent.
Preparing for Self-Training		
1	parallel English	2.4M
2	parallel Swahili	2.4M
3	monolingual English	4.4M
4	monolingual Swahili	0.4M
Self-Training Round 1		
5	synthetic parallel	9.6M
6	authentic parallel	2.4M
7	+ upsampling (4×)	9.6M
8	concat #5 and #7	19.2M
Self-Training Round 2		
9	refined parallel #8	19.2M
10	concat #8 and #9	38.4M
11	upsampled authentic parallel #6 (5×)	12.0M
12	concat #10 and #11	50.4M

Table 6: Data statistics for bidirectional self-training. Note that #5 “synthetic parall” comes from monolingual English (#3), monolingual Swahili (#4), parallel English (#1), and parallel Swahili (#2). In our preliminary experiments, 4× (#7) and 5× (#11) upsampling strategies perform best in their corresponding settings, thus leaving as the default settings.

enization and truecasing with Moses scripts (Koehn et al., 2007). In order to limit the size of vocabulary of NMT models, we adopted byte pair encoding (BPE) (Sennrich et al., 2016) with 32k operations. Larger beam size may worsen translation quality (Koehn and Knowles, 2017), thus we set beam_size=10 when performing n-best reranking (§2.9). All models were trained on 4 16GB NVIDIA V100 GPUs.

Main results Our main experiment is shown in Table 7, our baseline system is developed with the original parallel corpus and last-10 ensemble strategy. Unsurprisingly, the baseline system relatively performs the worst.

The proposed *Bidirectional Pretrain* in §2.2 and *Denoising Pretrain* in §2.3 could consistently and significantly improve the model performance, showing their effectiveness in low resource scenarios (Zhang and Tao, 2020). Clearly, combining *Bidirectional Pretrain* and *Denoising Pretrain*

#	Models	Valid	Test	Δ_{ave}	Off. Valid	Δ
1	Baseline (w/ Para. Data)	47.1	48.5	–	31.8	–
2	+Bidirectional Pretrain	48.5	49.9			
3	+Denoising Pretrain	48.6	49.6			
4	+Combination of #2 and #3	48.9	50.1	+1.7		
5	Bi. Self-Training (w/ Mono. & Para. Data)	49.4	50.8	+2.3		
6	+Combination of #2 and #3	50.1	51.6	+3.1		
7	Iterative Bi. Self-Training	49.7	50.9	+2.5		
8	+Combination of #2 and #3	50.5	51.8	+3.4	38.2	+6.4
9	#8 + Tagged Back Translation	52.4	53.1	+5.0	40.1	+8.3
10	#9 + Transductive Finetune	51.8	53.0	+4.6	41.5	+9.7
11	+Iterative + #10	51.6	52.8	+4.4	41.9	+10.1
12	#11 + Reranking	52.1	53.5	+5.0	42.3	+10.5
13	#12 + Post Processing	52.5	54.0	+5.5	42.6	+10.8

SacreBLEU of Final Submission (#13) on official test set **25.3**

Table 7: Sacrebleu of Sw→En on our randomly sampled “Valid/ Test” sets and official validation set “Off. Valid”, where “ Δ ” represents the performance gains compared with baseline #1. The submitted system is #13.

Data	Compl.	BLEU
Baseline	7.87	47.1
Bi. Self-Training	5.34	49.4
Iterative Bi. Self-Training	4.89	49.7

Table 8: Explanation of why Bidirectional Self-Training works. The data complexity “Compl.” is measured on their corresponding training sets and alignment information is trained with *fast-align* (Dyer et al., 2013). The BLEU scores are reported on our sampled validation set.

could achieve better results (averaged +1.7 BLEU scores), indicating their complementary.

As shown in #5 and #7, the proposed *Bidirectional Self-Training* and its refined iterative version, could consistently enhance the model. To explore why self-training improves model performance, we discuss it from the point view of data complexity. As shown in Table 8, with the *Bidirectional Self-Training* iteratively progresses, the data complexity becomes lower, leading to the better BLEU scores. Notably, the combination of our proposed two pretraining approaches push the SOTA performance up to higher points. We believe that the effect of our proposed two pretrain strategies are still under-investigated, which will leave as future works. Overall, with strategies #2 – 8, the model performance in terms of official validation test achieves surprisingly +6.4 BLEU scores.

The *Tagged Back Translation* (§2.7) with in-

domain monolingual data significantly improves the performance of both our sampled test set and official valid set by +5.0 and +8.3 against baseline, respectively.

We empirically show that *Transductive FineTune* (§2.8) indeed improves the official validation performance but harms the performance of our sampled valid& test set that co-distributed with the training set. This indicates that transductive learning is an effective practice to transfer a well-trained model across domains.

And the last two strategies *Reranking* (§2.9) and *Post Processing* (§2.10) could further improve the official validation BLEU score from 41.9 to 42.6, which substantially outperforms the baseline by +10.8 BLEU score.

5 Conclusion and Future Work

This paper presents the University of Sydney & JD’s speech machine translation system for IWSLT2021 Swahili→English task. The whole system is pipelined, containing ASR, case correction, punctuation generation and NMT tasks, and we mainly focused on NMT task.

We leveraged multi-dimensional strategies and frameworks to improve the translation qualities, which achieves surprisingly +10.8 BLEU scores improvement against baseline and ranks the **1st** among all the participants. We find that our proposed BIDIRECTIONAL PRETRAINING (§2.2) and DENOISING PRETRAINING (§2.3) can consistently

improves the competitive baselines. Also, we employ BIDIRECTIONAL SELF-TRAINING in §2.5 and TAGGED BT in §2.7 make the most of the existing parallel and monolingual data.

In the future, we would like to polish other components in the pipeline to achieve better performance. Also, it is worthy to try an end-to-end approach with cross-modal structures to incorporate audio and vision knowledge (Xu et al., 2021). For robust model training and data utilization, we would explore better strategies, e.g. adversarial training (Wu et al., 2021) and curriculum learning (Liu et al., 2020a; Zhou et al., 2021).

Acknowledgments

We would thank anonymous reviewers for their considerate proofreading and valuable comments.

References

- Dzmitry Bahdanau, KyungHyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of ICLR 2015*.
- Chao Bei, Hao Zong, Yiming Wang, Baoyong Fan, Shiqi Li, and Conghu Yuan. An empirical study of machine translation for the shared task of WMT18. In *WMT*.
- Ondřej Bojar, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Philipp Koehn, and Christof Monz. 2018. Findings of the 2018 conference on machine translation (WMT18). In *WMT*.
- Isaac Caswell, Ciprian Chelba, and David Grangier. 2019. Tagged back-translation. In *WMT*, Florence, Italy.
- Colin Cherry and George Foster. 2012. Batch tuning strategies for statistical machine translation. In *NAACL*.
- Trevor Cohn, Cong Duy Vu Hoang, Ekaterina Vymolova, Kaisheng Yao, Chris Dyer, and Gholamreza Haffari. 2016. Incorporating structural alignment biases into an attentional neural translation model. In *NAACL*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Liang Ding and Dacheng Tao. 2019. The University of Sydney’s machine translation system for WMT19. In *WMT*.
- Liang Ding, Longyue Wang, Xuebo Liu, Derek F. Wong, Dacheng Tao, and Zhaopeng Tu. 2021a. Progressive multi-granularity training for non-autoregressive translation. In *findings of ACL*.
- Liang Ding, Longyue Wang, Xuebo Liu, Derek F. Wong, Dacheng Tao, and Zhaopeng Tu. 2021b. Rejuvenating low-frequency words: Making the most of parallel data in non-autoregressive translation. In *ACL*.
- Liang Ding, Longyue Wang, Xuebo Liu, Derek F. Wong, Dacheng Tao, and Zhaopeng Tu. 2021c. Understanding and improving lexical choice in non-autoregressive translation. In *ICLR*.
- Liang Ding, Longyue Wang, and Dacheng Tao. 2020a. Self-attention with cross-lingual position representation. In *ACL*.
- Liang Ding, Longyue Wang, Di Wu, Dacheng Tao, and Zhaopeng Tu. 2020b. Context-aware cross-attention for non-autoregressive translation. In *COLING*.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. A simple, fast, and effective reparameterization of IBM model 2. In *NAACL*.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. Understanding back-translation at scale. In *Proceedings of EMNLP 2019*.
- Ahmed El-Kishky, Vishrav Chaudhary, Francisco Guzmán, and Philipp Koehn. 2020. CCAligned: A massive collection of cross-lingual web-document pairs. In *EMNLP 2020*.
- Yvette Graham, Barry Haddow, and Philipp Koehn. 2020. Statistical power and translationese in machine translation evaluation. In *EMNLP*.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. 2018. Non-autoregressive neural machine translation. In *ICLR*.
- Jiatao Gu, Changhan Wang, and Junbo Zhao. 2019. Levenshtein transformer. In *NeurIPS*.
- Fengxiang He, Tongliang Liu, and Dacheng Tao. 2019. Control batch size and learning rate to generalize well: Theoretical and empirical evidence. In *NeurIPS*.
- Tianyu He, Xu Tan, Yingce Xia, Di He, Tao Qin, Zhibo Chen, and Tie-Yan Liu. 2018. Layer-wise coordination between encoder and decoder for neural machine translation. In *NeurIPS*.
- Dan Hendrycks, Kimin Lee, and Mantas Mazeika. 2019. Using pre-training can improve model robustness and uncertainty. In *ICML*.
- Frederick Jelinek. 1997. *Statistical methods for speech recognition*. MIT press.
- Yoon Kim and Alexander M. Rush. 2016. Sequence-level knowledge distillation. In *EMNLP*.

- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Philipp Koehn. 2009. *Statistical machine translation*. Cambridge University Press.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *ACL*.
- Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. In *WNMT*.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*.
- Tomer Levinboim, Ashish Vaswani, and David Chiang. 2015. Model invertibility regularization: Sequence alignment with or without parallel data. In *NAACL*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *ACL*.
- P. Liang, D. Klein, and Michael I. Jordan. 2007. Agreement-based learning. In *NeurIPS*.
- Xuebo Liu, Houtim Lai, Derek F Wong, and Lidia S Chao. 2020a. Norm-based curriculum learning for neural machine translation. In *ACL*.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020b. Multilingual denoising pre-training for neural machine translation. *TACL*.
- Benjamin Marie, Raphael Rubino, and Atsushi Fujita. 2020. Tagged back-translation revisited: Why does it really work? In *ACL*.
- Alexander Mathis, Thomas Biasi, Steffen Schneider, Mert Yuksekgonul, Byron Rogers, Matthias Bethge, and Mackenzie W Mathis. 2021. Pretraining boosts out-of-domain robustness for pose estimation. In *WACV*.
- Robert C. Moore and William Lewis. 2010. Intelligent selection of language model training data. In *ACL*.
- Hermann Ney. 1999. Speech translation: Coupling of recognition and translation. In *ICASSP*.
- Xuan-Phi Nguyen, Shafiq Joty, Wu Kui, and Ai Ti Aw. 2020. Data diversification: A simple strategy for neural machine translation. In *NeurIPS*.
- Matt Post. 2018. A call for clarity in reporting BLEU scores. In *WMT*.
- Nils Reimers and Iryna Gurevych. 2020. Making monolingual sentence embeddings multilingual using knowledge distillation. In *EMNLP*.
- Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2019. Wiki-matrix: Mining 135m parallel sentences in 1620 language pairs from wikipedia. *arXiv*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving neural machine translation models with monolingual data. In *ACL*, Berlin, Germany.
- Andreas Stolcke. 2002. Srlm-an extensible language modeling toolkit. In *IWSLT*.
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in opus. In *LREC*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of NIPS 2017*.
- Longyue Wang, Zhaopeng Tu, Xing Wang, Li Ding, Liang Ding, and Shuming Shi. Tencent AI lab machine translation systems for WMT20 chat translation task. In *WMT*.
- Qiang Wang, Bei Li, Jiqiang Liu, Bojian Jiang, Zheyang Zhang, Yinqiao Li, Ye Lin, Tong Xiao, and Jingbo Zhu. 2018. The NiuTrans machine translation system for WMT18. In *WMT*.
- Di Wu, Yiren Chen, Liang Ding, and Dacheng Tao. 2021. Bridging the gap between clean data training and real-world inference for spoken language understanding. *arXiv*.
- Di Wu, Liang Ding, Fan Lu, and Jian Xie. 2020a. Slotrefine: A fast non-autoregressive model for joint intent detection and slot filling. In *EMNLP*.
- Lijun Wu, Yiren Wang, Yingce Xia, QIN Tao, Jianhuang Lai, and Tie-Yan Liu. 2019. Exploiting monolingual data at scale for neural machine translation. In *EMNLP*.
- Shuangzhi Wu, Xing Wang, Longyue Wang, Fangxu Liu, Jun Xie, Zhaopeng Tu, Shuming Shi, and Mu Li. 2020b. Tencent neural machine translation systems for the wmt20 news translation task. In *WMT*.
- Yufei Xu, Qiming Zhang, Jing Zhang, and Dacheng Tao. 2021. Vitae: Vision transformer advanced by exploring intrinsic inductive bias. *arXiv*.
- Jing Zhang and Dacheng Tao. 2020. Empowering things with intelligence: a survey of the progress, challenges, and opportunities in artificial intelligence of things. *IEEE Internet of Things Journal*.
- Mike Zhang and Antonio Toral. 2019. The effect of translationese in machine translation test sets. In *WMT*.

Zaixiang Zheng, Hao Zhou, Shujian Huang, Lei Li, Xin-Yu Dai, and Jiajun Chen. 2019. Mirror-generative neural machine translation. In *ICLR*.

Chunting Zhou, Graham Neubig, and Jiatao Gu. 2020. Understanding knowledge distillation in non-autoregressive machine translation. In *ICLR*.

Lei Zhou, Liang Ding, Kevin Duh, Ryohei Sasano, and Koichi Takeda. 2021. Self-guided curriculum learning for neural machine translation. In *IWSLT*.