

End-to-End Speech Translation with Knowledge Distillation

Yuchen Liu^{1,2}, Hao Xiong⁴, Zhongjun He⁴, Jiajun Zhang^{1,2}, Hua Wu⁴, Haifeng Wang⁴ and Chengqing Zong^{1,2,3}

¹NLPR, Institute of Automation, Chinese Academy of Sciences, Beijing, China

²University of Chinese Academy of Sciences, China

³CAS Center for Excellence in Brain Science and Intelligence Technology, Beijing, China

⁴Baidu Inc. No. 10, Shangdi 10th Street, Beijing, 100085, China

{yuchen.liu, jjzhang, cqzong}@nlpr.ia.ac.cn,
{xionghao05, hezhongjun, wu.hua, wanghaifeng}@baidu.com

Abstract

End-to-end speech translation (ST), which directly translates from source language speech into target language text, has attracted intensive attentions in recent years. Compared to conventional pipeline systems, end-to-end ST models have advantages of lower latency, smaller model size and less error propagation. However, the combination of speech recognition and text translation in one model is more difficult than each of these two tasks. In this paper, we propose a *knowledge distillation* approach to improve ST model by transferring the knowledge from text translation model. Specifically, we first train a text translation model, regarded as a teacher model, and then ST model is trained to learn output probabilities from teacher model through knowledge distillation. Experiments on English-French Augmented LibriSpeech and English-Chinese TED corpus show that end-to-end ST is possible to implement on both similar and dissimilar language pairs. In addition, with the instruction of teacher model, end-to-end ST model can gain significant improvements by over 3.5 BLEU points.

Index Terms: Speech recognition, Speech translation, Knowledge distillation, Transformer

1. Introduction

Conventional speech translation system is a pipeline of two main components: an automatic speech recognition (ASR) model which provides transcripts of source language utterances, and a text machine translation (MT) model which translates the transcripts to target language [1, 2, 3, 4, 5]. This pipeline system usually suffers from time delay, parameter redundancy and error accumulation. In contrast, end-to-end ST, based on an encoder-decoder architecture with attention mechanism, is more compact and efficient. It can directly generate translations from raw audio and jointly optimize parameters on the final goal. Therefore, this model has become a new trend in speech translation research studies [6, 7, 8, 9, 10, 11].

However, despite appealing advantages of end-to-end ST model, its performance is generally inferior. One of the important reasons is due to extremely scarce data which includes speech in source language paired with text in target language. Previous studies resort pretraining or multi-task learning approaches to improve the translation quality. They either pretrain ASR task on high-resource data [11], or use multi-task learning to train ST model with ASR or MT model simultaneously [9, 10]. Nevertheless, they only gain limited improvements and do not take full advantage of text data. We notice that the performance between end-to-end ST and MT model exists a huge

gap, thus how to utilize MT model to help instruct end-to-end ST model is of great significance.

It is a challenge to train an end-to-end ST model directly from speech signal without text guidance while achieving comparable performance as text translation model. Given that text translation models are superior to ST model, we consider ST model can be improved by leveraging *knowledge distillation*. In knowledge distillation, there is usually a big teacher model with a small student model. It has been shown that the output probabilities of teacher model are smooth, which are easier for student model to learn from than ground-truth text [12]. Thus, a student model can be taught by imitating the behaviour of teacher model, such as output probabilities [12, 13], hidden representation [14, 15], or generated sequence [16], and alleviate the performance gap between itself and the teacher model.

In this paper, we present a method based on knowledge distillation for end-to-end ST model to learn knowledge from text translation model. We first train a text translation model on parallel text data (regarded as teacher) and then an end-to-end ST model (regarded as student) is trained by learning from ground-truth translations and the outputs of teacher model simultaneously. Experiments conducted on 100h English-French Augmented LibriSpeech corpus and 542h English-Chinese TED corpus show that it is possible to train a compact end-to-end speech translation model on both similar and dissimilar language pairs. With the instruction of teacher model, end-to-end ST model can gain significant improvements, approaching to the traditional pipeline system.

2. Related Work

End-to-end model has already become a dominant paradigm in machine translation task, which adopts an encoder-decoder architecture and generates target words from left to right at each step [1, 3, 5]. This model has also achieved promising results in ASR fields [2, 4, 17]. Recent works purpose a further attempt to combine these two tasks together by building an end-to-end speech-to-text translation without the use of source language translation during learning or decoding.

Anastasopoulos et al. [6] use k -means clustering to cluster repeated audio patterns and automatically align spoken words with their translations. Duong et al. [7] focus on the alignment between speech and translated phrase but not to directly predict the final translations. Bérard et al. [8] give the first proof of the potential for end-to-end speech-to-text translation without using source language. They further conduct experiments on a larger English-to-French dataset and pre-train encoder and

decoder which improves performance [10]. Weiss et al. [9] also use multi-task learning and show that end-to-end model can outperform a cascade of independently trained pipeline system on Fisher Callhome Spanish-English speech translation task. Bansal et al. [11] find pretraining encoder on higher-resource language ASR training data can achieve gains in low-resource speech translation system. However, these work mainly focus on pretraining acoustic encoder and do not take full advantage of text data.

Knowledge distillation is first adopted to apply for model compression, the main idea of which is to train a smaller student model to mimic a larger teacher model, by minimizing the loss between the teacher and student predictions. It has soon been applied to a variety of tasks, like image classification [12, 18, 19, 20], speech recognition [12] and natural language processing [13, 16, 21]. The teacher and student model in conventional knowledge distillation usually handle the same task, while in our method the teacher model and student model have different input modalities where teacher uses text as input and student uses speech.

3. Models

In this paper, we apply end-to-end models with the same architecture for all three tasks (ASR, ST and MT). The model architecture is similar with *Transformer* [5], which is the state-of-art model in MT task. Recently, this model also begins to be used in ASR task, showing a decent performance [22, 23]. In this section, we first describe the core architecture of *Transformer* and then show how this model is applied to ASR/ST and MT task.

3.1. Core Module of Transformer

Transformer is an encoder-decoder architecture which entirely relies on self-attention mechanism including scaled dot-product attention and multi-head attention. It consists of N stacked encoder and decoder layers. Each encoder layer has two blocks, which is a self-attention block followed by a feed-forward block. Decoder layer has the same architecture with encoder layer except an extra encoder-decoder attention block to perform attention over the output of the top encoder layer. Residual connection and layer normalization are employed around each block. In addition, the self-attention block in the decoder is modified with mask to prevent present positions attending to future positions during training.

To be detailed, multi-head attention technique is applied in self-attention and encoder-decoder attention blocks to obtain information from different representation subspaces at different positions. Each head is corresponding to a scaled dot-product attention, which operates on query Q , key K and value V :

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}}\right)\mathbf{V} \quad (1)$$

where d_k is the dimension of the key. Then the output values are concatenated,

$$\text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)\mathbf{W}^O$$

where $\text{head}_i = \text{Attention}(\mathbf{Q}\mathbf{W}_i^Q, \mathbf{K}\mathbf{W}_i^K, \mathbf{V}\mathbf{W}_i^V)$ (2)

where the $\mathbf{W}_i^Q \in \mathbb{R}^{d_{model} \times d_q}$, $\mathbf{W}_i^K \in \mathbb{R}^{d_{model} \times d_k}$, $\mathbf{W}_i^V \in \mathbb{R}^{d_{model} \times d_v}$ and $\mathbf{W}^O \in \mathbb{R}^{d_v \times d_{model}}$ are projection matrices

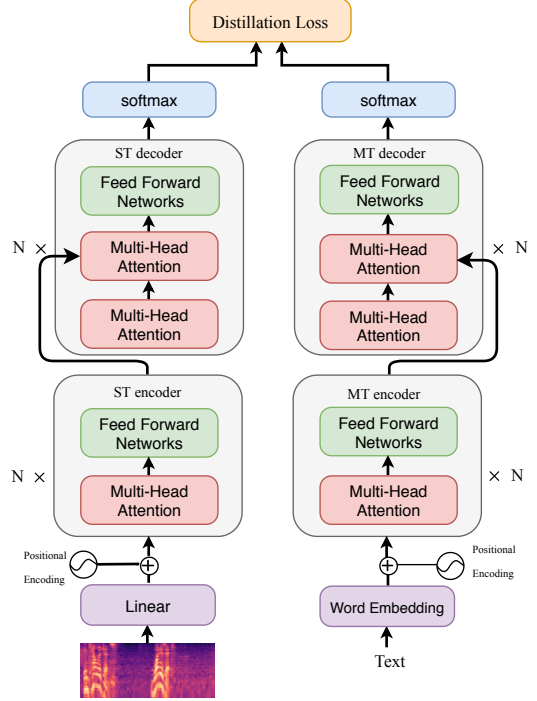


Figure 1: *Model architecture of our method. The left part is ST model, regarded as a student model, whose input is speech signal. The right part is MT model, regarded as a teacher model, whose input is the source sentence corresponding to the input of student model. The output of both student model and teacher model is target sentence. The top part is distillation loss, where the student model not only matches the ground-truth, but also the output probabilities of the teacher model.*

that are learned. $d_q = d_k = d_v = d_{model}/h$, h is the number of heads.

Position-wise feed-forward block is composed of two linear transformations with a ReLU activation in between.

$$\text{FFN}(x) = \max(0, x\mathbf{W}_1 + \mathbf{b}_1)\mathbf{W}_2 + \mathbf{b}_2 \quad (3)$$

where the weights $\mathbf{W}_1 \in \mathbb{R}^{d_{model} \times d_{ff}}$, $\mathbf{W}_2 \in \mathbb{R}^{d_{ff} \times d_{model}}$ and the biases $\mathbf{b}_1 \in \mathbb{R}^{d_{ff}}$, $\mathbf{b}_2 \in \mathbb{R}^{d_{model}}$.

3.2. ASR/ST Model

The ASR/ST model is shown in the left part of Figure 1, whose input is a series of discrete-time speech signal. We first use log-Mel filterbank to convert raw speech signal into a sequence of acoustic features and then apply mean and variance normalization. To prevent the GPU memory overflow and produce approximate hidden representation length against target length, we apply frame stack and downsample similar to [24, 25]. The final acoustic feature sequence is $S = (s_1, s_2, \dots, s_n)$ with dimension of $d_{filterbank} \times num_{stack}$. Then the feature sequence is fed into a linear transformation with a normalization layer to map with model dimension d_{model} . In addition, positional encodings are added to the feature sequence in order to enable the model to attend by relative positions. This sequence is finally treated as the input into *Transformer* model. Other parts are the same with *Transformer* model. For ASR the input to decoder is source language text, while the input to decoder in ST is target language text.

3.3. MT Model

We also use *Transformer* to train a baseline MT model, as shown in the right part of Figure 1. The difference between MT model and ASR/ST model is the input to the encoder. In MT model, $X = (x_1, x_2, \dots, x_n)$ is a sequence of tokens, representing source sentence. We embed the words in sequence X into a real continuous space with the dimension of \mathbb{R}_{model}^d , which can be fed into a neural network.

3.4. Knowledge Distillation

Training an end-to-end ST model is considerably difficult than MT model. The accuracy of the later model is usually much higher than the former. Therefore, we present MT model as a teacher to teach ST model. Here we give a description of the idea of knowledge distillation.

Denote $D = (s, x, y)$ as the corpus of triple data corresponding to speech signal, transcription in source language and its translation. The log-likelihood loss of ST model can be formulated as follows:

$$L_{ST}(D; \theta) = - \sum_{(s, y) \in D} \log P(y|s; \theta) \quad (4)$$

$$\log P(y|s; \theta) = \sum_{t=1}^N \sum_{k=1}^{|V|} \mathbb{1}(y_t = k) \log P(y_t = k|y_{<t}, s; \theta) \quad (5)$$

where s is the acoustic feature sequence of source speech signal, y is the target translated sentence, N is the length of the output sequence, $|V|$ is the vocabulary size of the output language, y_t is the t -th output token, $\mathbb{1}(y_t = k)$ is an indicator function which indicates whether the output token is equal to the ground-truth.

We denote the output distribution of teacher model for token y_t as $Q(y_t|y_{<t}, x; \theta_T)$, and x is the source transcribed sentence which corresponds to speech signal s . Then the cross entropy between the distributions of teacher and student is:

$$L_{KD}(D; \theta, \theta_T) = - \sum_{(x, y) \in D} \sum_{t=1}^N \sum_{k=1}^{|V|} Q(y_t = k|y_{<t}, x; \theta_T) \log P(y_t = k|y_{<t}, x; \theta) \quad (6)$$

In distillation loss, the student not only matches the output of ground-truth, but also the output probabilities of teacher model, which is more smooth and yields smaller variance in gradients [12]. Then the total loss function is,

$$L_{ALL}(D; \theta; \theta_T) = (1 - \lambda)L_{ST}(D; \theta) + \lambda L_{KD}(D; \theta, \theta_T) \quad (7)$$

where λ is a hyper-parameter to trade off these two loss terms.

4. Experiments

4.1. Datasets

We conduct experiments on Augmented LibriSpeech which is collected by [26] and available for free. This corpus is built by automatically aligning e-books in French with English utterances of LibriSpeech, which contains 236 hours of speech in total. They provide quadruplet: English speech signal, English transcription, French text translations from alignment of e-books and *Google Translate* references. Following [10], We only use the 100 hours clean train set for training, with 2 hours

development set and 4 hours test set, which corresponds to 47,271, 1071 and 2048 utterances respectively. To be consistent with their settings, we also double the training size by concatenating the aligned references with the *Google Translate* references.

To verify whether the end-to-end speech translation model can handle on dissimilar language pairs, we build a corpus in English-Chinese direction. The raw data (including video, subtitles and timestamps) are crawled from TED website¹. For each talk, we build a *wav* audio file extracted from video by *ffmpeg*². We also collect its corresponding transcript and save in *txt* format. We divide each audio file into small segments based on timestamps instead of voice activity detection (VAD), because it eliminates the influence of improper fragments and guarantees each utterance containing complete semantic information, which is important for translation. In the end, we totally get 317,088 utterances (~ 542 hours). Development and test sets are split according to the partition in IWSLT. We use dev2010 as development set and tst2015 as test set, which has 835 utterances (~ 1.48 hours) and 1,223 utterances (~ 2.37 hours) respectively. The remaining data are put into training set. We will release this dataset to public as a benchmark soon.

4.2. Experimental Setup

Our acoustic features are 80-dimensional log-Mel filterbanks extracted with a step size of 10ms and window size of 25ms and extended with mean subtraction and variance normalization. The features are stacked with 3 frames to the left and downsample to a 30ms frame rate. For text data, we lowercase all the texts, tokenize and apply normalize punctuations with the Moses scripts³. For Augmented LibriSpeech corpus, we apply BPE [27] on the combination of English and French text to obtain subword units. The number of merge operations in BPE is set to 8K, resulting in a shared vocabulary with 8,159 subwords. For TED English-Chinese, the merge number is 30K, and vocabulary size are 28,912 and 30,000, respectively. We report case-insensitive BLEU scores [28] by *multi-bleu.pl* script for the evaluation of ST and MT tasks and use word error rates (WER) to evaluate ASR task.

Because the size of Augmented LibriSpeech is relatively small, we set the hidden size $d_{model} = 256$, the filter size in feed-forward layer $d_{ff} = 1024$, the head number $h = 8$, the residual dropout and attention dropout are 0.1. For TED English-Chinese, we set the hidden size $d_{model} = 512$ with the filter size $d_{ff} = 2048$. MT model, as a teacher model, can use bigger parameters. We use 512 hidden sizes, 2048 filter sizes with 8 heads. The number of encoder layers and decoder layers in above models are all set to 6. We train our models with Adam optimizer [29] with $\beta_1 = 0.9$, $\beta_2 = 0.98$ and $\epsilon = 10^{-9}$ on 2 NVIDIA V100 GPUs.

4.3. Results

Table 1 shows the results for the ASR and MT tasks on Augmented LibriSpeech. It can be seen that *Transformer* model significantly outperforms in both ASR and MT tasks, with 0.92 WER reduction and 4.1 BLEU scores improvement in beam search compared to [10]. We contribute it to the superior performance of *Transformer* model which is good at modeling long distance in sequence-to-sequence tasks, especially for MT

¹<https://www.ted.com>

²<http://ffmpeg.org>

³<https://www.statmt.org/ Moses/>

Table 1: ASR and MT results on test set of Augmented LibriSpeech.

LibriSpeech	Method	WER(↓)	BLEU(↑)
Bérard [10]	greedy	19.9	19.2
	beam search	17.9	18.8
Ours	greedy	21.46	21.35
	beam search	16.98	22.91

Table 2: ST results on Augmented LibriSpeech test. KD denotes knowledge distillation.

LibriSpeech	Method	greedy	beam	ensemble
Bérard [10]	Pipeline	14.6	14.6	15.8
	End-to-end	12.3	12.9	15.5
	Pre-trained	12.6	13.3	
Ours	Pipeline	15.75	17.85	18.4
	End-to-end	10.19	13.15	17.8
	Pre-trained	13.89	14.30	
	KD	14.96	17.02	

tasks. Contrary to [10] which uses characters as output units, we consider subword units can also obtain improvements.

For ST task, we have four settings. The pipeline model uses ASR outputs as MT inputs, where ASR model and MT model are described above. The *end-to-end* model is directly trained on source speech signal paired with target text translations. The *pre-trained* model is identical to *end-to-end* model, but it is initialized with ASR and MT models. Knowledge distillation (KD) is our method which uses MT model as teacher model to instruct *end-to-end* ST model.

As shown in Table 2, all four settings surpass the results in [10]. Noticing that there exists a huge gap between the performance of the end-to-end ST model and MT model, even if the end-to-end ST model is pretrained, thus we conduct *knowledge distillation* to instruct ST model with MT model. The result shows that this method can bring significant improvement on the BLEU score which increases from 14.30 to 17.02. With the instruction of MT model, the performance gap is alleviated, approaching to the pipeline system. It demonstrates the effectiveness of our method.

We also conduct experiments on English-Chinese to verify our methods. Table 3 presents the results of MT and ST models. Pipeline model combines both the ASR (WER is 18.2%) and MT models. It is difficult to train end-to-end ST model from random initialization parameters, for the reordering between dissimilar language pairs is difficult to align with frame based speech representations. The end-to-end ST model here is pretrained with ASR. With knowledge distillation, it can obtain significant improvements, proving the generality of our method. Although end-to-end ST does not outperform pipeline system, it shows the potential to implement a compact end-to-end model even on dissimilar language pairs.

4.4. Analysis

To evaluate the effect of teacher model, we explore different hyper-parameters λ of the distillation loss on Augmented LibriSpeech. With λ increasing, ST will pay more attention to the teacher model. When λ equals 0, it is the pre-trained end-to-end model; when λ is 1, it will ignore ground-truth and only learn from the teacher. As Table 4 shows, the performance becomes better with the increasing of λ . End-to-end ST obtains

Table 3: MT and ST results on TED English-Chinese test.

TED	MT	Pipeline	End-to-end	KD
BLEU	27.08	22.28	16.80	19.55

Table 4: The effect of teacher model weight on ST results.

λ	0.0	0.2	0.4	0.6	0.8	1.0
BLEU	14.30	15.68	16.73	16.62	16.93	17.02

the best performance when it only learns the output distributions of teacher model.

We further analyze how knowledge from MT model helps ST through visualizations of the encoder-decoder attention. Figure 2 shows an example. The attentions of ASR (a) and MT (c) models have more confident than ST model. Each output token in the former two model concentrates on specific frames or tokens, especially for MT model, while the attention in ST (b) model tends to be smoothed out across many input frames. However, with the help of MT model, the attention of ST model with KD (d) becomes more concentrated. For example, the speech frames $l = 45 \sim 55$ are corresponding to “was talking” in ASR (a), which can be translated to “se parlait” in French (c). The attention in ST model with KD has more weights on frames $l = 45 \sim 55$ than that in original ST model.




Figure 2: The visualizations of attention in different models. (a), (b), (c), (d) are the encoder-decoder attention of ASR, end-to-end ST, MT and end-to-end ST with KD, respectively.

5. Conclusions

In this work, we present *knowledge distillation* method to improve the end-to-end ST model by transferring the knowledge from MT model. Experiments on two language pairs demonstrate that with the instruction of MT model, end-to-end ST model can gain significant improvements. Although the end-to-end ST does not outperform pipeline system, it shows the potential to come close in performance. In the future we will utilize other knowledges like the outputs from ASR model to further improve the performance of ST model.

6. References

- [1] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” in *Proc. ICLR*, 2015.

- [2] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *Proc. ICASSP*, 2016.
- [3] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, and e. a. Mohammad Norouzi, *Googles neural machine translation system: bridging the gap between human and machine translation*. arXiv preprint arXiv:1609.08144, 2016.
- [4] C.-C. Chiu, T. N. Sainath, Y. Wu, R. Prabhavalkar, P. Nguyen, Z. Chen, A. Kannan, R. J. Weiss, K. Rao, K. Gonina *et al.*, "State-of-the-art speech recognition with sequence-to-sequence models," in *Proc. ICASSP*, 2017.
- [5] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in Neural Information Processing Systems*, pp. 5998–6008, 2017.
- [6] A. Anastasopoulos, D. Chiang, and L. Duong, "An unsupervised probability model for speech-to-translation alignment of low-resource languages," in *Proc. EMNLP*, 2016.
- [7] L. Duong, A. Anastasopoulos, D. Chiang, S. Bird, and T. Cohn, "An attentional model for speech translation without transcription," in *Proc. NAACL*, 2016.
- [8] A. Bérard, O. Pietquin, C. Servan, and L. Besacier, "Listen and translate: A proof of concept for end-to-end speech-to-text translation," in *NeurIPS Workshop on End-to-end Learning for Speech and Audio Processing*, 2016.
- [9] R. J. Weiss, J. Chorowski, N. Jaitly, Y. Wu, and Z. Chen, "Sequence-to-sequence models can directly translate foreign speech," in *Proc. Interspeech*, 2017.
- [10] A. Bérard, L. Besacier, A. C. Kocabiyikoglu, and O. Pietquin, "End-to-end automatic speech translation of audiobooks," in *Proc. ICASSP*, 2018.
- [11] S. Bansal, H. Kamper, K. Livescu, A. Lopez, and S. Goldwater, "Pre-training on high-resource speech recognition improves low-resource speech-to-text translation," *arXiv preprint arXiv:1809.01431*, 2018.
- [12] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [13] M. Freitag, Y. Al-Onaizan, and B. Sankaran, "Ensemble distillation for neural machine translation," *arXiv preprint arXiv:1702.01802*, 2017.
- [14] J. Yim, D. Joo, J. Bae, and J. Kim, "A gift from knowledge distillation: Fast optimization, network minimization and transfer learning," in *Proc. CVPR*, pp. 4133–4141, 2017.
- [15] A. Romero, N. Ballas, S. E. Kahou, A. Chassang, C. Gatta, and Y. Bengio, "Fitnets: Hints for thin deep nets," in *Proc. ICLR*, 2015.
- [16] Y. Kim and A. M. Rush, "Sequence-level knowledge distillation," in *Proc. EMNLP*, 2016.
- [17] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, "End-to-end attention-based large vocabulary speech recognition," in *Proc. ICASSP*, 2016.
- [18] Y. Li, J. Yang, Y. Song, L. Cao, J. Luo, and L.-J. Li, "Learning from noisy labels with distillation," in *Proc. ICCV*, pp. 1910–1918, 2017.
- [19] C. Yang, L. Xie, S. Qiao, and A. Yuille, "Knowledge distillation in generations: More tolerant teachers educate better students," in *Proc. CVPR*, 2018.
- [20] R. Anil, G. Pereyra, A. Passos, R. Ormandi, G. E. Dahl, and G. E. Hinton, "Large scale distributed neural network training through online distillation," *arXiv preprint arXiv:1804.03235*, 2018.
- [21] X. Tan, Y. Ren, D. He, T. Qin, Z. Zhao, and T.-Y. Liu, "Multilingual neural machine translation with knowledge distillation," in *Proc. ICLR*, 2019.
- [22] L. Dong, S. Xu, and B. Xu, "Speech-transformer: a no-recurrence sequence-to-sequence model for speech recognition," in *Proc. ICASSP*, pp. 5884–5888, 2018.
- [23] S. Zhou, L. Dong, S. Xu, and B. Xu, "Syllable-based sequence-to-sequence speech recognition with the transformer in mandarin chinese," in *Proc. Interspeech*, 2018.
- [24] H. Sak, A. Senior, K. Rao, and F. Beaufays, "Fast and accurate recurrent neural network acoustic models for speech recognition," in *Proc. ICASSP*, 2015.
- [25] A. Kannan, Y. Wu, P. Nguyen, T. N. Sainath, Z. Chen, and R. Prabhavalkar, "An analysis of incorporating an external language model into a sequence-to-sequence model," in *Proc. ICASSP*, pp. 1–5828, 2018.
- [26] A. C. Kocabiyikoglu, L. Besacier, and O. Kraif, "Augmenting librispeech with french translations: A multimodal corpus for direct speech translation evaluation," *Language Resources and Evaluation*, 2018.
- [27] R. Sennrich, B. Haddow, and A. Birch, "Neural machine translation of rare words with subword units," in *Proc. ACL*, 2016.
- [28] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "Bleu: a method for automatic evaluation of machine translation," in *Proc. ACL*, pp. 311–318, 2002.
- [29] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.