

# M<sup>3</sup>ST: MIX AT THREE LEVELS FOR SPEECH TRANSLATION

Xuxin Cheng<sup>1†</sup>, Qianqian Dong<sup>2</sup>, Fengpeng Yue<sup>2</sup>, Tom Ko<sup>2</sup>, Mingxuan Wang<sup>2</sup>, Yuexian Zou<sup>1\*</sup>

<sup>1</sup>ADSPLAB, School of ECE, Peking University, China <sup>2</sup>ByteDance  
chengxx@stu.pku.edu.cn

{dongqianqian, yuefengpeng, tom.ko, wangmingxuan.89}@bytedance.com, zouyx@pku.edu.cn

## ABSTRACT

How to solve the data scarcity problem for end-to-end speech-to-text translation (ST)? It's well known that data augmentation is an efficient method to improve performance for many tasks by enlarging the dataset. In this paper, we propose Mix at three levels for Speech Translation (M<sup>3</sup>ST) method to increase the diversity of the augmented training corpus. Specifically, we conduct two phases of fine-tuning based on a pre-trained model using external machine translation (MT) data. In the first stage of fine-tuning, we mix the training corpus at three levels, including word level, sentence level and frame level, and fine-tune the entire model with mixed data. At the second stage of fine-tuning, we take both original speech sequences and original text sequences in parallel into the model to fine-tune the network, and use Jensen-Shannon divergence to regularize their outputs. Experiments on MuST-C speech translation benchmark and analysis show M<sup>3</sup>ST outperforms current strong baselines and achieves state-of-the-art results on eight directions with an average BLEU of 29.9.

**Index Terms**— Speech Translation, Mix At Three Levels, Data Augmentation

## 1. INTRODUCTION

Speech translation (ST), which has a wide range of applications, is a cross-modal task that the model receives acoustic speech signals and generates text translations in the target language. End-to-end ST system has attracted more and more attention and some work has demonstrated its superiority, even outperforming conventional cascaded systems [1, 2].

One of the major challenges in training an end-to-end ST model is the data scarcity. For example, there are only a few hundreds hours in MuST-C corpus whose size is far smaller than the data of automatic speech recognition (ASR) or machine translation (MT). How to make better use of limited labeled data and other parallel corpora from MT is promising but still a problem. Besides, data augmentation has been proven effective in alleviating this problem.

There are also some work investigating methods of data augmentation for ST, including DropDim [3], SkinAugment [4], STR [5], etc. However, existing work usually makes augmentations from a single level and has not been verified on strong baselines. Our work proposes an augmentation strategy in multiple levels, while demonstrating that there is a superposition of performance when used concurrently. Inspired by MixSpeech [6] for ASR, we introduce Mixup as data augmentation technique for ST and extend the original implementation.

In this paper, we propose M<sup>3</sup>ST, an effective method for speech translation. In order to leverage limited data more effectively, we mix the training corpus at three levels, including word level, sentence level and frame level. At word level, we replace one noun from transcription by its similar word and perform the same operation in its corresponding speech and translation. At sentence level, we concatenate two sentences from different speakers. At frame level, we combine two different speech features with weight  $\lambda$  as the input and use each target sequence to calculate the cross-entropy loss. Then we mix the two losses using the same weight  $\lambda$ . Experimental results show that our method achieves promising performance on the benchmark dataset MuST-C [7], and achieves state-of-the-art results on eight directions regardless of whether using external MT data or not.

## 2. METHOD

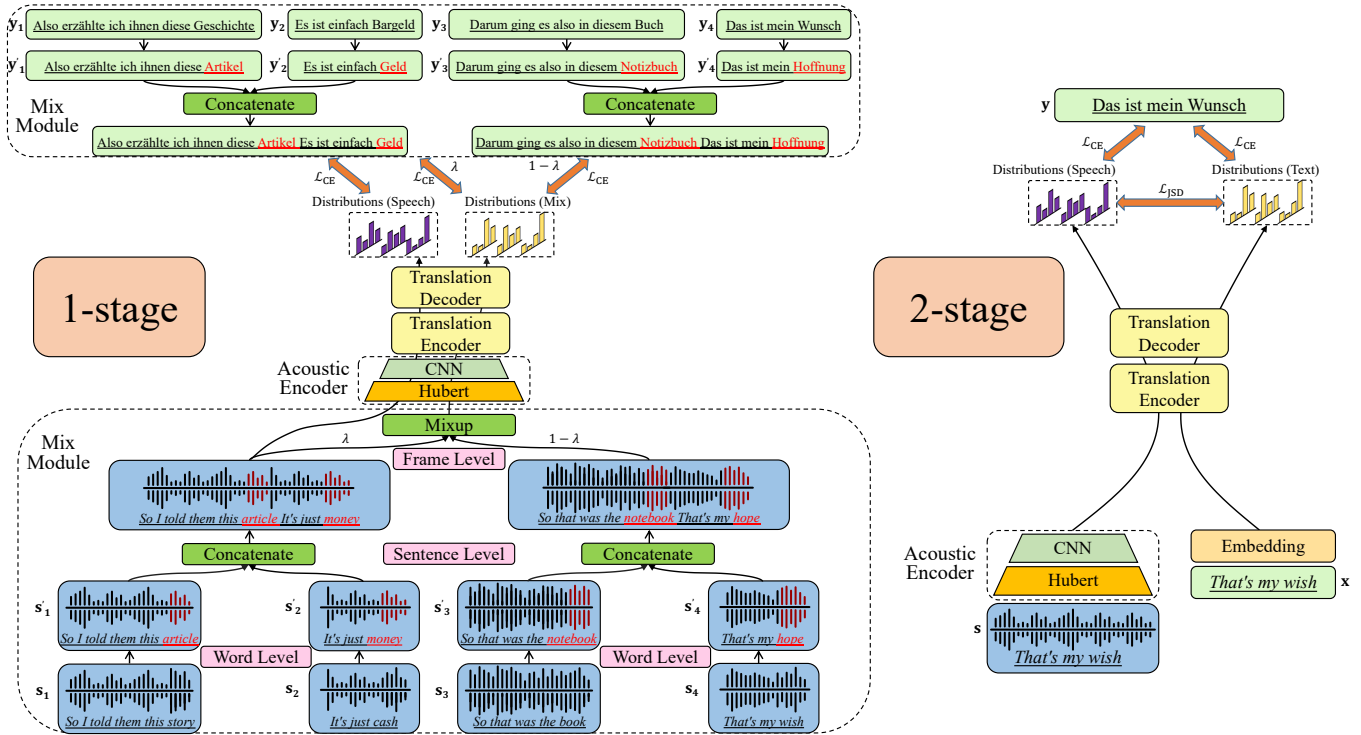
In this section, we first introduce the basic problem formulation of E2E-ST in Section 2.1 and show the overall model architecture in Section 2.2. Then, we introduce our proposed M<sup>3</sup>ST in Section 2.3. Finally we introduce fine-tune with Jensen-Shannon divergence in Section 2.4.

### 2.1. Problem Formulation

The corpus of speech translation usually contains *speech-transcription-translation* triples, which can be denoted as  $\mathcal{D} = \{(s, x, y)\}$ , where  $s$  is the sequence of audio,  $x$  is the corresponding transcription in the source language, and  $y$  is the translation in the target language. An end-to-end speech translation model aims to directly generate translation sequences  $y$  from speech sequences  $s$ , without generating intermediate transcription  $x$ . For the input sequence  $x$ , the

<sup>†</sup> Work was done while at ByteDance.

\* Corresponding author: zouyx@pku.edu.cn.



**Fig. 1:** Overview of our proposed framework. In the 1-stage of fine-tuning, we mix the training corpus at three levels, including word level, sentence level and frame level, and fine-tune the entire model with mixed data. In the 2-stage of fine-tuning, both original speech sequence and transcription sequence are fed into the shared translation module to predict the translation, and we regularize two output predictions with an additional JS Divergence loss.

standard training objective is to optimize the maximum likelihood estimation (MLE) loss of the training data.

$$\mathcal{L}_{CE}(\mathbf{x}, \mathbf{y}) = - \sum_{i=1}^{|\mathbf{y}|} \log p_{\theta}(\mathbf{y}_i | \mathbf{y}_{<i}, \mathbf{h}(\mathbf{x})) \quad (1)$$

## 2.2. Model Architecture

The overall structure of  $M^3ST$  is displayed in Fig.1, including four modules: (a) a mix module (see Section 2.3 for details) which mixes speech input and text input at three levels; (b) an acoustic encoder which extracts low-level features from acoustic input; (c) a translation encoder which extracts further semantic information from the output of acoustic encoder and word embedding; (d) a translation decoder which generates sentence tokens in the target language.

**Acoustic Encoder** Recent work has shown that Hubert [8] has greater potential for improving the performance of end-to-end speech translation than Wav2vec2.0 [9], so we use a pre-trained Hubert without any fine-tuning to extract acoustic representation. In addition, we add two convolutional layers to down sampling the extracted feature, since the length of the sequence of acoustic representations is much bigger than that of the sequence of sub-word embeddings.

**Translation Encoder** Our translation encoder follows the base configuration, which is composed of 6 transformer [10] encoder layers. Each layer includes a self-attention layer, a

feed-forward layer, residual connections and normalization layers. The input of the translation encoder is the output of the acoustic encoder for ST task and the embedding of the transcription for MT task.

**Translation Decoder** Our translation decoder also follows the base configuration, which is composed of 6 transformer [10] decoder layers. Compared with a transformer encoder layer, a transformer decoder layer includes an additional cross-attention layer. The output of the translation encoder is fed here and used to predict the translation.

**Two Stages of Fine-tuning** We first pre-train the translation encoder and translation decoder using parallel *transcription-translation* pairs from both the speech translation corpus and the external MT dataset. Then we fine-tune the entire model for the ST task using the mixed data in the first stage. Finally we fine-tune the model with original ST data without mixing in the second stage.

## 2.3. Mix Module

Mixup technique [11] is applied in many tasks to improve the generalization limitation of empirical risk minimization. The main idea of Mixup is to improve the generalization ability of the model and enhance the robustness of the model. As we mentioned in Section 1, to leverage limited data more effectively, we present  $M^3ST$  to use mix at three levels, including word level, sentence level and frame level.

**Word Level** For a *speech-transcription-translation* triple, denoted as  $\mathcal{D} = \{(\mathbf{s}, \mathbf{x}, \mathbf{y})\}$ , where  $\mathbf{s} = [s_1, s_2, \dots, s_{|\mathbf{s}|}]$ ,  $\mathbf{x} = [x_1, x_2, \dots, x_{|\mathbf{x}|}]$ ,  $\mathbf{y} = [y_1, y_2, \dots, y_{|\mathbf{y}|}]$ , we first randomly choose a noun denoted  $x_i$  from  $\mathbf{x}$  and one of its similar words denoted as  $x'_i$ . Then we get the corresponding speech  $s'_i$  and translation  $y'_i$  of  $x'_i$ . Finally we replace  $s_i$  by  $s'_i$  and  $y_i$  by  $y'_i$ . Note that although we also use word-level alignment between speech and text, we use it extremely infrequently compared to STEMM [12].

**Sentence Level** For two *speech-translation* pairs from different speakers, denoted as  $(s_i, x_i, y_i)$  and  $(s_j, x_j, y_j)$ , we just concatenate them. Through this method, we can enhance the insensitivity of the model to the speaker and thus further improve the performance of the model.

**Frame Level** We set a combination weight, denoted as  $\lambda$  and combine two different speech features with weight  $\lambda$  as the input before they were fed into Hubert. To maintain a balanced mix, we ensure that every mixed pair with weight  $\lambda$  will be mixed with weight  $1 - \lambda$ . Then we use each target sequence to calculate the cross-entropy loss and mix the two losses using the same weight  $\lambda$ .

$$\begin{aligned} \mathbf{s}_{mix} &= \lambda \mathbf{s}_i + (1 - \lambda) \mathbf{s}_j \\ \mathcal{L}_i &= \mathcal{L}_{CE}(\mathbf{s}_{mix}, \mathbf{y}_i) \\ \mathcal{L}_j &= \mathcal{L}_{CE}(\mathbf{s}_{mix}, \mathbf{y}_j) \\ \mathcal{L}_{MIX}(\mathbf{s}, \mathbf{y}) &= \lambda \mathcal{L}_i + (1 - \lambda) \mathcal{L}_j \end{aligned} \quad (2)$$

where  $\mathbf{s}_{mix}$  is the mixed speech sequence and  $\mathcal{L}_{MIX}$  is the mixed loss we propose. In the first stage of the fine-tuning, our training corpus consists of two parts, including original parallel *speech-translation* pairs which more focus on the ST task itself and the mixed parallel *speech-translation* pairs devoted to improve the generalization ability and enhance the robustness of the model. With the cross-entropy losses of two forward passes, the final training objective of the first stage of fine-tuning is as follows:

$$\mathcal{L}_1(\mathbf{s}, \mathbf{y}) = \mathcal{L}_{CE}(\mathbf{s}, \mathbf{y}) + \mathcal{L}_{MIX}(\mathbf{s}, \mathbf{y}) \quad (3)$$

#### 2.4. Fine-tune With Jensen-Shannon Divergence

In the second stage of fine-tuning, to enhance the knowledge transfer from MT task to ST task, we take both parallel *speech-translation* pairs for ST task and *transcription-translation* pairs for MT task as input. Besides, we introduce Jensen-Shannon Divergence (JSD) to regularize output predictions from the two tasks, which is

$$\mathcal{L}_{JSD}(\mathbf{s}, \mathbf{x}, \mathbf{y}) = \sum_{i=1}^{|\mathbf{y}|} \text{JSD} \{p_{\theta}(\mathbf{y}_i | \mathbf{y}_{<i}, \mathbf{h}(\mathbf{s})) \parallel p_{\theta}(\mathbf{y}_i | \mathbf{y}_{<i}, \mathbf{h}(\mathbf{x}))\} \quad (4)$$

where  $\mathbf{h}(\cdot)$  is the contextual representation output by the translation encoder,  $p_{\theta}(\mathbf{y}_i | \mathbf{y}_{<i}, \mathbf{h}(\mathbf{s}))$  is the probability distribution of the  $i$ -th target token predicted by the translation decoder of the speech sequence  $\mathbf{s}$ , and  $p_{\theta}(\mathbf{y}_i | \mathbf{y}_{<i}, \mathbf{h}(\mathbf{x}))$  is that of the transcription sequence  $\mathbf{x}$ . Our training loss

of the second stage of fine-tuning consists of the following elements.

$$\mathcal{L}_2(\mathbf{s}, \mathbf{y}) = \mathcal{L}_{CE}(\mathbf{s}, \mathbf{y}) + \mathcal{L}_{CE}(\mathbf{x}, \mathbf{y}) + \mathcal{L}_{JSD}(\mathbf{s}, \mathbf{x}, \mathbf{y}) \quad (5)$$

### 3. EXPERIMENTS

#### 3.1. Datasets

**ST datasets** We conduct our experiments on all the translation directions in a multilingual speech translation corpus called MuST-C [7], which contains translations from English (En) to German (De), Spanish (Es), French (Fr), Italian (It), Dutch (NI), Portuguese (Pt), Romanian (Ro) and Russian (Ru). Each direction includes a triplet of speech, transcription, and translation. As a benchmark ST dataset, the sizes of the eight directional TED talks range from 385 hours (Pt) to 504 hours (Es). In all experiments, model selection is based on the dev set and the final results are reported on the test-COMMON set.

**MT datasets** We also use external MT datasets to pre-train our translation model, including WMT for En-De/Es/Fr/Ro/Ru directions, and OPUS100 datasets [20] for En-It/NI/Pt directions, as an expanded setup. For En-De, the size of WMT16 data is 4.6M.

#### 3.2. Experimental setups

**Pre-processing** We use the raw 16-bit 16kHz mono-channel audios as speech input and remove the punctuation of the transcriptions. Both transcriptions and translations are case-sensitive. For each translation direction, a unigram SentencePiece<sup>1</sup>[21] model is used to learn a vocabulary whose size is 10K on the text data from ST dataset. We use the vocabulary to partition the text in the ST and MT corpora into subword units. Considering the grammatical differences between languages, we only replace nouns with the practice mentioned above. We first collect words from transcriptions, then use CoreNLP<sup>2</sup> to extract nouns from the words. To get word-level *speech-transcription* alignment pairs, we use Montreal Forced Aligner<sup>3</sup>[22] toolkit. To get word-level *transcription-translation* alignment pairs, we use fast\_align [23] toolkit. We also use word2vec [24] to search the five closest words of each noun from the transcriptions.

**Training Details** We use Adam optimizer [25] to optimize the parameters in our model. We use EarlyStop which can stop training if the loss on the dev set doesn't decrease within ten epochs, with no more than 16M source audio frames per batch. The mix ratio  $\lambda$  is set to 0.4.

**Baseline Models** We compare our model with strong end-to-end ST models, including XSTNet [1], TaskAware [19], ConST [17], STEMM [12], Fairseq ST [14], SATE [18], TDA [15], Revisit ST [16]. We also implement a strong baseline W2V2-Transformer based on Hubert which has the

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

<sup>2</sup><https://downloads.cs.stanford.edu/nlp/software/stanford-corenlp-4.5.1.zip>

<sup>3</sup><https://github.com/MontrealCorpusTools/Montreal-Forced-Aligner>

**Table 1:** Case-sensitive detokenized BLEU scores on MuST-C tst-COMMON set. “Speech” denotes unlabeled speech data. “Text” means unlabeled text data. † uses external 40M OpenSubtitles [13] MT data. \* and \*\* mean the improvements over W2V2-Transformer baseline are statistically significant ( $p < 0.05$  and  $p < 0.01$ , respectively).

Models	External Data				De	Es	Fr	It	BLEU				Avg.
	Speech	Text	ASR	MT					Nl	Pt	Ro	Ru	
<i>Pre-train w/o external MT data</i>													
Fairseq S2T’20 [14]	×	×	×	×	22.7	27.2	32.9	22.7	27.3	28.1	21.9	15.3	24.8
XSTNet’21 [1]	✓	×	×	×	25.5	29.6	36.0	25.5	30.0	31.3	25.1	16.9	27.5
TDA’21 [15]	×	×	×	×	25.4	29.6	36.1	25.1	29.6	31.1	23.9	16.4	27.2
Revisit ST’22 [16]	×	×	×	×	23.0	28.0	33.5	23.5	27.1	28.2	23.0	15.6	25.2
STEMM’22 [12]	✓	×	×	×	25.6	30.3	36.1	25.6	30.1	31.0	24.3	17.1	27.5
ConST’22 [17]	✓	×	×	×	25.7	30.4	36.8	26.3	30.6	32.0	24.8	17.3	28.0
W2V2-Transformer	✓	×	×	×	24.3	29.6	35.2	25.1	29.1	30.3	23.4	16.5	26.7
M <sup>3</sup> ST	✓	×	×	×	<b>26.4**</b>	<b>31.0**</b>	<b>37.2**</b>	<b>26.6**</b>	<b>30.9**</b>	<b>32.8**</b>	<b>25.4**</b>	<b>18.3**</b>	<b>28.6</b>
<i>Pre-train w/ external MT data</i>													
XSTNet’21 [1]	✓	×	×	✓	27.1	30.8	38.0	26.4	31.2	32.4	25.7	18.5	28.8
SATE’21 [18]	×	×	✓	✓	28.1 <sup>†</sup>	-	-	-	-	-	-	-	-
TaskAware’21 [19]	×	×	✓	✓	28.9	-	-	-	-	-	-	-	-
STEMM’22 [12]	✓	×	×	✓	28.7	31.0	37.4	25.8	30.5	31.7	24.5	17.8	28.4
ConST’22 [17]	✓	×	×	✓	28.3	32.0	38.3	27.2	31.7	33.1	25.6	18.9	29.4
W2V2-Transformer	✓	×	×	✓	27.2	30.2	36.8	25.6	29.7	30.9	24.1	17.5	27.8
M <sup>3</sup> ST	✓	×	×	✓	<b>29.3**</b>	<b>32.4**</b>	<b>38.5*</b>	<b>27.5**</b>	<b>32.5**</b>	<b>33.4**</b>	<b>25.9**</b>	<b>19.3**</b>	<b>29.9</b>

same model architecture as our proposed M<sup>3</sup>ST and is pre-trained in the same way, but we don’t conduct the first stage of fine-tuning with the mix module.

### 3.3. Main Results

Table 1 shows experiment results of different models on MuST-C. Without external MT data, our model gains a large improvement of 0.6 BLEU (average over 8 directions) over the previous best models, and even outperforms W2V2-Transformer with external MT data except En-De, which proves the effectiveness of mix module. When external MT data is introduced, our method also outperforms SOTA by an average of 0.5 BLEU, which further demonstrates the potential of our M<sup>3</sup>ST method.

### 3.4. Analysis

#### 3.4.1. Ablation Study

To study the effectiveness of mixing at each level in our approach, we perform ablation experiments on En-De with external MT data, whose results are shown in Table 2. We remove the mix at word level, at sentence level and at frame level in turn. From the results, we can observe the absence of mix at frame level leads to 0.5 BLEU drops, which is the largest of the three levels. This indicates that the mix at frame level can improve the performance of model significantly by introducing another speech, acting like a contrast signal to force the ST model to better translate the corresponding speech rather than being misled by other speech signals.

#### 3.4.2. Effect of the Layer of Mixing at Frame Level

We conduct a series of experiments to vary the layer of mixing at frame level, and the results are shown in Table 3. We can observe that the best BLEU can be obtained when we mix the speech at layer 0, i.e., when we mix the embedding of

**Table 2:** Ablation results on En-De with external MT data.

Models	BLEU
M <sup>3</sup> ST	29.3
- word	29.1
- word - sentence	28.8
- word - sentence - frame	28.3

them. We suppose the reason for this is that more acoustic information has been lost when mixing at a deeper level.

**Table 3:** Results of mixing at different layer on En-De with external MT data.

Layer	BLEU
0	29.3
2	29.1
4	28.6
6	28.3

## 4. CONCLUSION

In this paper, we propose the Mix at three levels for Speech Translation (M<sup>3</sup>ST) method to mix the training corpus at three levels, including word level, sentence level and frame level. Experiments and analysis demonstrate the effectiveness of our proposed method, which can leverage limited parallel pairs more efficiently. In the future, we will explore how to further solve the problem of data scarcity for ST task.

## 5. ACKNOWLEDGEMENTS

This paper was partially supported by Shenzhen Science & Technology Research Program (No: GXWD2020123116580 7007-20200814115301001) and NSFC (No: 62176008).

## 6. REFERENCES

- [1] R. Ye, M. Wang, and L. Li, "End-to-end speech translation via cross-modal progressive training," in *Proc. of INTERSPEECH*, 2021.
- [2] C. Han, M. Wang, H. Ji, and L. Li, "Learning shared semantic space for speech-to-text translation," in *Proc. of ACL Findings*, 2021.
- [3] H. Zhang, D. Qu, K. Shao, and X. Yang, "Dropdim: A regularization method for transformer networks," *IEEE Signal Processing Letters*, 2022.
- [4] A. D. McCarthy, L. Puzon, and J. Pino, "Skinaugment: Auto-encoding speaker conversions for automatic speech translation," in *Proc. of ICASSP*, 2020.
- [5] T. K. Lam, S. Schamoni, and S. Riezler, "Sample, translate, recombine: Leveraging audio alignments for data augmentation in end-to-end speech translation," in *Proc. of ACL*, 2022.
- [6] L. Meng, J. Xu, X. Tan, J. Wang, T. Qin, and B. Xu, "Mixspeech: Data augmentation for low-resource automatic speech recognition," in *Proc. of ICASSP*, 2021.
- [7] M. A. Di Gangi, R. Cattoni, L. Bentivogli, M. Negri, and M. Turchi, "MuST-C: a Multilingual Speech Translation Corpus," in *Proc. of NAACL*, 2019.
- [8] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [9] R. Ye, C. Zhao, T. Ko, C. Meng, T. Wang, M. Wang, and J. Cao, "Gigast: A 10,000-hour pseudo speech translation corpus," *arXiv preprint arXiv:2204.03939*, 2022.
- [10] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. of NeurIPS*, 2017.
- [11] H. Zhang, M. Cissé, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *Proc. of ICLR*, 2018.
- [12] Q. Fang, R. Ye, L. Li, Y. Feng, and M. Wang, "STEMM: Self-learning with speech-text manifold mixup for speech translation," in *Proc. of ACL*, 2022.
- [13] P. Lison and J. Tiedemann, "OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles," in *Proc. LREC*, 2016.
- [14] C. Wang, Y. Tang, X. Ma, A. Wu, D. Okhonko, and J. Pino, "Fairseq S2T: Fast speech-to-text modeling with fairseq," in *Proc. of AACL*, 2020.
- [15] Y. Du, Z. Zhang, W. Wang, B. Chen, J. Xie, and T. Xu, "Regularizing end-to-end speech translation with triangular decomposition agreement," in *Proc. of AAAI*, 2022.
- [16] B. Zhang, B. Haddow, and R. Sennrich, "Revisiting end-to-end speech-to-text translation from scratch," *arXiv preprint arXiv:2206.04571*, 2022.
- [17] R. Ye, M. Wang, and L. Li, "Cross-modal contrastive learning for speech translation," in *Proc. of NAACL*, 2022.
- [18] C. Xu, B. Hu, Y. Li, Y. Zhang, S. Huang, Q. Ju, T. Xiao, and J. Zhu, "Stacked acoustic-and-textual encoding: Integrating the pre-trained models into speech translation encoders," in *Proc. of ACL*, 2021.
- [19] S. Indurthi, M. A. Zaidi, N. K. Lakumarapu, B. Lee, H. Han, S. Ahn, S. Kim, C. Kim, and I. Hwang, "Task aware multi-task learning for speech to text tasks," in *Proc. of ICASSP*, 2021.
- [20] B. Zhang, P. Williams, I. Titov, and R. Sennrich, "Improving massively multilingual neural machine translation and zero-shot translation," in *Proc. of ACL*, 2020.
- [21] T. Kudo and J. Richardson, "SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing," in *Proc. of EMNLP*, 2018.
- [22] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using kaldif," in *Proc. of INTERSPEECH*, 2017.
- [23] C. Dyer, V. Chahuneau, and N. A. Smith, "A simple, fast, and effective reparameterization of IBM model 2," in *Proc. of NAACL*, 2013.
- [24] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
- [25] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.