

CMOT: Cross-modal Mixup via Optimal Transport for Speech Translation

Yan Zhou¹, Qingkai Fang^{1,2}, Yang Feng^{1,2*}

¹Key Laboratory of Intelligent Information Processing
Institute of Computing Technology, Chinese Academy of Sciences (ICT/CAS)

²University of Chinese Academy of Sciences, Beijing, China
2001zhouyan@gmail.com, {fangqingkai21b, fengyang}@ict.ac.cn

Abstract

End-to-end speech translation (ST) is the task of translating speech signals in the source language into text in the target language. As a cross-modal task, end-to-end ST is difficult to train with limited data. Existing methods often try to transfer knowledge from machine translation (MT), but their performances are restricted by the modality gap between speech and text. In this paper, we propose **Cross-modal Mixup via Optimal Transport (CMOT)** to overcome the modality gap. We find the alignment between speech and text sequences via optimal transport and then mix up the sequences from different modalities at a token level using the alignment. Experiments on the MuST-C ST benchmark demonstrate that CMOT achieves an average BLEU of 30.0 in 8 translation directions, outperforming previous methods. Further analysis shows CMOT can adaptively find the alignment between modalities, which helps alleviate the modality gap between speech and text. ¹

1 Introduction

Speech translation (ST) is the task to translate speech signals in one language into text in another. Conventional ST systems work in a cascaded mode (STENTIFORD and STEER, 1988; Waibel et al., 1991), by combining automatic speech recognition (ASR) and machine translation (MT), which may suffer from error propagation and high latency. To solve these problems, end-to-end ST is proposed (Bérard et al., 2016; Duong et al., 2016). End-to-end ST directly generates the target text directly from the source audio without the intermediate transcript, which has drawn more attention in recent years (Vila et al., 2018; Salesky et al., 2019; Di Gangi et al., 2019b,c; Inaguma et al., 2020;

Wang et al., 2020a; Zhao et al., 2021a; Dinh et al., 2022; Duquenne et al., 2022a).

End-to-end ST has to perform cross-lingual and cross-modal transformation simultaneously, hence bringing more challenges than ASR or MT. Unfortunately, it is harder to collect parallel speech-to-text pairs than text-to-text pairs, which makes ST suffer from the under-fitting problem. To solve the data shortage problem, some researchers propose methods to introduce text-to-text parallel data of MT into ST, including pretraining (Alinejad and Sarkar, 2020; Zheng et al., 2021; Xu et al., 2021), multi-task learning (Le et al., 2020; Vydana et al., 2021; Tang et al., 2021), knowledge distillation (Liu et al., 2019; Gaido et al., 2020; Inaguma et al., 2021), etc. In this way, the knowledge of MT can be transferred to ST, helping learn a better cross-lingual transformation.

Although these methods of transfer learning have shown improvements in the translation quality, ST is joint learning of cross-lingual and cross-modal transformation, and the effectiveness of transfer learning is also based on the assumption that the cross-modal source will be projected to the common representation space. Therefore, the mixup strategy is employed to further improve ST by reducing the gap between the speech and text representation spaces. Previous methods with mixup strategy usually perform word-level mixup between aligned tokens in different modalities (Fang et al., 2022), hence needing the help of external alignment tools. However, the alignment tools are not always available and require large amounts of additional annotated data for training, so how to obtain such alignments remains a problem for most languages.

Illuminated by the application of the optimal transport theory in finding the alignment between languages or modalities in recent years (Chen et al., 2020; Gu and Feng, 2022), we propose **Cross-modal Mixup via Optimal Transport (CMOT)**. We

* Corresponding author: Yang Feng.

¹Code is publicly available at <https://github.com/ictnlp/CMOT>.

use optimal transport to adaptively find the alignment between speech and text, and use this alignment to achieve token-level mixup. The mixup sequence will serve as a medium between the speech sequence and text sequence to realize cross-modal knowledge transfer. Experiments on the MuST-C dataset show that CMOT achieves an average BLEU of 30.0. In addition, we prove that optimal transport can help find cross-modal alignments and CMOT can help alleviate the modality gap in ST.

2 Method

In this section, we will describe our proposed Cross-modal Mixup via Optimal Transport (CMOT). We find the alignment between speech and text sequences via OT, and mix up the two unimodal sequences to get a cross-modal mixed sequence. We predict the translation with all these sequences and regularize their outputs. Figure 1 illustrates the overview of our proposed method.

2.1 Problem Formulation

Generally, a speech translation corpus contains triplets of source speech s , its transcript x and translation y , which can be denoted as $\mathcal{D} = \{(s, x, y)\}$. Given the corpus, an end-to-end ST system directly converts speech signals s into text translation y without generating the intermediate transcript x .

2.2 Model Architecture

Our model consists of four modules: the speech encoder, the text embedding layer, the translation encoder and the translation decoder. We first get parallel speech and text sequences with the speech encoder and the text embedding layer, and feed them into the translation encoder separately. Then we mix up the two encoder outputs with CMOT, and feed both the unimodal sequences and cross-modal mixed sequence to predict the translation.

Speech Encoder The speech encoder extracts the low-level features from the raw speech input. We use HuBERT (Hsu et al., 2021) model stacked with a sub-sampler as the speech encoder. The HuBERT model is a self-supervised pretrained speech model and can improve the performance of ST systems as shown in recent works (Zhang et al., 2022a; Feng et al., 2022). It consists of a convolutional feature extractor and a BERT (Devlin et al., 2019) style encoder. The sub-sampler consists of two convolution layers and is designed to reduce the length of the speech sequence.

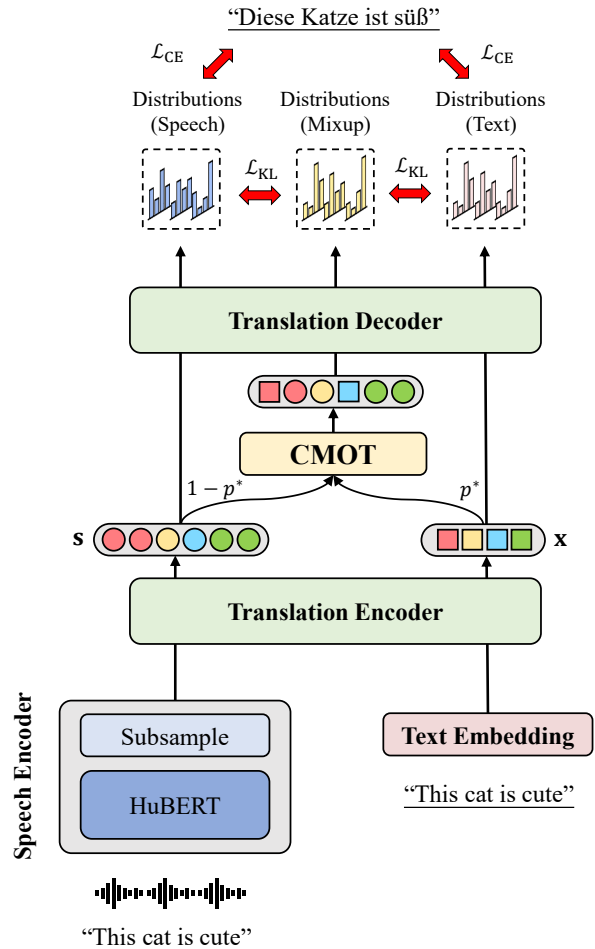


Figure 1: Overview of our proposed architecture with CMOT.

Text Embedding The text embedding layer embeds tokenized text into a sequence of embeddings.

Translation Encoder The joint translation encoder receives either speech or text embeddings and learns further semantic information. It is composed of N_e transformer (Vaswani et al., 2017) encoder layers.

Translation Decoder The joint translation decoder receives either unimodal sequences of speech and text or the cross-modal mixed sequence and predicts the translation y . It is composed of N_d transformer decoder layers.

2.3 Cross-modal Mixup via Optimal Transport (CMOT)

We mix up the speech sequence and text sequence with CMOT. For a speech sequence \mathbf{H}^s and the corresponding text sequence \mathbf{H}^x , we first get the alignment A between them via optimal transport (OT), and then mix up \mathbf{H}^s and \mathbf{H}^x with the help of A to get a mixed sequence \mathbf{H}' .

Optimal Transport The Optimal Transport (OT)

problem is a classical mathematical problem and is nowadays used as a geometric tool to compare probability distributions (Peyré et al., 2019). It focuses on the scheme with the minimum cost of transferring one distribution to another under a given transfer cost. Consider the following discrete transport problem: for two independent probability distributions P and Q ,

$$\begin{aligned} P &= \{(w_i, m_i)\}_{i=1}^n, \quad s.t. \sum_i m_i = 1; \\ Q &= \{(\hat{w}_j, \hat{m}_j)\}_{j=1}^{\hat{n}}, \quad s.t. \sum_j \hat{m}_j = 1, \end{aligned} \quad (1)$$

where each point $w_i \in \mathbb{R}^d$ has a weight $m_i \in [0, \infty)$. Given a transfer cost function $c(w_i, \hat{w}_j)$ and denoting the mass transported from w_i to \hat{w}_j by \mathbf{T}_{ij} , the transport cost can be defined as:

$$\begin{aligned} \mathcal{D}(P, Q) &= \min_{\mathbf{T} \geq 0} \sum_{i,j} \mathbf{T}_{ij} c(w_i, \hat{w}_j) \\ s.t. \sum_{j=1}^{\hat{n}} \mathbf{T}_{ij} &= m_i, \forall i \in \{1, \dots, n\}, \\ \sum_{i=1}^n \mathbf{T}_{ij} &= \hat{m}_j, \forall j \in \{1, \dots, \hat{n}\}. \end{aligned} \quad (2)$$

Relaxed OT with Window Strategy If we regard the speech sequence and the text sequence as two independent distributions, we can use OT to measure the distance between them. For a speech sequence $\mathbf{H}^s = (h_1^s, \dots, h_n^s)$, a text sequence $\mathbf{H}^x = (h_1^x, \dots, h_{\hat{n}}^x)$, and a cost function c , we define the transport problem as:

$$\begin{aligned} \mathcal{D}(\mathbf{H}^s, \mathbf{H}^x) &= \min_{\mathbf{T} \geq 0} \sum_{i,j} \mathbf{T}_{ij} c(h_i^s, h_j^x) \\ s.t. \sum_{j=1}^{\hat{n}} \mathbf{T}_{ij} &= m_i^s, \forall i \in \{1, \dots, n\}, \\ \sum_{i=1}^n \mathbf{T}_{ij} &= m_j^x, \forall j \in \{1, \dots, \hat{n}\}. \end{aligned} \quad (3)$$

Here we use the Euclidean distance as the cost function c , and norm as the mass m^s and m^x inspired by Schakel and Wilson (2015) and Yokoi et al. (2020)'s discovery that important tokens have larger norms. The solutions to the OT problem include some accurate algorithms like Sinkhorn (Cun- turi, 2013) algorithm and IPOT (Xie et al., 2020), which will bring great time complexity. Kusner et al. (2015) proposed a relaxed moving distance

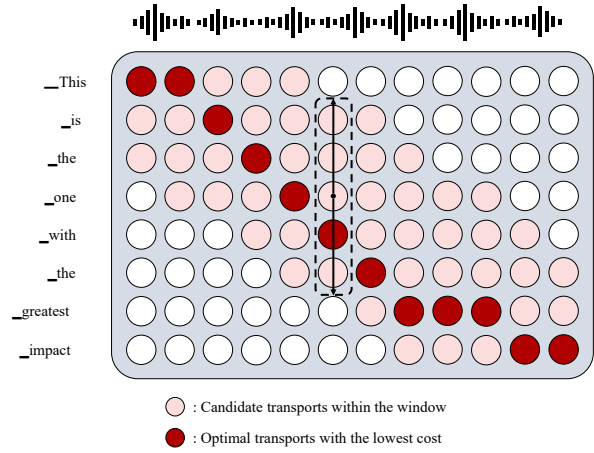


Figure 2: An example of OT from speech sequence to text sequence with window strategy. The x-axis is the speech sequence and the y-axis is the text sequence. White circles represent the transports that cannot be selected, pink ones represent the candidate transports and red ones represent the selected optimal transports. The dashed box represents the window, and we use $W = 2$ as an example here.

that removes one of the two constraints to obtain a lower bound of the accurate solution. Following this work, we remove the second constraint, and our transport problem thus becomes:

$$\begin{aligned} \mathcal{D}^*(\mathbf{H}^s, \mathbf{H}^x) &= \min_{\mathbf{T} \geq 0} \sum_{i,j} \mathbf{T}_{ij} c(h_i^s, h_j^x) \\ s.t. \sum_{j=1}^{\hat{n}} \mathbf{T}_{ij} &= m_i^s, \forall i \in \{1, \dots, n\}. \end{aligned} \quad (4)$$

This relaxed problem yields a lower bound of the original problem. We find that this relaxed OT improves the training speed without degrading the performance, as shown in Section 4.1. Now the optimal solution for each speech token h_i^s is to move all its mass to the closest text token h_j^x , so the transportation matrix becomes:

$$\mathbf{T}_{ij}^* = \begin{cases} m_i^s & \text{if } j = \arg \min_j c(h_i^s, h_j^x) \\ 0 & \text{otherwise} \end{cases}, \quad (5)$$

This transport matrix \mathbf{T}^* implies the alignment between \mathbf{H}^s and \mathbf{H}^x . We define the cross-modal alignment as $\mathbf{A} = (a_1, \dots, a_n)$, where

$$a_i = \operatorname{argmax}_{1 \leq j \leq \hat{n}} \mathbf{T}_{ij}^*. \quad (6)$$

Considering the monotonicity and locality of the alignment from the speech sequence to the text sequence, we apply a window strategy for this alignment. We constrain the index of the aligned target

to the neighborhood of the diagonal. The window size is set as a constant W both to left and to right:

$$a_i = \underset{\max(1, \lambda i - W) \leq j \leq \min(\hat{n}, \lambda i + W)}{\operatorname{argmax}} \mathbf{T}_{ij}^*, \quad (7)$$

$$\lambda = \hat{n}/n.$$

Figure 2 shows how the relaxed OT with window strategy works. Through the above steps, we get the alignment A in an adaptive manner.

Cross-modal Mixup via OT With the help of the alignment A , we will be able to mixup the speech and text sequences at the token level. For a mixup probability p^* , we generate the mixup sequence $\mathbf{H}' = (h'_1, \dots, h'_n)$ by:

$$h'_i = \begin{cases} h_i^s, & p > p^*, \\ h_{a_i}^x, & p \leq p^*. \end{cases} \quad (8)$$

where p is sampled from a uniform distribution $U(0, 1)$. We feed \mathbf{H}^s , \mathbf{H}^x , \mathbf{H}' into the joint transformer to predict the translation. Note that the positions of OT and mixup can be different. Our scheme is calculating OT before the translation encoder and performing mixup with the encoder outputs, which brings better performance in the experiment. See more details in Section 4.3.

2.4 Training Strategy

We follow the paradigm of pretraining-finetuning to train our model.

Pretraining We pretrain the text embedding and the joint transformer with an MT task. The training objective during this phase is the cross entropy loss:

$$\mathcal{L}_{\text{MT}} = -\mathbb{E}_{\mathbf{x}, \mathbf{y}} \log P(\mathbf{y}|\mathbf{x}). \quad (9)$$

Multitask Finetuning During finetuning, we train the model with multiple tasks. We add the cross entropy loss of ST and MT to the final loss:

$$\begin{aligned} \mathcal{L}_{\text{ST}} &= -\mathbb{E}_{(\mathbf{s}, \mathbf{x}, \mathbf{y}) \in \mathcal{D}} \log P(\mathbf{y}|\mathbf{s}), \\ \mathcal{L}_{\text{MT}} &= -\mathbb{E}_{(\mathbf{s}, \mathbf{x}, \mathbf{y}) \in \mathcal{D}} \log P(\mathbf{y}|\mathbf{x}). \end{aligned} \quad (10)$$

Meanwhile, we regularize the predictions of mixup and speech sequences by minimizing the Kullback-Leibler (KL) Divergence between the token-level probability distributions of them, and we do the same for the predictions of mixup and text sequences. For a speech prediction distribution S , a text prediction distribution T and mixup prediction distribution M , we have:

$$\begin{aligned} \mathcal{L}_{\text{KL}}(M, S) &= \frac{1}{2} (\mathbb{D}_{\text{KL}}(M||S) + \mathbb{D}_{\text{KL}}(S||M)), \\ \mathcal{L}_{\text{KL}}(M, T) &= \frac{1}{2} (\mathbb{D}_{\text{KL}}(M||T) + \mathbb{D}_{\text{KL}}(T||M)). \end{aligned} \quad (11)$$

Hence the final training objective is:

$$\mathcal{L} = \mathcal{L}_{\text{ST}} + \mathcal{L}_{\text{MT}} + \lambda \mathcal{L}_{\text{KL}}(M, S) + \lambda \mathcal{L}_{\text{KL}}(M, T), \quad (12)$$

where λ is the weight to control \mathcal{L}_{KL} .

3 Experiments

We conduct experiments under two settings, the base and the expanded. The former only uses ST data, while the latter uses external MT data additionally. We will describe the datasets we use, our experimental setups and our main results in this section.

3.1 Datasets

ST Datasets We conduct experiments on the MuST-C² (Di Gangi et al., 2019a) dataset, a widely used benchmark for ST. It consists of translations from English (En) to 8 languages: German (De), French (Fr), Russian (Ru), Spanish (Es), Italian (It), Romanian (Ro), Portuguese (Pt) and Dutch (Nl). For each direction, it comprises at least 385 hours of audio recorded from TED Talks. We use dev set for validation and tst-COMMON set for evaluation.

External MT Datasets We incorporate external MT data to pretrain the text embedding and the joint transformer. For directions of En-De, En-Fr, En-Ru, En-Es and En-Ro we use data from WMT (Bojar et al., 2016). For En-It, En-Pt and En-Nl we use data from OPUS100³ (Zhang et al., 2020).

Table 2 shows the detailed statistics of the MuST-C, WMT and OPUS100 datasets we use.

3.2 Experimental Setups

Model Configurations For the speech encoder, we use HuBERT⁴ (Hsu et al., 2021) base model, which is pretrained on LibriSpeech (Panayotov et al., 2015) with no finetuning and is one of the state-of-the-art pretrained audio models. Following the usual practice for ST, we stack two convolution layers of kernel size 5, stride size 2, padding 2 and hidden size 1024 after the HuBERT model as a sub-sampler. The text embedding is an embedding layer of dimension 512. For the joint transformer, we use $N_e = 6$ encoder layers and $N_d = 6$ decoder layers with 512 hidden units, 2048 feed-forward hidden units and 8 attention heads. We use fairseq⁵ (Ott et al., 2019) for implementation.

²Here we refer to v1.0. <https://ict.fbk.eu/must-c/>

³<http://opus.nlpl.eu/opus-100.php>

⁴<https://github.com/facebookresearch/fairseq/tree/main/examples/hubert>

⁵<https://github.com/facebookresearch/fairseq/>

Models	Speech Pretraining	BLEU								
		En-De	En-Fr	En-Ru	En-Es	En-It	En-Ro	En-Pt	En-Nl	Avg.
Base setting (w/o external MT data)										
Fairseq ST (Wang et al., 2020a)	×	22.7	32.9	15.3	27.2	22.7	21.9	28.1	27.3	24.8
XSTNet (Ye et al., 2021)	✓	25.5	36.0	16.9	29.6	25.5	25.1	31.3	30.0	27.5
STEMM (Fang et al., 2022)	✓	25.6	36.1	17.1	30.3	25.6	24.3	31.0	30.1	27.5
ConST (Ye et al., 2022)	✓	25.7	36.8	17.3	30.4	26.3	24.8	32.0	30.6	28.0
HuBERT-Transformer	✓	25.4	36.5	16.8	30.5	26.0	24.2	31.7	29.9	27.6
CMOT	✓	27.0**	37.3**	17.9**	31.1**	26.9**	25.3**	32.7**	31.2**	28.7
Expanded setting (w/ external MT data)										
Chimera (Han et al., 2021)	✓	27.1	35.6	17.4	30.6	25.0	24.0	30.2	29.2	27.4
XSTNet (Ye et al., 2021)	✓	27.1	38.0	18.5	30.8	26.4	25.7	32.4	31.2	28.8
STEMM (Fang et al., 2022)	✓	28.7	37.4	17.8	31.0	25.8	24.5	31.7	30.5	28.4
ConST (Ye et al., 2022)	✓	28.3	38.3	18.9	32.0	27.2	25.6	33.1	31.7	29.4
HuBERT-Transformer	✓	27.5	38.3	18.8	32.5	26.9	25.3	32.3	30.8	29.1
CMOT	✓	29.0**	39.5**	19.2**	32.8*	27.5**	26.0**	33.5**	32.1**	30.0

Table 1: BLEU scores on MuST-C tst-COMMON set. “Speech Pretraining” denotes using pretrained speech models. *, ** indicate the improvements over HuBERT-Transformer are statistically significant with $p < 0.05$ and $p < 0.01$.

En→	ST (MuST-C)		MT	
	hours	#sents	name	#sents
De	408	234K	WMT16	4.6M
Fr	492	280K	WMT14	40.8M
Ru	489	270K	WMT16	2.5M
Es	504	270K	WMT13	15.2M
It	465	258K	OPUS100	1.0M
Ro	432	240K	WMT16	0.6M
Pt	385	211K	OPUS100	1.0M
Nl	442	253K	OPUS100	1.0M

Table 2: Statistics of the datasets.

Pre-processing For speech input, we use the 16-bit 16kHz mono-channel raw audio, and we filter out samples with frames greater than 480k or less than 1k to ensure training efficiency. For transcripts and translations, we tokenize them using a unigram SentencePiece⁶ (Kudo and Richardson, 2018) model with a vocabulary size of 10k, which is shared by the source language and the target language.

Training Our training process follows the pretraining-finetuning paradigm. We pretrain the text embedding layer, the translation encoder and decoder with an MT task. For the base setting, we only use the transcript-translation pairs with a learning rate of $2e-3$, and warm-up steps of 8k. For the expanded setting we first pretrain them on external MT data with a learning rate of $7e-4$ and warm-up steps of 8k, and then pretrain them on transcript-translation pairs with a learning rate of $1e-4$ and warm-up steps of 4k. For both settings,

we allow at most 33k input tokens per batch.

We set the learning rate to $1e-4$ and warm-up steps to 10k to finetune the whole model. We train the entire model for 60k updates with a batch of 16M audio frames. We use an Adam optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.9$, $\beta_2 = 0.98$, dropout of 0.1 and the label smoothing value of 0.1. For CMOT, we set the weight λ of KL loss to 2.0 (see more in Appendix 4.3) and the mixup probability p^* to 0.2 (see more in Section 4.3), and we use the window strategy for OT by default with window size $W = 10$. All models are trained on 4 Nvidia GeForce RTX 3090 GPUs.

Inference We average the checkpoints of the last 10 epochs for evaluation. We use beam search with a beam size of 8. We evaluate case-sensitive detokenized BLEU on MuST-C tst-COMMON set using sacreBLEU⁷ (Post, 2018). We use sacreBLEU to test the significance of the results.

Baselines We compare our method with the Fairseq ST (Wang et al., 2020a) baseline implemented with the same framework (Ott et al., 2019). We implement HuBERT-Transformer as a strong baseline, which has the same architecture as our method. The only difference is that HuBERT-Transformer takes speech as input and performs the ST task during fine-tuning. We train this model for 40k steps, except for the En-Ru direction where we use a training step of 100k due to specific language complexities.

We also compare our method with several robust end-to-end ST systems: Chimera (Han et al., 2021) which learns a shared semantic space for speech and text, XSTNet (Ye et al., 2021) which uses a

⁶<https://github.com/google/sentencepiece>

⁷<https://github.com/mjpost/sacrebleu>

progressive multi-task training strategy, STEMM (Fang et al., 2022) which applies word-level mixup strategy and ConST (Ye et al., 2022) which applies contrastive learning strategy. These baselines all adopt the multi-task learning mode and have similar model architectures.

3.3 Main Results

Table 1 shows the main results. HuBERT-Transformer, the baseline model we implement, achieves relatively high BLEU scores. For the base setting without external MT data, CMOT outperforms HuBERT-Transformer by 1.1 BLEU in average of 8 directions. For the expanded setting, CMOT outperforms HuBERT-Transformer by 0.9 BLEU in average. Besides, CMOT also outperforms other strong baselines.

4 Analysis

4.1 Effect of OT

We use CMOT to learn the alignment between speech and text, and to improve the performance of the ST system. In this section, we compare CMOT with other models without OT or with other types of OT. All the models have the same architecture. MTL denotes the multi-task learning (ST and MT) model. Relaxed OT denotes the model trained with the same process of CMOT but without the window strategy in Section 2.3. IPOT denotes the model trained with the same process as CMOT but using IPOT (Xie et al., 2020) for OT calculation. The A-score is defined as the accuracy of our alignment $A = (a_1, \dots, a_n)$ compared to the reference alignment $A^* = (a_1^*, \dots, a_n^*)$:

$$\text{A-score} = \sum_i \mathbb{1}(a_i = a_i^*)/n, \quad (13)$$

Here A^* is obtained with Montreal Forced Aligner (MFA)⁸ (McAuliffe et al., 2017). As shown in Table 3, we can see the result of relaxed OT is slightly lower than that of IPOT, but when the window strategy is used, its performance exceeds that of IPOT with a faster speed. CMOT achieves the best BLEU score, along with the best A-score of alignment, proving that CMOT not only takes advantage of this alignment but optimizes it as well.

In addition, we observe the cost matrices of OT to verify that the cost matrix can be used for cross-modal alignment. We compare the cost matrices in

⁸<https://montreal-forced-aligner.readthedocs.io/en/latest/>

Model	BLEU \uparrow	A-score \uparrow
HuBERT-Transformer	27.5	0.48
MTL	27.6	0.50
Relaxed OT	28.8	0.54
IPOT	28.9	0.55
CMOT	29.0	0.57

Table 3: ST performances and alignment scores of different models on MuST-C En-De tst-COMMON set. These experiments are conducted under the expanded setting.

MTL and CMOT, and analyze a typical example, as shown in Figure 3. The redder grids indicate the distance between the speech token and the text token is smaller, while the bluer grids are the opposite. To better observe the word-level alignment, we use MFA to cut the sentences of speech and text into words, and the white solid lines denote the segmentation. For CMOT, we can see that the closer speech-text pairs basically lie on the diagonal, which is consistent with the alignment of the two sequences. But for MTL without the help of OT, the distance between speech and the corresponding word is not that close. Comparing the two green boxes in the figure, we can see CMOT has better alignment on the word “pointed”.

4.2 Effect of Cross-modal Mixup

To examine the effectiveness of the mixup strategy in bridging the modality gap, we compare the modality gap of CMOT and baseline models. We calculate the mean sentence-level distances and mean word-level distances of these models. We focus on the modality gap at the encoder input and output, and the results are shown in Figure 4. We can see that compared with the baseline model, CMOT has a smaller distance between modalities, which means our method effectively reduces the modality gap and improves the performance of ST.

4.3 Ablation Study

Training Objectives As a multi-task learning framework, our ST system’s performance can be affected by the training objectives. Table 4 shows the impact of training objectives on ST performance. Comparing Exp.I and II, We can empirically conclude that multi-task learning strategy alone helps little, while Exp.III suggests regularizing the output distributions is effective. Exp.IV, V, and VI show that calculating KL divergence for both mixup and speech outputs, and mixup and text outputs can bring the greatest improvement. In Exp.VII, we

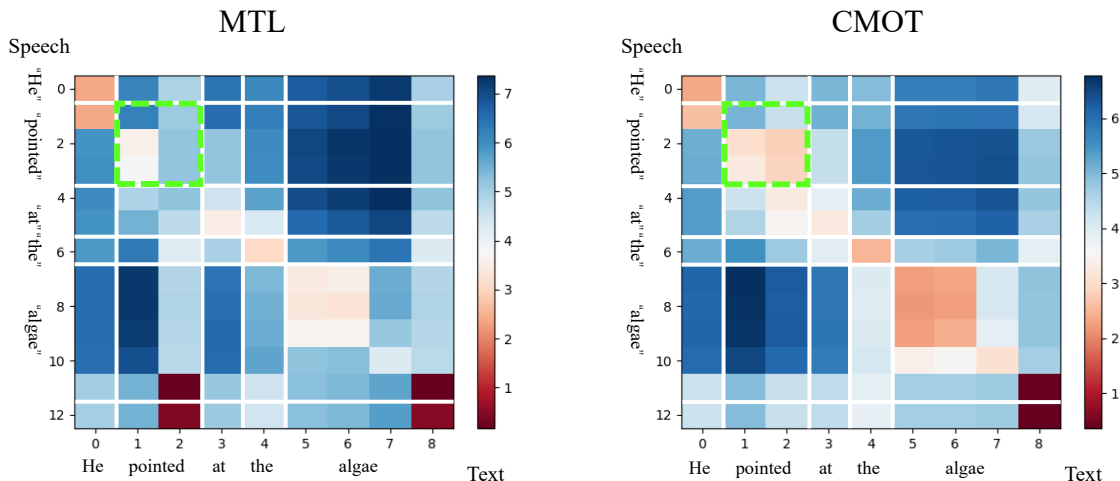


Figure 3: An example of how our CMOT can learn better alignment between speech and text. We draw the cost matrices in OT calculation for MTL and CMOT models.




Figure 4: Modality gaps of different models. The data are calculated on the MuST-C En-De dev set with models of expanded setting. “in” / “out” denotes the encoder input / output sequences, while “sentence” / “word” denotes sentence-level / word-level Euclidean distance.

add the translation loss of the mixup sequence:

$$\mathcal{L}_{\text{Mixup}} = -\mathbb{E}_{(s,x,y) \in \mathcal{D}} \log P(y|\mathbf{H}'), \quad (14)$$

where \mathbf{H}' is the mixup sequence calculated as Equation (8). Comparing Exp.VI and VII, we can see $\mathcal{L}_{\text{Mixup}}$ does not help. We also inspect the effect of OT loss which takes the value of \mathcal{D}^* in Equation (4) in Exp.VIII, and see a slight performance decline. Here OT is only used to help find the alignment, so OT loss is not indispensable during fine-tuning, and we conjecture adding OT loss may make it harder for the model to converge to good parameters, as the training objective is too complex.

Positions of OT and Mixup We investigate the impact of positions of OT and the mixup on our

#Exp.	Training Objective	BLEU
I	\mathcal{L}_{ST}	27.5
II	$\mathcal{L}_{\text{ST}} + \mathcal{L}_{\text{MT}}$	27.6
III	$\mathcal{L}_{\text{ST}} + \mathcal{L}_{\text{MT}} + \lambda \mathcal{L}_{\text{KL}}(S, T)$	28.5
IV	$\mathcal{L}_{\text{ST}} + \mathcal{L}_{\text{MT}} + \lambda \mathcal{L}_{\text{KL}}(M, S)$	28.5
V	$\mathcal{L}_{\text{ST}} + \mathcal{L}_{\text{MT}} + \lambda \mathcal{L}_{\text{KL}}(M, T)$	28.2
VI	$\mathcal{L}_{\text{ST}} + \mathcal{L}_{\text{MT}} + \lambda \mathcal{L}_{\text{KL}}(M, S) + \lambda \mathcal{L}_{\text{KL}}(M, T)$	29.0
VII	$\mathcal{L}_{\text{ST}} + \mathcal{L}_{\text{MT}} + \mathcal{L}_{\text{Mixup}} + \lambda \mathcal{L}_{\text{KL}}(M, S) + \lambda \mathcal{L}_{\text{KL}}(M, T)$	28.7
VIII	$\mathcal{L}_{\text{ST}} + \mathcal{L}_{\text{MT}} + \lambda \mathcal{L}_{\text{KL}}(M, S) + \lambda \mathcal{L}_{\text{KL}}(M, T) + \mu \mathcal{L}_{\text{OT}}$	28.6

Table 4: ST performances with different training objectives on MuST-C En-De tst-COMMON set. These experiments are conducted under the expanded setting. $\mathcal{L}_{\text{Mixup}}$ denotes the translation loss of the mixup sequence. We set the weights $\lambda = 2$ and $\mu = 0.1$.

OT	Mixup	
	encoder in	encoder out
encoder in	28.6	29.0
encoder out	28.7	28.7

Table 5: BLEU scores of different positions of OT and mixup on MuST-C En-De tst-COMMON set. These experiments are conducted under the expanded setting.

system. As shown in Table 5, the best scheme is to calculate OT for encoder inputs and perform the mixup for encoder outputs. We believe there is more original alignment information at the lower layers of the model, which is better for OT calculation, while the hidden states of higher layers are more suitable for the mixup, as they have passed several layers of joint transformer and are closer in the representation space.

Mixup Probability We explore the impact of mixup probability on ST performance and the results are shown in Figure 5. Here p^* denotes the proportion of text tokens in the mixup sequence.




Figure 5: BLEU scores on MuST-C En-De tst-COMMON set with different mixup probability p^* . These experiments are conducted under the expanded setting.

Speech Encoder	Baseline	CMOT
wav2vec 2.0 (Baevski et al., 2020)	26.9	28.5
HuBERT (Hsu et al., 2021)	27.5	29.0

Table 6: BLEU scores with different speech encoders on MuST-C En-De tst-COMMON set. Here "baseline" represents the model with the same structure as CMOT but without mixup during training. These experiments are conducted under the expanded setting.

Previous works find a slightly larger probability is better, while we find $p^* = 0.2$ improves ST performance the most. We think this is because CMOT needs to find the alignment between speech and text adaptively, so a larger proportion of text tokens in the mixup sequence might bring some noise.

Speech Encoder We investigate the impact of different speech encoders on our ST system. HuBERT (Hsu et al., 2021) and wav2vec 2.0 (Baevski et al., 2020) are both commonly used speech pretrained models. Table 6 indicates that using HuBERT as a speech encoder can achieve higher performances than using wav2vec 2.0.

In addition, in order to provide a fair comparison with other works using wav2vec 2.0, we conduct experiments on CMOT with wav2vec 2.0 as the speech encoder in more directions. As shown in Table 7, CMOT using wav2vec 2.0 still performs well, surpassing previous works in these three language pairs.

Weight of KL Loss Our finetuning process follows a multi-task learning approach, where the weight of each training objective can affect the overall performance. For the KL loss in Equation 12, we choose several different weight values λ ranging from 1.0 to 3.0, and find when $\lambda = 2.0$ or $\lambda = 2.5$ the performance is best, as shown in Figure 6. When the weight of KL loss is too small, the

Model	En-De	En-Fr	En-Es
Chimera	27.1	35.6	30.6
XSTNet	27.1	38.0	30.8
STEMM	28.7	37.4	31.0
ConST	28.3	38.3	32.5
W2V2-Transformer	26.9	36.6	30.0
CMOT-W2V2	28.5	39.1	32.6

Table 7: BLEU scores with different models using wav2vec 2.0 on MuST-C tst-COMMON set. "CMOT-W2V2" refers to our CMOT model with wav2vec 2.0 as the speech encoder. "W2V2-Transformer" refers to the model with the same structure as CMOT-W2V2 but without mixup during training, and we directly report the results given in Fang et al. (2022). Here the experiments of W2V2-Transformer and CMOT-W2V2 are both conducted under the expanded setting.




Figure 6: BLEU scores with different KL weight λ on MuST-C En-De tst-COMMON set. Here the x-axis is the weight of KL loss. These experiments are all conducted under the expanded setting. We set the mixup probability $p^* = 0.2$.

regularization effect of the KL objective becomes little, while a large weight of KL loss can lead to performance degradation of the main tasks ST and MT. Hence, we suggest the weight be at a moderate value, and we choose $\lambda = 2.0$ for our system.

Window Size for OT We examine the impact of window size W on our ST system. As shown in Figure 7, a small window size may restrict the model's ability to identify appropriate alignments, while a large window size may diminish its ability to capture local dependencies.

5 Related Work

End-to-end ST Traditional ST systems work in a cascaded mode (STENTIFORD and STEER, 1988; Waibel et al., 1991), with problems of error propagation and high latency. Bérard et al.




Figure 7: BLEU scores with different window sizes W on MuST-C En-De tst-COMMON set. These experiments are all conducted under the expanded setting.

(2016); Duong et al. (2016) proposed end-to-end ST without intermediate transcription, which has become the mainstream paradigm of ST in recent years (Vila et al., 2018; Salesky et al., 2019; Di Gangi et al., 2019b,c; Inaguma et al., 2020; Wang et al., 2020a; Zhao et al., 2021a; Dinh et al., 2022; Duquenne et al., 2022a). An end-to-end ST system is difficult to train due to its cross-modal cross-lingual characteristics and the scarcity of parallel data, hence some techniques like pretraining (Bansal et al., 2019; Stoian et al., 2020; Wang et al., 2020b; Alinejad and Sarkar, 2020; Zheng et al., 2021; Xu et al., 2021; Zhang et al., 2022b), self-training (Pino et al., 2020; Wang et al., 2021), curriculum learning (Kano et al., 2017; Wang et al., 2020c), multi-task learning (Le et al., 2020; Vydana et al., 2021; Tang et al., 2021), data augmentation (Lam et al., 2022; Fang and Feng, 2023a), and knowledge distillation (Liu et al., 2019; Gaido et al., 2020; Inaguma et al., 2021) are applied to improve its performance. To overcome the modality gap between speech and text, various methods are used. Zhao et al. (2021b) proposed a mutual learning scenario. Han et al. (2021) and Duquenne et al. (2022b) projected speech and text into a common representation space. Chen et al. (2022) proposed the method of modality matching. Ye et al. (2022) applied the contrastive learning strategy. Fang and Feng (2023b) proposed cross-modal regularization with scheduled sampling.

Mixup Our work draws on the strategy of mixup. Mixup is first proposed in Zhang et al. (2018) as a form of data augmentation to improve generalization and increase the robustness of neural networks by creating “convex combinations of pairs

of examples and their labels”. Verma et al. (2019) proposed the concept of manifold mixup, going a step further than surface-level mixup. In recent years, mixup is used to reduce the gap of representation spaces, such as cross-modal transfer (So et al., 2022; Ye and Guo, 2022) and cross-lingual transfer (Yang et al., 2021). Fang et al. (2022) proposed STEMM, applying mixup to end-to-end ST to overcome modality gap. Although both STEMM and our work use the mixup strategy, STEMM needs external alignment tools, while CMOT finds the alignment adaptively.

Optimal Transport Optimal transport is a classical mathematical problem. It is now commonly used to describe the transfer cost between two distributions. Villani (2009) systematically demonstrates the theories and proofs of optimal transport. Peyré et al. (2019) outlines the main theoretical insights that support the practical effectiveness of OT. There are some approximation algorithms for the OT problem, like the Sinkhorn algorithm (Cuturi, 2013) and IPOT (Xie et al., 2020) algorithm. Kusuner et al. (2015) proposed a relaxed form of OT in order to measure the similarity between documents, which reduced the computational cost. Chen et al. (2020) used OT in image-text pretraining to “explicitly encourage fine-grained alignment between words and image regions”. Gu and Feng (2022) applied the relaxed form of OT to machine translation and managed to “bridge the gap between the semantic-equivalent representations of different languages”.

6 Conclusion

We propose Cross-modal Mixup via Optimal Transport (CMOT) to adaptively find the alignment between speech and text sequences, and to mix up the sequences of different modalities at the token level. Experiments on the MuST-C benchmark demonstrate the effectiveness of our method, which helps overcome the modality gap and thus improves the performance of end-to-end ST systems. We will continue to explore the application of this adaptive cross-modal alignment method in the future.

Limitations

The CMOT method for speech translation is based on the idea of cross-modal knowledge transfer and the paradigm of multi-task learning, so it requires the transcripts of speech during training, which

may not be applicable to some languages without transcribed text. Besides, this work mainly focuses on the finetuning of speech translation, and applications on the pretraining phase or on other tasks are in need of exploration in the future.

Ethics Statement

Our proposed CMOT can help build a strong end-to-end ST system. It has prospects for application in many scenarios requiring speech translation, which helps people understand speech in a foreign language. Nevertheless, the results generated by end-to-end ST systems are not necessarily perfect, so people may not fully rely on the results for the time being.

Acknowledgements

This work was supported by National Key R&D Program of China(NO. 2018AAA0102502). We thank all the anonymous reviewers for their insightful and valuable comments.

References

- Ashkan Alinejad and Anoop Sarkar. 2020. Effectively pretraining a speech translation decoder with machine translation data. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8014–8020.
- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems*, 33:12449–12460.
- S Bansal, H Kamper, K Livescu, A Lopez, and S Goldwater. 2019. Pre-training on high-resource speech recognition improves low-resource speech-to-text translation. *NAACL 2019*.
- Alexandre Bérard, Olivier Pietquin, Laurent Besacier, and Christophe Servan. 2016. Listen and translate: A proof of concept for end-to-end speech-to-text translation. In *NIPS Workshop on end-to-end learning for speech and audio processing*.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurélie Névélol, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. 2016. *Findings of the 2016 conference on machine translation*. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 131–198, Berlin, Germany. Association for Computational Linguistics.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Uniter: Universal image-text representation learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX*, pages 104–120.
- Zhehuai Chen, Yu Zhang, Andrew Rosenberg, Bhuvana Ramabhadran, Pedro J. Moreno, Ankur Bapna, and Heiga Zen. 2022. MAESTRO: Matched Speech Text Representations through Modality Matching. In *Proc. Interspeech 2022*, pages 4093–4097.
- Marco Cuturi. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Mattia A Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019a. Must-c: a multilingual speech translation corpus. In *2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2012–2017. Association for Computational Linguistics.
- Mattia A Di Gangi, Matteo Negri, and Marco Turchi. 2019b. Adapting transformer to end-to-end spoken language translation. In *INTERSPEECH 2019*, pages 1133–1137. International Speech Communication Association (ISCA).
- Mattia Antonino Di Gangi, Matteo Negri, Roldano Cattoni, Roberto Dessi, and Marco Turchi. 2019c. Enhancing transformer for end-to-end speech-to-text translation. In *Proceedings of Machine Translation Summit XVII: Research Track*, pages 21–31.
- Tu Anh Dinh, Danni Liu, and Jan Niehues. 2022. Tackling data scarcity in speech translation using zero-shot multilingual machine translation techniques. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6222–6226. IEEE.
- Long Duong, Antonios Anastasopoulos, David Chiang, Steven Bird, and Trevor Cohn. 2016. An attentional model for speech translation without transcription. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 949–959.
- Paul-Ambroise Duquenne, Hongyu Gong, Benoît Sagot, and Holger Schwenk. 2022a. T-modules: Translation modules for zero-shot cross-modal machine translation. *arXiv preprint arXiv:2205.12216*.

- Paul-Ambroise Duquenne, Hongyu Gong, Benoît Sagot, and Holger Schwenk. 2022b. **T-Modules: Translation Modules for Zero-Shot Cross-Modal Machine Translation**. In *EMNLP 2022 - 2022 Conference on Empirical Methods in Natural Language Processing*, Abu Dhabi, United Arab Emirates.
- Qingkai Fang and Yang Feng. 2023a. Back translation for speech-to-text translation without transcripts. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*.
- Qingkai Fang and Yang Feng. 2023b. Understanding and bridging the modality gap for speech translation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*.
- Qingkai Fang, Rong Ye, Lei Li, Yang Feng, and Mingxuan Wang. 2022. **STEMM: Self-learning with speech-text manifold mixup for speech translation**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7050–7062, Dublin, Ireland. Association for Computational Linguistics.
- Tzu-hsun Feng, Annie Dong, Ching-Feng Yeh, Shu-wen Yang, Tzu-Quan Lin, Jiatong Shi, Kai-Wei Chang, Zili Huang, Haibin Wu, Xuankai Chang, et al. 2022. Superb@ slt 2022: Challenge on generalization and efficiency of self-supervised speech representation learning. *arXiv preprint arXiv:2210.08634*.
- Marco Gaido, Mattia A Di Gangi, Matteo Negri, and Marco Turchi. 2020. End-to-end speech-translation with knowledge distillation: Fbk@ iwslt2020. In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 80–88.
- Shuhao Gu and Yang Feng. 2022. Improving zero-shot multilingual translation with universal representations and cross-mappings. *arXiv preprint arXiv:2210.15851*.
- Chi Han, Mingxuan Wang, Heng Ji, and Lei Li. 2021. Learning shared semantic space for speech-to-text translation. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2214–2225.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460.
- Hirofumi Inaguma, Tatsuya Kawahara, and Shinji Watanabe. 2021. Source and target bidirectional knowledge distillation for end-to-end speech translation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1872–1881.
- Hirofumi Inaguma, Shun Kiyono, Kevin Duh, Shigeki Karita, Nelson Yalta, Tomoki Hayashi, and Shinji Watanabe. 2020. Espnet-st: All-in-one speech translation toolkit. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 302–311.
- Takatomo Kano, Sakriani Sakti, and Satoshi Nakamura. 2017. **Structured-Based Curriculum Learning for End-to-End English-Japanese Speech Translation**. In *Proc. Interspeech 2017*, pages 2630–2634.
- Diederik P. Kingma and Jimmy Ba. 2015. **Adam: A method for stochastic optimization**. In *ICLR (Poster)*.
- Taku Kudo and John Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71.
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From word embeddings to document distances. In *International conference on machine learning*, pages 957–966. PMLR.
- Tsz Kin Lam, Shigehiko Schamoni, and Stefan Riezler. 2022. **Sample, translate, recombine: Leveraging audio alignments for data augmentation in end-to-end speech translation**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 245–254, Dublin, Ireland. Association for Computational Linguistics.
- Hang Le, Juan Pino, Changhan Wang, Jiatao Gu, Didier Schwab, and Laurent Besacier. 2020. Dual-decoder transformer for joint automatic speech recognition and multilingual speech translation. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3520–3533.
- Yuchen Liu, Hao Xiong, Jiajun Zhang, Zhongjun He, Hua Wu, Haifeng Wang, and Chengqing Zong. 2019. **End-to-End Speech Translation with Knowledge Distillation**. In *Proc. Interspeech 2019*, pages 1128–1132.
- Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. 2017. Montreal forced aligner: Trainable text-speech alignment using kaldi. *Proc. Interspeech 2017*, pages 498–502.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: an asr corpus based on public domain audio books. In *2015*

- IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5206–5210. IEEE.
- Gabriel Peyré, Marco Cuturi, et al. 2019. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607.
- Juan Pino, Qiantong Xu, Xutai Ma, Mohammad Javad Dousti, and Yun Tang. 2020. [Self-Training for End-to-End Speech Translation](#). In *Proc. Interspeech 2020*, pages 1476–1480.
- Matt Post. 2018. A call for clarity in reporting bleu scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191.
- Elizabeth Salesky, Matthias Sperber, and Alexander H Waibel. 2019. Fluent translations from disfluent speech in end-to-end speech translation. In *NAACL-HLT (1)*.
- Adriaan M. J. Schakel and Benjamin J. Wilson. 2015. [Measuring word significance using distributed representations of words](#). *CoRR*, abs/1508.02297.
- Junhyuk So, Changdae Oh, Yongtaek Lim, Hoyoon Byun, Minchul Shin, and Kyungwoo Song. 2022. Geodesic multi-modal mixup for robust fine-tuning. *arXiv preprint arXiv:2203.03897*.
- FWM STENTIFORD and MG STEER. 1988. Machine translation of speech. *British Telecom technology journal*, 6(2):116–123.
- Mihaela C Stoian, Sameer Bansal, and Sharon Goldwater. 2020. Analyzing asr pretraining for low-resource speech-to-text translation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7909–7913. IEEE.
- Yun Tang, Juan Pino, Changhan Wang, Xutai Ma, and Dmitriy Genzel. 2021. [A general multi-task learning framework to leverage text data for speech to text tasks](#). In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6209–6213.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. 2019. Manifold mixup: Better representations by interpolating hidden states. In *International Conference on Machine Learning*, pages 6438–6447. PMLR.
- Laura Cross Vila, Carlos Escolano, José AR Fonollosa, and Marta R Costa-Jussa. 2018. End-to-end speech translation with the transformer. In *IberSPEECH*, pages 60–63.
- Cédric Villani. 2009. *Optimal transport: old and new*, volume 338. Springer.
- Hari Krishna Vydana, Martin Karafiát, Katerina Zmolkova, Lukáš Burget, and Honza Černocký. 2021. [Jointly trained transformers models for spoken language translation](#). In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7513–7517.
- A Waibel, AN Jain, AE McNair, H Saito, AG Hauptmann, and J Tebelskis. 1991. Janus: a speech-to-speech translation system using connectionist and symbolic processing strategies. In *[Proceedings] ICASSP 91: 1991 International Conference on Acoustics, Speech, and Signal Processing*, pages 793–796. IEEE.
- Changhan Wang, Yun Tang, Xutai Ma, Anne Wu, Dmytro Okhonko, and Juan Pino. 2020a. Fairseq s2t: Fast speech-to-text modeling with fairseq. In *Proceedings of the 1st Conference of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 33–39.
- Changhan Wang, Anne Wu, Juan Pino, Alexei Baevski, Michael Auli, and Alexis Conneau. 2021. [Large-Scale Self- and Semi-Supervised Learning for Speech Translation](#). In *Proc. Interspeech 2021*, pages 2242–2246.
- Chengyi Wang, Yu Wu, Shujie Liu, Zhenglu Yang, and Ming Zhou. 2020b. Bridging the gap between pre-training and fine-tuning for end-to-end speech translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9161–9168.
- Chengyi Wang, Yu Wu, Shujie Liu, Ming Zhou, and Zhenglu Yang. 2020c. Curriculum pre-training for end-to-end speech translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3728–3738.
- Yujia Xie, Xiangfeng Wang, Ruijia Wang, and Hongyuan Zha. 2020. A fast proximal point method for computing exact wasserstein distance. In *Uncertainty in artificial intelligence*, pages 433–453. PMLR.
- Chen Xu, Bojie Hu, Yanyang Li, Yuhao Zhang, Shen Huang, Qi Ju, Tong Xiao, and Jingbo Zhu. 2021. Stacked acoustic-and-textual encoding: Integrating the pre-trained models into speech translation encoders. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2619–2630.

- Huiyun Yang, Huadong Chen, Hao Zhou, and Lei Li. 2021. Enhancing cross-lingual transfer by manifold mixup. In *International Conference on Learning Representations*.
- Junjie Ye and Junjun Guo. 2022. Dual-level interactive multimodal-mixup encoder for multi-modal neural machine translation. *Applied Intelligence*, pages 1–10.
- Rong Ye, Mingxuan Wang, and Lei Li. 2021. [End-to-End Speech Translation via Cross-Modal Progressive Training](#). In *Proc. Interspeech 2021*, pages 2267–2271.
- Rong Ye, Mingxuan Wang, and Lei Li. 2022. [Cross-modal contrastive learning for speech translation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5099–5113, Seattle, United States. Association for Computational Linguistics.
- Sho Yokoi, Ryo Takahashi, Reina Akama, Jun Suzuki, and Kentaro Inui. 2020. Word rotator’s distance. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2944–2960.
- Biao Zhang, Philip Williams, Ivan Titov, and Rico Senrich. 2020. Improving massively multilingual neural machine translation and zero-shot translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1628–1639.
- Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. 2018. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*.
- Yuhao Zhang, Chen Xu, Bojie Hu, Chunliang Zhang, Tong Xiao, and Jingbo Zhu. 2022a. Improving end-to-end speech translation by leveraging auxiliary speech and text data. *arXiv preprint arXiv:2212.01778*.
- Ziqiang Zhang, Long Zhou, Junyi Ao, Shujie Liu, Lirong Dai, Jinyu Li, and Furu Wei. 2022b. [Speechcut: Bridging speech and text with hidden-unit for encoder-decoder based speech-text pre-training](#).
- Chengqi Zhao, Mingxuan Wang, Qianqian Dong, Rong Ye, and Lei Li. 2021a. Neurst: Neural speech translation toolkit. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 55–62.
- Jiawei Zhao, Wei Luo, Boxing Chen, and Andrew Gilman. 2021b. Mutual-learning improves end-to-end speech translation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3989–3994.
- Renjie Zheng, Junkun Chen, Mingbo Ma, and Liang Huang. 2021. Fused acoustic and text encoding for multimodal bilingual pretraining and speech translation. In *International Conference on Machine Learning*, pages 12736–12746. PMLR.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section 7 Limitations
- A2. Did you discuss any potential risks of your work?
Section 8 Ethic Statement
- A3. Do the abstract and introduction summarize the paper's main claims?
Section 0 Abstract, Section 1 Introduction
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 2 Method, Section 3 Experiments, Section 4 Analysis

- B1. Did you cite the creators of artifacts you used?
Section 2 Method, Section 3 Experiments, Section 4 Analysis
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
The datasets and models we use are for research purposes only, not for any commercial purposes. Our use does not violate their licenses.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Our use is consistent with their intended use, and there is no special need to emphasize this as the datasets and models we use have been used in the same way in recent years.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
The speech translation dataset we use is collected from public TED talks. Since they are already open on the Internet, we cannot anonymize them. The machine translation datasets we use are already anonymized, so there is no need to discuss their problems.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 3 Experiments
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Appendix A Statistics of Datasets

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

C Did you run computational experiments?

Section 3 Experiments

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

Section 3 Experiments

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Section 3 Experiments, Section 4 Analysis

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section 3 Experiments

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Section 3 Experiments

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.