

Listen, Understand and Translate: Triple Supervision Decouples End-to-end Speech-to-text Translation

Qianqian Dong,^{1,2*} Rong Ye,³ Mingxuan Wang,³ Hao Zhou,³ Shuang Xu,¹ Bo Xu,^{1,2} Lei Li³

¹ Institute of Automation, Chinese Academy of Sciences (CASIA), Beijing, China

² School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

³ ByteDance AI Lab, China

{dongqianqian2016, shuang.xu, xubo}@ia.ac.cn, {yerong, wangmingxuan.89, zhouhao.nlp, lileilab}@bytedance.com

Abstract

An end-to-end speech-to-text translation (ST) takes audio in a source language and outputs the text in a target language. Existing methods are limited by the amount of parallel corpus. Can we build a system to fully utilize signals in a parallel ST corpus? We are inspired by human understanding system which is composed of auditory perception and cognitive processing. In this paper, we propose Listen-Understand-Translate, (LUT), a unified framework with triple supervision signals to decouple the end-to-end speech-to-text translation task. LUT is able to guide the acoustic encoder to extract as much information from the auditory input. In addition, LUT utilizes a pre-trained BERT model to enforce the upper encoder to produce as much semantic information as possible, without extra data. We perform experiments on a diverse set of speech translation benchmarks, including Librispeech English-French, IWSLT English-German and TED English-Chinese. Our results demonstrate LUT achieves the state-of-the-art performance, outperforming previous methods. The code is available at <https://github.com/dqccasia/st>.

Introduction

Processing audio in one language and translating it into another language has been requested in many applications. Traditional speech translation (ST) systems are cascaded by connecting separately trained automatic speech recognition (ASR), spoken language normalization processing and machine translation (MT) subsystems (Sperber et al. 2017; Wang et al. 2018; Dong et al. 2019; Sperber et al. 2019b; Zhang et al. 2019; Beck, Cohn, and Haffari 2019; Cheng et al. 2019). However, such cascaded ST systems have drawbacks including higher latency, larger memory footprint, and potential error propagation in the subsystems. In contrast, an end-to-end ST system has a single unified model, which is beneficial in deployment. While very promising, existing end-to-end ST models still cannot outperform cascaded systems in terms of translation accuracy.

Cascaded ST systems usually have intermediate stages which extract acoustic features and source-text semantic features, before translating to the target text, like humans with

perception systems and cognitive systems to process different information. Ideally, a neural encoder-decoder network should also benefit from imitating these intermediate steps. The challenges are: a) there is no sufficient supervision to guide the internals of an encoder-decoder to process the audio input and obtain acoustic and semantic information properly. b) the training corpus for ST with pairs of source audio and target text is much smaller than those typically used for ASR and MT. Previous works attempt to relieve these challenges using pre-training and fine-tuning approaches. They usually initialize the ST model with the encoder trained on ASR data to mimic the speech recognizing process and then fine-tune on a speech translation dataset to make the cross-lingual translation. However, pre-training and fine-tuning are still not sufficient enough to train an effective ST system, for the following reasons: a) the encoder for speech recognition is mainly used to extract acoustic information, while the ST model requires to encode both acoustic and semantic information. b) previous studies (Battenberg et al. 2017) have proved that the learned alignments between input and output units in the ASR models are local and monotonic, which is not conducive to modeling long-distance dependencies for translation models. c) the gap of length between the input audio signals (typically ~ 1000 frames) and target sentences (typically ~ 20 tokens) renders the association from the encoder to decoder difficult to learn.

Based on the above analysis, we explore decoupled model structure, LUT, with an acoustic encoder (to Listen), a semantic encoder (to Understand), and a translation decoder (to Translate) to imitate the intermediate steps for effective end-to-end speech translation. In addition to the widely-used translation loss with cross-entropy, we exploit two additional auxiliary supervising signals. We utilize the monotonic alignment for transcriptions to ensure the acoustic encoder capture necessary acoustic information from the input audio spectrum sequence. In this way, the local relations among the nearby audio frames are preserved. We utilize the pre-trained embedding, to guide the semantic encoder to capture a proper semantic representation. Notice that neither of the two auxiliary supervision is required during the inference, and therefore our method is efficient.

The contributions of the paper are as follows: 1) We design LUT, a unified framework augmented with additional components and supervisions to decouple and guide

*The work was done while QD was interning at ByteDance AI Lab.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

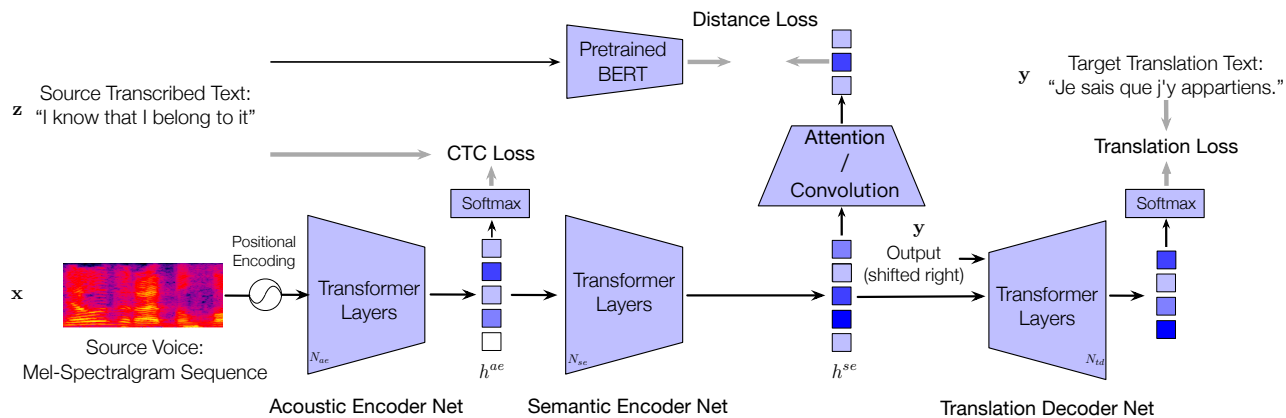


Figure 1: The architecture of LUT. It contains three modules, an acoustic encoder, a semantic encoder, and a translation decoder.

the speech translation. 2) Our proposed method can extract semantic knowledge from the pre-trained language model and utilize external ASR corpus to enhance acoustic modeling more effectively benefiting from the flexibly designed structure. 3) We conduct experiments and do analysis on three mainstream speech translation datasets, LibriSpeech (English-French), IWSLT2018 (English-German) and TED (English-Chinese), to verify the effectiveness of our model.

Methodology

In this section, we illustrate how we design the speech-to-text translation model, LUT, whose architecture allows a flexible configuration of the backbone network structure in each module. One can freely choose convolutional layers, recurrent neural networks, or *Transformer* network as the main building structure. Figure 1 illustrates the overall architecture of the LUT, using *Transformer* (Vaswani et al. 2017) as the backbone network. Our proposed LUT consists of three modules: *a)* an *acoustic encoder* network that encodes the audio input sequence into hidden features corresponding to the source text; *b)* a *semantic encoder* network that extracts hidden semantic representation for translation, which behaves like a normal machine translation encoder; *c)* a *translation decoder* network that outputs sentence tokens in the target language. Notice an input sequence typically has a length of more than 1000, while a target sentence has tens of tokens.

Problem Formulation The training corpus for speech translation contains speech-transcription-translation triples, denoted as $\mathcal{S} = \{(\mathbf{x}, \mathbf{z}, \mathbf{y})\}$. Specially, $\mathbf{x} = (x_1, \dots, x_{T_x})$ is a sequence of acoustic features. $\mathbf{z} = (z_1, \dots, z_{T_z})$ and $\mathbf{y} = (y_1, \dots, y_{T_y})$ represents the corresponding text sequence in source language and target language, respectively. Meanwhile, $\mathcal{A} = \{(\mathbf{x}', \mathbf{z}')\}$ represents the external ASR corpus. Usually, the amount of ST corpus is much smaller than that of ASR, i.e. $|\mathcal{S}| \ll |\mathcal{A}|$.

Acoustic Encoder

The acoustic encoder of LUT takes the input of low-level audio features and outputs a series of vectors corresponding to the transcribed text in the source language. The original audio signal is transformed into mel-frequency cepstrum (Mermelstein 1976), which is the standard preprocessing in speech recognition. The sequence of frames (\mathbf{x}) are processed by a feed-forward linear layer, and N_{ae} layers of *Transformer* sub-network, which includes a multi-head attention layer, a feed-forward layer, normalization layers, and residual connections. The output of acoustic encoder is denoted as \mathbf{h}^{ae} . They are further projected linearly with a softmax layer to obtain auxiliary output probability p for each token in the vocabulary. Note here the vocabulary is augmented with one extra blank symbol “_” (“<blk>” in the vocabulary). We employ the alignment-free Connectionist Temporal Classification (CTC) (Graves et al. 2006) algorithm to align the acoustic encoder output and the expected supervision sequence \mathbf{z} .

CTC Loss Given the ground truth transcribed token sequence \mathbf{z} , there can be multiple raw predicted label sequences from the acoustic encoder. Let g denote the mapping from the raw label sequence to the ground truth, which is based on a deterministic rule by removing the blank symbols and consecutive duplicate tokens. For example, $g(aa_abb_)=g(a_abb_)=aab$. $\mathcal{B}^{-1}(\mathbf{z})$ is the set of all raw label sequences corresponding to the ground truth transcription. Then the conditional probability of a ground truth token sequence \mathbf{z} can be modeled by marginalizing over all raw label sequences:

$$P(\mathbf{z}|\mathbf{x}) = \sum_{\pi \in \mathcal{B}^{-1}(\mathbf{z})} P(\pi|\mathbf{x}) \quad (1)$$

Where each raw label probability $p(\pi_t|\mathbf{x})$ for a sequence π is calculated from the acoustic encoder using the following equation:

$$P(\pi|\mathbf{x}) = \prod_{t=1}^{T_x} p(\pi_t|\mathbf{x}) \quad (2)$$

Finally, the acoustic encoder loss is defined as

$$\mathcal{L}_{ae}(\theta; \mathbf{x}, \mathbf{z}) = -\log P(\mathbf{z}|\mathbf{x}) \quad (3)$$

Semantic Encoder

The second module of LUT is motivated by the commonly used encoder for a neural machine translation (NMT) model. LUT’s semantic encoder aims to extract semantic and contextual information for translation. However, unlike the encoder in the NMT model taking the input of source sentence tokens, LUT’s semantic encoder takes the hidden representation \mathbf{h}_{ae} computed from the acoustic encoder as the input. Since we do not have explicit supervision of the semantic representation, we utilize a pre-trained BERT (Devlin et al. 2019) model to calculate sentence embeddings for the source sentence \mathbf{z} and then further employ these embeddings to supervise the training of this encoder module. This approach of self-supervision is advantageous because it enables training using a very large independent monolingual corpus in the source language.

The semantic encoder contains N_{se} *Transformer* layers at the core and then connects to two branches. The output of this module is denoted as \mathbf{h}^{se} . One branch is to compute an overall semantic vector of the input, marked as “Seq-level Distance”. It is realized using a 2D convolutional layer to reduce dimension, a normalization layer, and an average pooling layer to shrink the vectors into one. The output of this branch is denoted as v_0^{se} , which is a single vector. This is to be compared with the class-label representation h_c^{BERT} calculated by a BERT model. Another branch is aimed to match the semantic representation of the transcribed source sentence, marked as “Word-level Distance”. This branch is connected to an auxiliary layer to calculate the length-synchronized semantic representation, note as v_1^{se} , a sequence of vectors of the size T_z , equivalent to that of the transcribed source sentence. To this end, we first use a separately pre-trained BERT model to calculate the sentence embedding vectors, excluding the class-label vector h_c^{BERT} . These vectors are organized into T_z time steps, denoted as \mathbf{h}^{BERT} . Suppose the N_{se} -layer transformer outputs a sequence of vectors at length T_x , denoted as $\mathbf{v} = \mathbf{h}^{se}$. Note each of these vectors are split into J heads, i.e. $\mathbf{h}^{\text{BERT}} = (\mathbf{h}_1^{\text{BERT}}, \dots, \mathbf{h}_J^{\text{BERT}})$ and $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_J)$. These BERT vectors are used as queries to compute the attention weights for the branch input hidden vectors.

$$\text{head}_i = \text{Attn}(\mathbf{h}_i^{\text{BERT}} W_i^Q, \mathbf{v}_i W_i^K, \mathbf{v}_i W_i^V) \quad (4)$$

Where the W_i^Q, W_i^K, W_i^V are parameters for the attention of i -th head. The attention is calculated by scaled dot-product layer, as follows:

$$\text{Attn}(Q, K, V) = \text{Softmax}\left(\frac{Q \cdot K^T}{\sqrt{d_k}}\right)V \quad (5)$$

where d_k is the dimension of the key K . With this layer, the output can be reduced to the same length as source text by concatenating the heads.

$$\mathbf{v}_1^{se} = \text{Concat}(\text{head}_1, \dots, \text{head}_J) \quad (6)$$

Algorithm 1 Semi-supervised Training Strategy

- 1: **while** not converged **do**
 - 2: **Step 1:** sample a training batch from $\{(\mathbf{x}, \mathbf{z})\}$ (or \mathcal{A} in expanded setting) to optimize \mathcal{L}_{ae} and \mathcal{L}_{se} .
 - 3: **Step 2:** sample a training batch from \mathcal{S} to optimize the total loss in Equation 9.
 - 4: **end while**
-

Finally, the distance loss is defined as the mean-squared error loss (MSE) between the calculated hidden representations and the BERT embeddings.

$$\mathcal{L}_{se}(\theta; \mathbf{x}, \mathbf{z}) = \begin{cases} \text{MSE}(v_0^{se} - h_c^{\text{BERT}}), & \text{Seq-level} \\ \text{MSE}(\mathbf{v}_1^{se} - \mathbf{h}^{\text{BERT}}), & \text{Word-level} \end{cases} \quad (7)$$

The key insight of our formulation is that the semantic encoder needs to behave like a text encoder of the NMT model, with only source language text data in the training. The distance loss (Equation 7) ensures that the semantic encoder could produce similar semantic embeddings close to the BERT representation trained on a separate large text corpus. During the inference, the output of this module is \mathbf{h}^{se} , therefore no additional source transcription text is needed and the BERT calculation is saved.

Translation Decoder

As with the normal machine translation model, our proposed LUT uses N_{td} layers of *Transformer* network as the decoder. Additional attention from the decoder to the semantic encoder output \mathbf{h}^{se} is added. The cross entropy loss is defined in Equation 8.

$$\mathcal{L}_{td}(\theta; \mathbf{x}, \mathbf{y}) = -\sum_{i=1}^{T_y} \log p_{\theta}(y_i | y_{<i}, \mathbf{h}^{se}) \quad (8)$$

Where the decoder probability p_{θ} is calculated from the final softmax layer based on the output of the decoder.

Training The overall objective function for end-to-end training is the weighted sum for three supervision modules:

$$\begin{aligned} \mathcal{L}(\theta; \mathbf{x}, \mathbf{z}, \mathbf{y}) \\ = \alpha \mathcal{L}_{ae}(\theta; \mathbf{x}, \mathbf{z}) + \beta \mathcal{L}_{se}(\theta; \mathbf{x}, \mathbf{z}) + \gamma \mathcal{L}_{td}(\theta; \mathbf{x}, \mathbf{y}) \end{aligned} \quad (9)$$

Where θ is the model parameter, and α, β and γ are hyper-parameters to balance among the acoustic encoder loss \mathcal{L}_{ae} , the semantic encoder loss \mathcal{L}_{se} , and the translation decoder loss \mathcal{L}_{td} .

We have also used the semi-supervised training strategy (Algorithm 1) to take advantages of the parallel `<speech, transcription>` supervisions and external ASR resources. In this way, the end-to-end ST model can avoid the complicated process of pre-training and then fine-tuning, while being able to better converge and maintain acoustic modeling.

Experiments

Data

We conduct experiments on three popular publicly available datasets, including Augmented LibriSpeech English-French dataset (Kocabiyikoglu, Besacier, and Kraif 2018), IWSLT201 English-German dataset (Jan et al. 2018) and TED English-Chinese dataset (Liu et al. 2019).

Augmented LibriSpeech Dataset Augmented LibriSpeech is built by automatically aligning e-books in French with English utterances of LibriSpeech. The dataset includes four types of information: English speech signal, English transcription, French text translations from the alignment of e-books with augmented references via Google Translate. Following the previous work (Liu et al. 2019), we also conduct experiments on the 100 hours clean train set for training, with 2 hours development set and 4 hours test set, corresponding to 47271, 1071, and 2048 utterances respectively.

IWSLT2018 English-German Dataset IWSLT2018 English-German is the KIT end-to-end speech translation corpus, which is built automatically by aligning English audios with SRT transcripts for English and German from lectures online. The raw data, including long wave files, English transcriptions, and the corresponding German translations, is segmented into chunks with the attached timestamps. It should be noted that some transcriptions are not aligned with their corresponding audio well. Noisy data is harmful to models’ performance, which can be avoided by data filtering, re-alignment, and re-segmentation (Liu et al. 2018). In this paper, the original data is used directly as training data to verify our method, with a size of 272 hours and 171121 segmentations. We use *dev2010* as validation set and *tst2013* as test set, corresponding to 653 and 793 utterances respectively¹.

TED English-Chinese Dataset English-Chinese TED is crawled from TED website² and released by (Liu et al. 2019) as a benchmark for speech translation from English audio to Chinese text. Following the previous work (Liu et al. 2019), we use *dev2010* as development set and *tst2015* as test set. The raw long audio is segmented based on timestamps for complete semantic information. Finally, we get 524 hours train set, 1.5 hours validation set and 2.5 hours test set, corresponding to 308,660, 835, 1223 utterances respectively.

LIUM2 Dataset We use LIUM2 as the external ASR parallel corpus ($\in \mathcal{A}$) used in the expanded experimental setting for broad reproducibility. LIUM2 (Rousseau, Deléglise, and Esteve 2014) is composed of segments of public talks extracted from the TED lecture website with 207 hours of speech data. Speed perturbation is performed on the raw signals with speed factors 0.9 and 1.1.

¹We use the data segmentation in the *iwslt - corpus/parallel* folder.

²<https://www.ted.com>

Data Preprocessing Following the previous work (Liu et al. 2019), we use acoustic features that are 80 dimensional log-Mel filterbanks. The features are extracted with a step size of 10ms and a window size of 25ms and extended with mean subtraction and variance normalization. The features are stacked with 5 frames to the right and downsampled to a 30ms frame rate. For target language text data, we lowercase all the texts, tokenize and apply normalize punctuations with the Moses scripts³. For source language text data, we lowercase all the texts, tokenize and remove the punctuation to make the data more consistent with the output of ASR. We apply BPE⁴ on the combination of source and target text to obtain shared subword units. The number of merge operations in BPE for ASR and MT systems is set to 8k and 30k, respectively. For strategies using BERT features, we apply the same pre-processing tool as BERT does to text data for ST models and regenerate the vocabulary. For English-French and English-German corpora, we report case-insensitive BLEU scores by `multi-bleu.pl`⁵ script for the evaluation of ST and MT tasks. And for English-Chinese corpus, we report character-level BLEU scores. We use word error rates (WER) to evaluate ASR tasks.

Baselines

We conduct our experiments in different settings.

Base Setting with only Speech-translation Data Our main purpose is to compare our method with conventional end-to-end speech translation models. In the setting, the training data is restricted to only the triple data.

Expanded Setting with External Data In the context of expanded setting, Bahar et al. (2019) apply the SpecAugment (Park et al. 2019) on Librispeech English-French ST task, which uses a total of 236 hours of speech for ASR pre-training. Inaguma et al. (2019) combine three ST datasets of 472 hours training data to train a multilingual ST model for both Librispeech English-French ST task and IWSLT2013 English-German ST task. And Wang et al. (2020a) introduce an additional 207 hours ASR corpus and 40M parallel data from WMT18 to enhance the ST performance. We mainly explored additional ASR data in this work.

MT System The input of the MT system is the manual transcribed text. The result of MT system can be regarded as the upper bound of ST models.

Details of the Model and Experiments

For ST tasks, we use a similar hyper-parameter setting with the *Transformer* base model (Vaswani et al. 2017) for the stack of *Transformer* layers, in which we set the hidden size $d_{model} = 768$ (512 for English-Chinese TED set) to match

³<https://github.com/moses-smt/mosesdecoder>

⁴<https://github.com/rsennrich/subword-nmt>

⁵<https://github.com/moses-smt/mosesdecoder/scripts/generic/multi-bleu.perl>

Method	Enc Pre-train (speech data)	Dec Pre-train (text data)	greedy	beam
MT system				
Transformer MT (Liu et al. 2019)	-	-	21.35	22.91
Base ST setting				
LSTM ST (Bérard et al. 2018)	✗	✗	12.30	12.90
+pre-train+multitask (Bérard et al. 2018)	✓	✓	12.60	13.40
LSTM ST+pre-train (Inaguma et al. 2020)	✓	✓	-	16.68
Transformer+pre-train (Liu et al. 2019)	✓	✓	13.89	14.30
+knowledge distillation (Liu et al. 2019)	✓	✓	14.96	17.02
TCEN-LSTM (Wang et al. 2020a)	✓	✓	-	17.05
Transformer+ASR pre-train (Wang et al. 2020b)	✓	✗	-	15.97
Transformer+curriculum pre-train (Wang et al. 2020b)	✓	✗	-	17.66
LUT	✗	✗	16.70	17.75
Expanded ST setting				
LSTM+pre-train+SpecAugment (Bahar et al. 2019)	✓(236h)	✓	-	17.00
Multilingual ST+PT (Inaguma et al. 2019)	✓(472h)	✗	-	17.60
Transformer+ASR pre-train (Wang et al. 2020b)	✓(960h)	✗	-	16.90
Transformer+curriculum pre-train (Wang et al. 2020b)	✓(960h)	✗	-	18.01
LUT	✓(207h)	✗	17.55	18.34

Table 1: Performance on Augmented Librispeech English-French test set. LUT achieves the best performance on both the base ST and the expanded settings.

the output of BERT. We experiment with the officially released BERT ($H = 768$) on English-French and English-German and BERT ($H = 512$) on English-Chinese. Learning from speech-transformer (Dong, Xu, and Xu 2018), one third of the layer is used for the decoder ($N_{td} = 4$). N_{ae} and N_{se} are both set to 4 for our best performance. For ASR and MT tasks, the standard *Transformer* base model is adopted. All samples are batched together with 20000-frame features by approximate feature sequence length during training. We train our models on 1 NVIDIA V100 GPU with a maximum number of training steps 400k, using Adam optimizer (Kingma and Ba 2015) with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e^{-8}$, and warmup-decay learning rate schedule with peak lr = $4e^{-4}$, 25k warmup steps, 0.5 decay rate and 50k decay steps. We use a greedy search and beam search (as default) with a max beam size of 8 for our experimental settings. The maximum decoding length for ASR and ST (MT) is set to 200 and 250, respectively. The hyper-parameters in Equation 9, α , β and γ are set to 0.5, 0.05, 0.45 (details in Table 8). For experiments in the base setting, the ST model is trained from scratch. For experiments in the expanded setting, the ST model is trained as the following two steps: *a*) pre-training the acoustic encoder with CTC loss with $(\mathbf{x}', \mathbf{z}') \in \mathcal{A}$, as Equation 3. *b*) fine-tuning the overall ST model with $(\mathbf{x}, \mathbf{z}, \mathbf{y}) \in \mathcal{S}$ and $(\mathbf{x}', \mathbf{z}') \in \mathcal{A}$, as Algorithm 1. SpecAugment strategy (Park et al. 2019) is adopted to avoid overfitting with frequency masking ($F = 30$, $mF = 2$) and time masking ($T = 40$, $mT = 2$). The final model is averaged on the last 10 checkpoints.

Results

Main Results

Librispeech English-French For En-Fr experiments, we compare the performance with existing end-to-end meth-

ods in Table 1. Clearly, LUT outperforms the previous best results in base setting and expanded setting respectively. Specifically, in the base setting, the model we propose outperforms ESPnet, which is equipped with both a well-trained encoder and decoder. We also achieve better results than a knowledge distillation augmented baseline in which an MT model is introduced to teach the ST model (Liu et al. 2019). Different from previous work, our work focuses on reducing the modeling burdens of the encoder by suggesting that auxiliary supervision signals make it easier to learn both the acoustic and semantic information. This proposal promises great potential for the application of the double supervised encoder. Compared to Wang et al. (2020a,b) which also includes tandem encoders, LUT is simple and flexible, without the need to introduce additional computational cost for data preprocessing (such as training the noisy MT model or the alignment model). Simple yet effective, LUT achieves the best performance in this benchmark dataset in terms of BLEU.

IWSLT2018 English-German For En-De experiments, we compare the performance with existing end-to-end methods in Table 2. Unlike that of Librispeech English-French, this dataset is noisy, and the transcriptions do not align well with the corresponding audios. As a result, there is a wide gap between the performance of the end to end ST and the upper bound of the ST. Overall, our method outperforms ESPnet on tst2013 by 2.81 BLEU in the base setting and has an advantage of 0.44 BLEU compared with Wang et al. (2020b) in the expanded setting. To be notice, LUT does not need any pretraining tricks, and achieves the state-of-the-art performance in the base setting. This trend is consistent with that on the Librispeech dataset.

Method	Enc Pre-train (speech data)	Dec Pre-train (text data)	tst2013
MT system			
RNN MT (Inaguma et al. 2020)	-	-	24.90
Base ST setting			
ESPnet (Inaguma et al. 2020)	✗	✗	12.50
+enc pre-train	✓	✗	13.12
+enc dec pre-train	✓	✓	13.54
Transformer+ASR pre-train (Wang et al. 2020b)	✓	✗	15.35
+curriculum pre-train (Wang et al. 2020b)	✓	✗	16.27
LUT	✗	✗	16.35
Expanded ST setting			
Multilingual ST (Inaguma et al. 2019)	✓(472h)	✗	14.60
CL-fast* (Kano, Sakti, and Nakamura 2017)	✓(479h)	✗	14.33
TCEN-LSTM (Wang et al. 2020a)	✓(479h)	✓(40M)	17.67
Transformer+curriculum pre-train (Wang et al. 2020b)	✓(479h)	✓(4M)	18.15
LUT	✓(207h)	✗	18.59

Table 2: Performance on IWSLT2018 English-German test set. *: re-implemented by Wang et al. (2020b). LUT achieves the best performance.

Method	Enc Pre-train (speech data)	Dec Pre-train (text data)	BLEU
MT system			
Transformer MT (Liu et al. 2019)	-	-	27.08
Base setting			
Transformer+pre-train (Liu et al. 2019)	✓	✓	16.80
+knowledge distillation (Liu et al. 2019)	✓	✓	19.55
Multi-task+pre-train* (Inaguma et al. 2019)(re-implemented)	✓	✗	20.45
LUT	✗	✗	20.84

Table 3: Performance on TED English-Chinese test set. LUT achieves the best performance.

TED English-Chinese For En-Zh experiments, we compared the performance with existing end-to-end methods in Table 3. Under the base setting, LUT exceeded the Transformer-based ST model augmented by knowledge distillation with 1.29 BLEU, proving the validity of our method.

Comparison with Cascaded Baselines Table 4 shows the comparison with cascaded ST systems on Augmented Librispeech En-Fr test set, IWSLT2018 En-De tst2013 set and TED En-Zh test set. For a fair comparison, we do the experiments on the base settings of English-French/German/Chinese translation. Results show that LUT either receives the equivalent performance or outperforms with cascaded methods on three datasets, thus displaying great potential for the end-to-end approach. This indicates our flexible structure can make good use of additional ASR corpus and learn valuable linguistic knowledge.

Analysis

Effects of Auxiliary Supervision We first study the effects of two auxiliary supervision for LUT. The results in Table 5 show that all the auxiliary supervision indicate positive results that can be superimposed. Models that use su-

	Method	BLEU
$En \rightarrow Fr$	cascaded	17.58
	LUT	17.75
$En \rightarrow De$	cascaded	15.38
	LUT	16.35
$En \rightarrow Zh$	cascaded	21.36
	LUT	20.84

Table 4: LUT versus cascaded systems on Augmented Librispeech En-Fr test set, IWSLT2018 En-De tst2013 set and TED En-Zh test set. “cascaded” systems consist of separate ASR and MT models trained independently.

perision only from the acoustic encoder can be regarded as a method of multi-task learning, which has a significant performance improvement compared to the model of direct pre-training and fine-tuning (seen in Table 1). This reflects the catastrophic forgetting problem that occurs in the sequential transfer learning based on the pre-training method.

Balance of Acoustic and Semantic Modeling Experimental results, shown in Table 6, prove that the performance

	Dev Bleu	Test Bleu
LUT	18.51	17.75
w/o Semantic Encoder Loss	17.72	16.81
w/o Acoustic Encoder Loss *	16.91	15.48
w/o Acoustic Encoder Loss	12.05	11.24

Table 5: Ablation study on En-Fr validation and test set. “**” means using ASR pre-training as initialization.

is better when the two modules are balanced. In order to determine which module has a more significant impact on performance, we conducted experiments on the layer number allocation of the two modules, in which the total number of layers of the acoustic encoder and semantic encoder is fixed, and the number of layers for one module is adjusted from 2 to 6. As the number of layers decreases, the two modules will both result in worse performance degradation, thus explaining that using enough layers to extract acoustic features and encode semantic representation is equally essential to the speech translation model.

N_{ae}	N_{se}	Dev BLEU	Test BLEU
2	6	14.81	13.09
3	5	17.01	15.50
4	4	17.93	16.70
5	3	17.07	16.21
6	2	16.47	15.49

Table 6: Performance on En-Fr corpus: LUT with varying layers in its acoustic encoder (N_{ae}) and semantic encoder (N_{se}). Greedy decoding is employed.

Sequence-level Distance Versus Word-level Distance

For this part, we conduct experiments with different branches described in previous sections. We conduct an experimental comparison of the performance differences caused by the pre-training features extracted by different layers of BERT for semantic encoder’s supervision. We finally adopt the pre-trained features from the last layer of BERT as our default setting. The results, as shown in Table 7, prove that the word-level distance benefit more from the BERT pre-trained features because of its finer and grainer regulation.

	Dev BLEU	Test BLEU
Seq-level Distance	17.64	16.61
Word-level Distance	17.93	16.70

Table 7: Performance on En-Fr corpus: LUT with different losses for semantic encoder. “Seq-level” and “Word-level” losses are described in Eq. (7). Greedy decoding is employed.

No.	α	β	γ	dev-BLEU	tst-BLEU
I	0.50	0.05	0.45	18.61	17.55
II	0.40	0.20	0.40	18.56	17.00
III	0.30	0.40	0.30	18.70	17.34
IV	0.20	0.05	0.75	18.75	17.41
V	0.20	0.60	0.20	18.94	17.17
VI	0.80	0.05	0.15	18.38	17.31

Table 8: Ablations on the loss weight on Augmented Librispeech En-Fr dataset, with greedy decoding under the expanded setting.

Effect of Hyper-parameters in the Loss We vary the loss weights of the three terms in loss function, i.e. α , β and γ in Equation 9. The results are listed in Table 8. For different settings, we want to emphasize the influence of :

- Translation loss \mathcal{L}_{td} (No. IV);
- CTC loss \mathcal{L}_{ae} (No. VI);
- Distance loss \mathcal{L}_{se} (No. V);
- Both CTC loss \mathcal{L}_{ae} and translation loss \mathcal{L}_{td} (No. I & II),
- All of them (No. III).

The performances do not vary much (± 0.5 BLEU) after convergence. We also find that distance loss \mathcal{L}_{se} is the easiest term to converge while CTC loss \mathcal{L}_{ae} is the hardest during the training.

Shallower or Deeper Encoders In Table 6, we have concluded that both acoustic and semantic encoders are equally important to the entire system, while in this section, we explore the performance of the deeper or shallower balanced encoders (Table 9). The model with shallower encoders performs worse due to the limited expressive capacity. While the encoders go deeper, the 5-layer-encoder model translates slightly better. However the 6-layer-encoder model performs worse than the 4-layer one. This is probably due to the overfitting problem, also mentioned in (Wang et al. 2020b). In addition, we guess that a deeper model with a smaller hidden size can achieve better performance, which is reserved for future exploration.

N_{ae}	N_{se}	N_{td}	test BLEU
6	6	6	< 5
6	6	4	17.15
5	5	4	17.58
4	4	4	17.55
3	3	4	16.39
2	2	4	14.13

Table 9: Ablations on the number of encoder and decoder layers on Augmented Librispeech dataset, with greedy decoding under the expanded setting.

Acoustic or Semantic In Table 10, we design auxiliary probing tasks to further analyze the learned representation (Lugosch et al. 2019). **SpeakerVer** is designed to identify the speaker, therefore it benefits more from acoustic information. **IntentIde** is focused on intention recognition, so it needs more linguistic knowledge. We use the Fluent Speech Commands dataset (Lugosch et al. 2019) for experiments which contains 30,043 utterances, 97 speakers, and 31 intents. The utterances are randomly divided into train, valid, and test splits in such a way that each split contains all possible wordings for each intent and all speakers. For the train split, we extract the hidden output of each layer of our well-trained LUT encoder, average and freeze it, followed by a fully connected layer. We then fine-tune for 20,000 steps on the two probing tasks respectively. We report the accuracy of the test split. It can be seen that during the modeling process of ST, acoustic information is modeled at low-level layers and semantic information is captured at high-level layers.

	SpeakerVer	IntentIde
AE Output \mathbf{h}^{ae}	97.6	91.0
SE output \mathbf{h}^{se}	46.3	93.1

Table 10: Classification accuracy on speaker verification and intent identification, using LUT’s acoustic encoder output (\mathbf{h}^{ae}) and semantic encoder output (\mathbf{h}^{se}).

Model Sizes for ST Systems We make a detailed comparison between the performance of end-to-end systems and cascaded systems with different model parameter sizes. Compared with the original pipeline system, the pipeline (small) system is a cascade of speech recognition models and machine translation models with halved layers. Results in Table 11 prove that end-to-end models have advantages in balancing performance and parameter size of the system.

	En-Fr			Parameter
	ASR↓	MT↑	ST↑	
Pipeline (small)	21.3	19.5	15.9	$\approx 129M$
Pipeline	16.6	20.98	16.38	$\approx 209M$
LUT	-	-	16.70	$\approx 144M$

Table 11: Performance with greedy search for ASR, MT and ST tasks on En-Fr test set with different model sizes. Pipeline: consists of independently trained ASR and MT systems (see details in the experimental settings).

Attention Visualization We analyze the learned representation through visualizations of the acoustic and semantic modeling’s attention between layers. Figure 2 shows an example of the distribution of attention weights. The attention of acoustic modeling is local and monotonous from the first layer to the fourth layer, matching the behavior of ASR. The attention of semantic encoder gradually tends to be smoothed out across the global context, which is beneficial to modeling semantic information. The observation is in line with our hypothesis.




Figure 2: The visualization of attention for different module layers. (a), (b) visualize the attention of the last layer of acoustic encoder and the first layer of semantic encoder respectively. Both the horizontal and vertical coordinates represent the same sequence of speech frames.

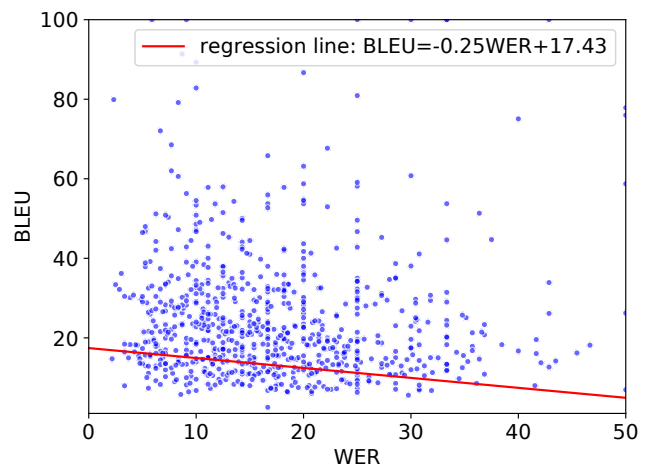


Figure 3: Relationship between WER and BLEU on En-Fr test set.

Correlation Analysis The quality of the hidden state obtained in our encoder depends largely on the accuracy of the acoustic modeling. Using the CTC loss function introduced in the acoustic encoder, we can also predict recognition results while predicting translation results. We can diagnose whether the wrong prediction for translation is caused by the wrong acoustic modeling in this way. We use samples from the test set to analyze the relationship between translation quality and acoustic modeling, which are evaluated by BLEU and WER respectively. We draw scatter plots of WER and BLEU on the test set, as can be seen in Figure 3. It can be seen that samples with a higher WER can usually obtain a translation result with a lower BLEU. Statistically, the Pearson correlation coefficient between BLEU and WER is $-0.205 < 0$ (with p-value = $2e^{-16} \ll 0.05$), which indicates the significant negative relation between them. At the same time, a minority of samples with a higher WER can obtain translation results with a higher BLEU, which indicates that our ST model has a certain degree of robustness to recognition errors.

Case Study Table 12 shows our case study analysis, proving that the end-to-end speech translation system can alleviate the problem of error propagation caused by upstream speech recognition errors. LUT can obtain the intermediate results of speech recognition by the way of CTC decoding, so it can perform a certain degree of interpretable diagnosis on the translation results. Benefiting from the ability of the end-to-end system to directly obtain the original audio information, our method is fault-tolerant in the case of incorrect recognition, missing recognition, repeated recognition, and so on during the acoustic modeling.

Speech #1	766-144485-0090.wav
Transcription	
<i>reference</i>	it was mister jack <u>maldon</u>
<i>hypothesis</i>	it was mister jack mal
Translation	
<i>reference</i>	c'était m. jack <u>maldon</u>
<i>hypothesis</i>	c'était m. jack <u>maldon</u>
Phenomenon	Incorrect Recognition
Speech #2	1257-122442-0101.wav
Transcription	
<i>reference</i>	cried the <u>old</u> soldier
<i>hypothesis</i>	cried the soldier
Translation	
<i>reference</i>	s'écria le <u>vieux</u> soldat,
<i>hypothesis</i>	s'écria le <u>vieux</u> soldat,
Phenomenon	Missing Recognition
Speech #3	1184-121026-0000.wav
Transcription	
<i>reference</i>	chapter <u>seventeen</u> the abbes chamber
<i>hypothesis</i>	chapter <u>seventeen</u> teen the abbey chamber
Translation	
<i>reference</i>	chapitre <u>xvii</u> la chambre de l'abbé.
<i>hypothesis</i>	chapitre <u>xvii</u> la chambre de l'abbé.
Phenomenon	Repeated Recognition

Table 12: Examples of transcription and translation on En-Fr test set generated by LUT. The underlined text means the desired output, and the bold text represents the incorrect predictions.

Related Work

End-to-end ST Previous works (Bérard et al. 2016; Duong et al. 2016) have proved the potential for end-to-end ST, which has attracted intensive attentions (Salesky et al. 2018; Di Gangi, Negri, and Turchi 2019; Bahar, Bieschke, and Ney 2019; Di Gangi et al. 2019; Inaguma et al. 2020). It's proved that pre-training (Weiss et al. 2017; Bérard et al. 2018; Bansal et al. 2019; Stoian, Bansal, and Goldwater 2020) and multi-task learning (Vydan et al. 2020) can significantly improve the performance. Two-pass decoding (Sung et al. 2019) and attention-passing (Anastasopoulos and Chiang 2018; Sperber et al. 2019a) techniques are proposed to handle deeper relationships and alleviate error propagation in end-to-end models. Data augmentation techniques (Jia et al. 2019; Pino et al. 2019; Bahar et al. 2019) are proposed to utilize ASR and MT corpora to generate fake data.

Pre-training has brought great gains to end-to-end models, such as knowledge distillation (Liu et al. 2019), modality agnostic meta-learning (Indurthi et al. 2019), curriculum learning (Kano, Sakti, and Nakamura 2017; Wang et al. 2020b) and so on (Wang et al. 2020a). Liu et al. (2020); Liu, Spanakis, and Niehues (2020) optimize the decoding strategy to achieve low-latency end-to-end ST. Salesky, Sperber, and Black (2019) explored additional features to enhance end-to-end models.

Knowledge Distillation for MT Representations learned by pre-training have been widely applied in the field of machine translation. There're different modeling granularities for representation learning, which can be utilized by fine-tuning, freezing, or distillation. Many knowledge distillation methods have been extended to transfer the rich knowledge using teacher-student architecture (Gou et al. 2020). Kim and Rush (2016) extended word-level knowledge distillation into sequence-level knowledge distillation for directing the sequence distribution of the student model. Sequence-level knowledge distillation was further explained from the perspective of data augmentation and regularization in Gordon and Duh (2019). Zhou, Gu, and Neubig (2019) studied how knowledge distillation affects the non-autoregressive MT models by empirical analysis. Lample and Conneau (2019); Edunov, Baevski, and Auli (2019) feeded the last layer of ELMo or BERT to the encoder of MT model for learning better representation. Yang et al. (2020) firstly leveraged asymptotic distillation to transfer the pre-training information to MT model.

Conclusion

In this paper, we propose LUT, a novel and unified training framework to decouple the end-to-end speech translation task into three phases: "Listen", "Understand" and "Translate", supervised by the acoustic, semantic, linguistic supervision, respectively. We empirically demonstrate the effectiveness of our approach as compared to previous methods on three benchmark datasets, and empirical analysis suggests that LUT is capable of capturing all acoustic, semantic and cross-lingual information properly.

Acknowledgements

We would like to thank the anonymous reviewers for their valuable comments. We would also like to thank Chengqi Zhao, Cheng Yi and Chi Han for their useful suggestions and helps. This work was supported by the Key Programs of the Chinese Academy of Sciences under Grant No.ZDBS-SSW-JSC006-2.

References

- Anastasopoulos, A.; and Chiang, D. 2018. Tied Multitask Learning for Neural Speech Translation. In *Proc. of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics*.
- Bahar, P.; Bieschke, T.; and Ney, H. 2019. A comparative study on end-to-end speech to text translation. In *Proc. of the 2019 IEEE Automatic Speech Recognition and Understanding Workshop*.

- Bahar, P.; Zeyer, A.; Schlüter, R.; and Ney, H. 2019. On using specaugment for end-to-end speech translation. In *Proc. of the 2019 International Workshop on Spoken Language Translation*.
- Bansal, S.; Kamper, H.; Livescu, K.; Lopez, A.; and Goldwater, S. 2019. Pre-training on high-resource speech recognition improves low-resource speech-to-text translation. In *Proc. of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*.
- Battenberg, E.; Chen, J.; Child, R.; Coates, A.; Li, Y. G. Y.; Liu, H.; Satheesh, S.; Sriram, A.; and Zhu, Z. 2017. Exploring neural transducers for end-to-end speech recognition. In *Proc. of the 2017 IEEE Automatic Speech Recognition and Understanding Workshop*.
- Beck, D.; Cohn, T.; and Haffari, G. 2019. Neural Speech Translation using Lattice Transformations and Graph Networks. In *Proc. of the 2019 Workshop on Graph-Based Methods for Natural Language Processing*.
- Bérard, A.; Besacier, L.; Kocabiyyikoglu, A. C.; and Pietquin, O. 2018. End-to-end automatic speech translation of audiobooks. In *Proc. of the 2018 IEEE International Conference on Acoustics, Speech and Signal Processing*.
- Bérard, A.; Pietquin, O.; Servan, C.; and Besacier, L. 2016. Listen and translate: A proof of concept for end-to-end speech-to-text translation. *NIPS Workshop on End-to-End Learning for Speech and Audio Processing*.
- Cheng, Q.; Fang, M.; Han, Y.; Huang, J.; and Duan, Y. 2019. Breaking the Data Barrier: Towards Robust Speech Translation via Adversarial Stability Training. In *Proc. of the 2019 International Workshop on Spoken Language Translation*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proc. of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*.
- Di Gangi, M. A.; Negri, M.; Cattoni, R.; Roberto, D.; and Turchi, M. 2019. Enhancing transformer for end-to-end speech-to-text translation. In *Machine Translation Summit XVII*.
- Di Gangi, M. A.; Negri, M.; and Turchi, M. 2019. Adapting transformer to end-to-end spoken language translation. In *Proc. of the 2019 Conference of the International Speech Communication Association*.
- Dong, L.; Xu, S.; and Xu, B. 2018. Speech-transformer: a no-recurrence sequence-to-sequence model for speech recognition. In *Proc. of the 2018 IEEE International Conference on Acoustics, Speech and Signal Processing*.
- Dong, Q.; Wang, F.; Yang, Z.; Chen, W.; Xu, S.; and Xu, B. 2019. Adapting Translation Models for Transcript Disfluency Detection. In *Proc. of the 2019 AAAI Conference on Artificial Intelligence*, volume 33.
- Duong, L.; Anastasopoulos, A.; Chiang, D.; Bird, S.; and Cohn, T. 2016. An attentional model for speech translation without transcription. In *Proc. of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics*.
- Edunov, S.; Baevski, A.; and Auli, M. 2019. Pre-trained language model representations for language generation. In *Proc. of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*.
- Gordon, M. A.; and Duh, K. 2019. Explaining Sequence-Level Knowledge Distillation as Data-Augmentation for Neural Machine Translation. *arXiv preprint arXiv:1912.03334*.
- Gou, J.; Yu, B.; Maybank, S. J.; and Tao, D. 2020. Knowledge Distillation: A Survey. *International Journal of Computer Vision*.
- Graves, A.; Fernández, S.; Gomez, F.; and Schmidhuber, J. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proc. of the 2016 International Conference on Machine Learning*.
- Inaguma, H.; Duh, K.; Kawahara, T.; and Watanabe, S. 2019. Multilingual end-to-end speech translation. In *Proc. of the 2019 IEEE Automatic Speech Recognition and Understanding Workshop*.
- Inaguma, H.; Kiyono, S.; Duh, K.; Karita, S.; Yalta, N.; Hayashi, T.; and Watanabe, S. 2020. ESPnet-ST: All-in-One Speech Translation Toolkit. In *Proc. of the 2020 Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.
- Indurthi, S.; Han, H.; Lakumarapu, N. K.; Lee, B.; Chung, I.; Kim, S.; and Kim, C. 2019. Data Efficient Direct Speech-to-Text Translation with Modality Agnostic Meta-Learning. *ICASSP*.
- Jan, N.; Cattoni, R.; Sebastian, S.; Cettolo, M.; Turchi, M.; and Federico, M. 2018. The iwslt 2018 evaluation campaign. In *Proc. of the 2018 International Workshop on Spoken Language Translation*.
- Jia, Y.; Johnson, M.; Macherey, W.; Weiss, R. J.; Cao, Y.; Chiu, C.-C.; Ari, N.; Laurenzo, S.; and Wu, Y. 2019. Leveraging weakly supervised data to improve end-to-end speech-to-text translation. In *Proc. of the 2019 IEEE International Conference on Acoustics, Speech and Signal Processing*.
- Kano, T.; Sakti, S.; and Nakamura, S. 2017. Structured-Based Curriculum Learning for End-to-End English-Japanese Speech Translation. *Proc. of the 2017 Conference of the International Speech Communication Association*.
- Kim, Y.; and Rush, A. M. 2016. Sequence-Level Knowledge Distillation. In *Proc. of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A method for stochastic optimization. *Proc. of the 3rd International Conference on Learning Representations*.
- Kocabiyyikoglu, A. C.; Besacier, L.; and Kraif, O. 2018. Augmenting Librispeech with French Translations: A Multimodal Corpus for Direct Speech Translation Evaluation. In *Proc. of the 2018 International Conference on Language Resources and Evaluation*.

- Lample, G.; and Conneau, A. 2019. Cross-lingual Language Model Pretraining. *Advances in Neural Information Processing Systems (NeurIPS)* .
- Liu, D.; Liu, J.; Guo, W.; Xiong, S.; Ma, Z.; Song, R.; Wu, C.; and Liu, Q. 2018. The USTC-NEL Speech Translation system at IWSLT 2018. In *Proc. of the 2018 International Workshop on Spoken Language Translation* .
- Liu, D.; Spanakis, G.; and Niehues, J. 2020. Low-Latency Sequence-to-Sequence Speech Recognition and Translation by Partial Hypothesis Selection. *Proc. Interspeech 2020* 3620–3624.
- Liu, Y.; Xiong, H.; Zhang, J.; He, Z.; Wu, H.; Wang, H.; and Zong, C. 2019. End-to-End Speech Translation with Knowledge Distillation. In *Proc. of the 2019 Conference of the International Speech Communication Association* .
- Liu, Y.; Zhang, J.; Xiong, H.; Zhou, L.; He, Z.; Wu, H.; Wang, H.; and Zong, C. 2020. Synchronous speech recognition and speech-to-text translation with interactive decoding. In *Proc. of the 2020 AAAI Conference on Artificial Intelligence* .
- Lugosch, L.; Ravanelli, M.; Ignoto, P.; Tomar, V. S.; and Bengio, Y. 2019. Speech Model Pre-Training for End-to-End Spoken Language Understanding. *Proc. of the 2019 Conference of the International Speech Communication Association* .
- Mermelstein, P. 1976. Distance measures for speech recognition, psychological and instrumental. *Pattern Recognition and Artificial Intelligence* .
- Park, D. S.; Chan, W.; Zhang, Y.; Chiu, C.-C.; Zoph, B.; Cubuk, E. D.; and Le, Q. V. 2019. SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition. *Proc. of the 2019 Conference of the International Speech Communication Association* .
- Pino, J.; Puzon, L.; Gu, J.; Ma, X.; McCarthy, A. D.; and Gopinath, D. 2019. Harnessing indirect training data for end-to-end automatic speech translation: Tricks of the trade. In *Proc. of the 2019 International Workshop on Spoken Language Translation* .
- Rousseau, A.; Deléglise, P.; and Esteve, Y. 2014. Enhancing the TED-LIUM corpus with selected data for language modeling and more TED talks. In *LREC* .
- Salesky, E.; Burger, S.; Niehues, J.; and Waibel, A. 2018. Towards fluent translations from disfluent speech. In *Proc. of the 2018 IEEE Spoken Language Technology Workshop* .
- Salesky, E.; Sperber, M.; and Black, A. W. 2019. Exploring Phoneme-Level Speech Representations for End-to-End Speech Translation. In *Proc. of the 2019 Annual Meeting of the Association for Computational Linguistics* .
- Sperber, M.; Neubig, G.; Niehues, J.; and Waibel, A. 2017. Neural Lattice-to-Sequence Models for Uncertain Inputs. In *Proc. of the 2017 Conference on Empirical Methods in Natural Language Processing* .
- Sperber, M.; Neubig, G.; Niehues, J.; and Waibel, A. 2019a. Attention-Passing Models for Robust and Data-Efficient End-to-End Speech Translation. *Transactions of the Association for Computational Linguistics* .
- Sperber, M.; Neubig, G.; Pham, N.-Q.; and Waibel, A. 2019b. Self-Attentional Models for Lattice Inputs. In *Proc. of the 2019 Annual Meeting of the Association for Computational Linguistics* .
- Stoian, M. C.; Bansal, S.; and Goldwater, S. 2020. Analyzing ASR pretraining for low-resource speech-to-text translation. In *Proc. of the 2020 IEEE International Conference on Acoustics, Speech and Signal Processing* .
- Sung, T.-W.; Liu, J.-Y.; Lee, H.-y.; and Lee, L.-s. 2019. Towards end-to-end speech-to-text translation with two-pass decoding. In *Proc. of the 2019 IEEE International Conference on Acoustics, Speech and Signal Processing* .
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is All you Need. In *Proc. of the 2017 International Conference on Neural Information Processing Systems* .
- Vydana, H. K.; Karafi'at, M.; Zmolikova, K.; Burget, L.; and Cernocky, H. 2020. Jointly Trained Transformers models for Spoken Language Translation. *arXiv preprint arXiv:2004.12111* .
- Wang, C.; Wu, Y.; Liu, S.; Yang, Z.; and Zhou, M. 2020a. Bridging the gap between pre-training and fine-tuning for end-to-end speech translation. In *Proc. of the 2020 AAAI Conference on Artificial Intelligence* .
- Wang, C.; Wu, Y.; Liu, S.; Zhou, M.; and Yang, Z. 2020b. Curriculum Pre-training for End-to-End Speech Translation. In *Proc. of the 2020 Annual Meeting of the Association for Computational Linguistics* .
- Wang, F.; Chen, W.; Yang, Z.; Dong, Q.; Xu, S.; and Xu, B. 2018. Semi-supervised disfluency detection. In *Proc. of the 27th International Conference on Computational Linguistics* .
- Weiss, R. J.; Chorowski, J.; Jaitly, N.; Wu, Y.; and Chen, Z. 2017. Sequence-to-Sequence Models Can Directly Translate Foreign Speech. In *Proc. of the 2017 Conference of the International Speech Communication Association* .
- Yang, J.; Wang, M.; Zhou, H.; Zhao, C.; Zhang, W.; Yu, Y.; and Li, L. 2020. Towards making the most of bert in neural machine translation. In *Proc. of the 2020 AAAI Conference on Artificial Intelligence* .
- Zhang, P.; Ge, N.; Chen, B.; and Fan, K. 2019. Lattice Transformer for Speech Translation. In *Proc. of the 2019 Annual Meeting of the Association for Computational Linguistics* .
- Zhou, C.; Gu, J.; and Neubig, G. 2019. Understanding Knowledge Distillation in Non-autoregressive Machine Translation. In *Proc. of the 2019 International Conference on Learning Representations* .