

Stacked Acoustic-and-Textual Encoding: Integrating the Pre-trained Models into Speech Translation Encoders

Chen Xu¹, Bojie Hu², Yanyang Li³, Yuhao Zhang¹,
Shen Huang², Qi Ju^{2*}, Tong Xiao^{1,4*}, Jingbo Zhu^{1,4}

¹NLP Lab, School of Computer Science and Engineering,
Northeastern University, Shenyang, China

²Tencent Minority-Mandarin Translation, Beijing, China

³The Chinese University of Hong Kong, Hong Kong, China

⁴NiuTrans Research, Shenyang, China

{xuchen, blamedrlee, yoohaozhang}@outlook.com,

{bojiehu, springhuang, damonju}@tencent.com,

{xiaotong, zhujingbo}@mail.neu.edu.cn

Abstract

Encoder pre-training is promising in end-to-end Speech Translation (ST), given the fact that speech-to-translation data is scarce. But ST encoders are not simple instances of Automatic Speech Recognition (ASR) or Machine Translation (MT) encoders. For example, we find that ASR encoders lack the global context representation, which is necessary for translation, whereas MT encoders are not designed to deal with long but locally attentive acoustic sequences. In this work, we propose a *Stacked Acoustic-and-Textual Encoding* (SATE) method for speech translation. Our encoder begins with processing the acoustic sequence as usual, but later behaves more like an MT encoder for a global representation of the input sequence. In this way, it is straightforward to incorporate the pre-trained models into the system. Also, we develop an adaptor module to alleviate the representation inconsistency between the pre-trained ASR encoder and MT encoder, and develop a multi-teacher knowledge distillation method to preserve the pre-training knowledge. Experimental results on the LibriSpeech En-Fr and MuST-C En-De ST tasks show that our method achieves state-of-the-art BLEU scores of 18.3 and 25.2. To our knowledge, we are the first to develop an end-to-end ST system that achieves comparable or even better BLEU performance than the cascaded ST counterpart when large-scale ASR and MT data is available¹.

1 Introduction

End-to-end Speech Translation (E2E ST) has become popular recently for its ability to free designers from cascading different systems and shorten

*Corresponding author

¹The source code is available at <https://github.com/xuchen-neu/SATE>

Setting	Model	BLEU
Restricted	Cascaded	23.3
	E2E+Pre-training	23.1
Unrestricted	Cascaded	28.1
	E2E+Pre-training	25.6

Table 1: BLEU scores [%] of a cascaded ST model and an end-to-end ST model with pre-training on the MuST-C En-De corpus. Restricted = training is restricted to the ST data, and Unrestricted = additional training data is allowed for ASR and MT.

the pipeline of translation (Duong et al., 2016; Bernard et al., 2016; Weiss et al., 2017). Promising results on small-scale tasks are generally favorable. However, speech-to-translation paired data is scarce. Researchers typically use pre-trained Automatic Speech Recognition (ASR) and Machine Translation (MT) models to boost ST systems (Bernard et al., 2018). For example, one can initialize the ST encoder using a large-scale ASR model (Bansal et al., 2019). But we note that, despite significant development effort, our end-to-end ST system with pre-trained models was not able to outperform the cascaded ST counterpart when the ASR and MT data size was orders of magnitude larger than that of ST (see Table 1).

In this paper, we explore reasons why pre-training has been challenging in ST, and how pre-trained ASR and MT models might be used together to improve ST. We find that the ST encoder plays both roles of acoustic encoding and textual encoding. This makes it problematic to view an ST encoder as either an individual ASR encoder or an individual MT encoder. More specifically, there are two problems.

- Modeling deficiency: the MT encoder tries to capture long-distance dependency structures of language, but the ASR encoder focuses more on local dependencies in the input sequence. Since the ST encoder is initialized by the pre-trained ASR encoder (Berard et al., 2018), it fails to model large contexts in the utterance. But a large scope of representation learning is necessary for translation (Yang et al., 2018).
- Representation inconsistency: on the decoder side of ST, the MT decoder is in general used to initialize the model. The assumption here is that the upstream component is an MT-like encoder, whereas the ST encoder actually behaves more like an ASR encoder.

We address these problems by marrying the world of ASR encoding with the world of MT encoding. We propose a *Stacked Acoustic-and-Textual Encoding* (SATE) method to cascade the ASR encoder and the MT encoder. It first reads and processes the sequence of acoustic features as a usual ASR encoder. Then an adaptor module passes the acoustic encoding output to an MT encoder with two principles: informative and adaptive. In this way, pre-trained ASR and MT encoders can work for what we would originally design them, and the incorporation of pre-trained models into ST is more straightforward. In addition, we develop a multi-teacher knowledge distillation method to robustly train the ST encoder and preserve the pre-trained knowledge during fine-tuning (Yang et al., 2020).

We test our method in a Transformer-based end-to-end ST system. Experimental results on the LibriSpeech En-Fr and MuST-C En-De speech translation benchmarks show that it achieves the state-of-the-art performance of 18.3 and 25.2 BLEU points. Under a more challenging setup, where the large-scale ASR and MT data is available, SATE achieves comparable or even better performance than the cascaded ST counterpart. We believe that we are the first to present an end-to-end system that can beat the strong cascaded system in unrestricted speech translation tasks.

2 Related Work

Speech translation aims at learning models that can predict, given some speech in the source language, the translation into the target language. The earliest

of these models were cascaded: they treated ST as a pipeline of running an ASR system and an MT system sequentially (Ney, 1999; Mathias and Byrne, 2006; Schultz et al., 2004). This allows the use of off-the-shelf models, and was (and is) popular in practical ST systems. However, these systems were sensitive to the errors introduced by different component systems and the high latency of the long pipeline.

As another stream in the ST area, end-to-end methods have been promising recently (Berard et al., 2016; Weiss et al., 2017; Berard et al., 2018). The rise of end-to-end ST can be traced back to the success of deep neural models (Duong et al., 2016). But, unlike other well-defined tasks in deep learning, annotated speech-to-translation data is scarce, which prevents well-trained ST models. A simple solution to this issue is data augmentation (Pino et al., 2019, 2020). This method is model-free but generating large-scale synthetic data is time consuming. As an alternative, researchers used multi-task learning (MTL) to robustly train the ST model so that it could benefit from additional guide signals (Weiss et al., 2017; Anastasopoulos and Chiang, 2018; Berard et al., 2018; Sperber et al., 2019; Dong et al., 2021). Generally, MTL requires a careful design of the loss functions and more complicated architectures.

In a similar way, more recent work pre-trains different components of the ST system, and consolidates them into one. For example, one can initialize the encoder with an ASR model, and initialize the decoder with the target-language side of an MT model (Berard et al., 2018; Bansal et al., 2019; Stoian et al., 2020). More sophisticated methods include better training and fine-tuning (Wang et al., 2020a,b), the shrink mechanism (Liu et al., 2020), the adversarial regularizer (Alinejad and Sarkar, 2020), and etc. Although pre-trained models have quickly become dominant in many NLP tasks, they are still found to underperform the cascaded model in ST. This motivates us to explore the reasons why this happens and methods to solve the problems accordingly.

3 Why is ST Encoding Difficult?

Following previous work in end-to-end models (Berard et al., 2016; Weiss et al., 2017), we envision an encoding-decoding process in which an input sequence is encoded into a representation vector, and the vector is then decoded into an output sequence.




Figure 1: (a) Localness in each layer of the ST, ASR, and MT encoders, (b) the impact of CTC position on localness, and (c) the impact of CTC position on performance of ST and ASR models.

In such a scenario, all end-to-end ST, ASR and MT systems can be viewed as instances of the same architecture. Then, components of these systems can be pre-trained and re-used across them.

An underlying assumption here is that the ST encoder is doing something quite similar to what the MT (or ASR) encoder is doing. However, [Sperber et al. \(2018\)](#) find that the ASR model benefits from a small attention window, which is inconsistent with the MT model ([Yang et al., 2018](#)). To verify this, we compare the behavior of ST, ASR and MT encoders. We choose Transformer as the base architecture ([Vaswani et al., 2017](#)) and run experiments on the MuST-C En-De corpus. We report the results on the MuST-C En-De tst-COMMON test data. For stronger systems, we use Connectionist Temporal Classification (CTC) ([Graves et al., 2006](#)) as the auxiliary loss on the encoders when we train the ASR and ST systems ([Watanabe et al., 2017](#); [Karita et al., 2019](#); [Bahar et al., 2019](#)). The CTC loss forces the encoders to learn alignments between speech and transcription. It is necessary for the state-of-the-art performance ([Watanabe et al., 2018](#)).

Here we define the *localness* of a word as the sum of the attention weights to the surrounding words (or features) within a fixed small window². The window size is 10% of the sequence length. Figure 1(a) shows the localness of the attention weights for different layers of the encoders. We see that the ST and ASR encoders prefer local attention which indicates a kind of short-distance dependencies in processing acoustics feature sequences. Whereas the MT encoder generates a

²Here we treat the attention weight of Transformer as a distribution over all positions.

more global distribution of attention weights for word sequences, especially when we stack more layers. This result arises a new question: Is local attention sufficient for speech translation?

Then, we design another experiment to examine if the high localness in attention weights of the ASR and ST encoders is due to the bias imposed by CTC. In Figure 1(b), we use the CTC loss in the intermediate layer and show the average localness of the layers above or below CTC. The CTC loss demonstrates strong preference for locally attentive models. The upper-level layers act more like an MT encoder, that is, the layers with no CTC loss generates more global distributions. Taking this further, Figure 1(c) demonstrates a slightly higher BLEU score when we free more upper-level layers from the guide of CTC. Meanwhile, the word error rate (WER) increases because only lower parts of the model are learned in a standard manner of ASR.

Now we have some hints: the ST encoder is not a simple substitution of the ASR encoder or the MT encoder. Rather, they are complementary to each other, that is, we need the ASR encoder to deal with the acoustic input, and the MT encoder to generate the representation vector that can work better with the decoder.

4 The Method

In speech translation, we want the encoder to represent the input speech to some sort of decoder-friendly representations. We also want the encoder to be “natural” for pre-training. In the following, we describe, *Stacked Acoustics-and-Textual Encoding* (SATE), a new ST encoding method to meet these requirements, and improvements of it.




Figure 2: The overall architecture of stacked acoustic-and-textual encoding.

4.1 Stacked Acoustic-and-Textual Encoding

Unlike previous work, the SATE method does not rely on a single encoder to receive the signal from both the CTC loss and the feedback of the decoder. Instead, it is composed of two encoders: the first does exactly the same thing as the ASR encoder (call it *acoustic encoder*), and the other generates a higher-level globally-attentive representation on top of the acoustic encoder (call it *textual encoder*).

See Figure 2 for the architecture of SATE. The acoustic encoder is trained by CTC in addition to the supervision signal from the translation loss. Let (x, y^s, y^t) be an ST training sample, where x is the input feature sequence of the speech, y^s is the transcription of x , and y^t is the translation in the target language. We define the output of the acoustic encoder as:

$$h^s = E_s(x) \quad (1)$$

where $E_s(\cdot)$ is the encoding function. Then, we add a Softmax layer on h^s to predict the CTC label path $\pi = (\pi_1, \dots, \pi_T)$, where T is the length of the input sequence. The probability of path $P(\pi|h^s)$ is the product of the probability $P(\pi_t|h_t^s)$ at every time t based on conditionally independent assumption:

$$P(\pi|h^s) \approx \prod_t P(\pi_t|h_t^s) \quad (2)$$

CTC works by summing over the probability of all possible alignment paths $\Phi(y^s)$ between x and y^s , as follows:

$$P_{\text{CTC}}(y^s|h^s) = \sum_{\pi \in \Phi(y^s)} P(\pi|h^s) \quad (3)$$

Then, the CTC loss is defined as:

$$\mathcal{L}_{\text{CTC}} = -\log P_{\text{CTC}}(y^s|h^s; \theta_{\text{CTC}}) \quad (4)$$

where θ_{CTC} is the model parameters of the acoustic encoder and the CTC output layer.

The acoustic encoder is followed by an adaptor. It receives h^s and $P(\pi|h^s)$, and produces a new representation required by the textual encoder. Let $A(\cdot, \cdot)$ be the adaptor module. Its output is defined as:

$$\hat{h}^s = A(h^s, P(\pi|h^s)) \quad (5)$$

We leave the design of the adaptor to Section 4.2. Furthermore, we stack the textual encoder on the adaptor. The output h^t is defined as:

$$h^t = E_t(\hat{h}^s) \quad (6)$$

where $E_t(\cdot)$ is the textual encoder. h^t is fed into the decoder for computing the translation probability $P_{\text{Trans}}(y^t|h^t)$, as in standard MT systems. We define the translation loss as:

$$\mathcal{L}_{\text{Trans}} = -\log P_{\text{Trans}}(y^t|h^t; \theta_{\text{ST}}) \quad (7)$$

where θ_{ST} is all model parameters except for the CTC output layer.

Finally, we interpolate \mathcal{L}_{CTC} and $\mathcal{L}_{\text{Trans}}$ (with coefficient α) for the loss of the entire model:

$$\mathcal{L} = \alpha \cdot \mathcal{L}_{\text{CTC}} + (1 - \alpha) \cdot \mathcal{L}_{\text{Trans}} \quad (8)$$

Since the textual encoder works for the decoder only, it is trained as an MT encoder. In this way, the acoustic and textual encoders can do what we would originally expect them to do: the acoustic encoder deals with the acoustic input (i.e., ASR encoding), and the textual encoder generates a representation for translation (i.e., MT encoding). Also, SATE is friendly to pre-training. One can simply use an ASR encoder as the acoustic encoder, and use an MT encoder as the textual encoder. Note that SATE is in general a cascaded model, in response to the pioneering work in ST (Ney, 1999). It can be seen as cascading the ASR and MT systems in an end-to-end fashion.

4.2 The Adaptor

Now we turn to the design of the adaptor. Note that the pre-trained MT encoder assumes that the input is a word embedding sequence. Simply stacking the MT encoder and the ASR encoder obviously

does not work well. For this reason, the adaptor fits the output of the ASR encoder (i.e., the acoustic encoder) to what an MT encoder would like to see. We follow two principles in designing the adaptor: *adaptive* and *informative*.

We need an adaptive representation to make the input of the textual encoder similar to that of the MT encoder. To this end, we generate the soft contextual representation that shares the same latent space with the embedding layer of the MT encoder.

As shown in Eq. (2), the CTC output $P(\pi_t|h_t^s)$ indicates the alignment probability over the vocabulary at time t . Instead of replacing the representation by the embedding of the most-likely token (Liu et al., 2020), we employ a soft token which is the expectation of the embedding over the distribution from CTC. Let W^e be the embedding matrix of the textual encoder, we define the soft representation h_{soft}^s as:

$$h_{\text{soft}}^s = P(\pi|h^s) \cdot W^e \quad (9)$$

Also, an informative representation should contain information in the original input (Peters et al., 2018). The output acoustic representation of the ASR encoder generally involves paralinguistic information, such as emotion, accent, and emphasis. They are not expressed in the form of text explicitly but might be helpful for translation. For example, the generation of the declarative or exclamatory sentences depends on the emotions of the speakers.

We introduce a single-layer neural network to learn to map the acoustic representation to the latent space of the textual encoder, which preserves the acoustic information:

$$h_{\text{map}}^s = \text{ReLU}(W^{\text{map}} \cdot h^s + b^{\text{map}}) \quad (10)$$

where W^{map} and b^{map} are the trainable parameters.

The final output of the adaptor is defined to be:

$$A(h^s, P(\pi|h^s)) = \lambda \cdot h_{\text{map}}^s + (1 - \lambda) \cdot h_{\text{soft}}^s \quad (11)$$

where λ is the weight of h_{map}^s and set to 0.5 by default. Figure 3 shows the architecture of the adaptor.

Note that, in the adaptor, we do not change the sequence length for textual encoding because such a way is simple for implementation and shows satisfactory results in our experiments. Although there is a length inconsistency issue, the sequence representation of the speech should be similar with the

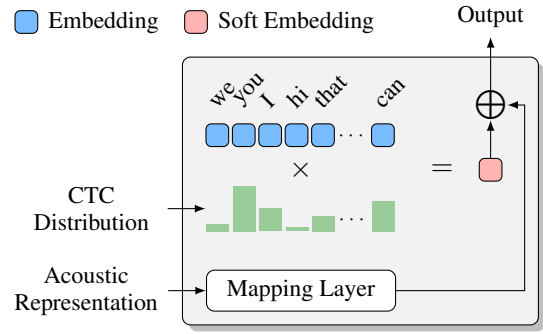


Figure 3: The architecture of the adaptor.

correspond transcription. Shrinking the sequence simply results in information incompleteness. We will investigate this issue in the future.

4.3 Multi-teacher Knowledge Distillation

Another improvement here is that we develop a multi-teacher knowledge distillation (MTKD) method to preserve the pre-trained knowledge during fine-tuning (Hinton et al., 2015).

The ST model mimics the teacher distribution by minimizing the cross-entropy loss between the teacher and student (Liu et al., 2019). For a training sample (x, y^s, y^t) , we define two loss functions:

$$\mathcal{L}_{\text{KD_CTC}} = - \sum_{m=1}^T \sum_{k=1}^{|V|} Q(\pi_m = v_k|x; \theta_{\text{ASR}}) \times \log P(\pi_m = v_k|x; \theta_{\text{CTC}}) \quad (12)$$

$$\mathcal{L}_{\text{KD_Trans}} = - \sum_{n=1}^{|y^t|} \sum_{k=1}^{|V|} Q(y_n^t = v_k|y^s; \theta_{\text{MT}}) \times \log P(y_n^t = v_k|x; \theta_{\text{ST}}) \quad (13)$$

where v_k is the word indexed by k and V is the vocabulary shared among the ST, ASR, and MT models. $Q(\cdot|\cdot)$ is the teacher distribution and $P(\cdot|\cdot)$ is the student distribution. θ_{ASR} , θ_{CTC} , θ_{MT} and θ_{ST} are the model parameters.

We can rewrite Eq. (8) to obtain a new loss:

$$\mathcal{L} = \alpha \cdot (\beta \cdot \mathcal{L}_{\text{CTC}} + (1 - \beta) \cdot \mathcal{L}_{\text{KD_CTC}}) + (1 - \alpha) \cdot (\gamma \cdot \mathcal{L}_{\text{Trans}} + (1 - \gamma) \cdot \mathcal{L}_{\text{KD_Trans}}) \quad (14)$$

where both β and γ are the hyper-parameters that balance the preference between the teacher distribution and the ground truth.

5 Experiments

5.1 Datasets and Preprocessing

We consider restricted and unrestricted settings on speech translation tasks. We run experiments on the LibriSpeech English-French (En-Fr) (Kocabişikoglu et al., 2018) and MuST-C English-German (En-De) (Gangi et al., 2019) corpora, which correspond to the low-resource and high-resource datasets respectively. Available ASR and MT data is only from the ST data under the restricted setting. For comparison in practical scenarios, the unrestricted setting allows the additional data for ASR and MT models.

LibriSpeech En-Fr Followed previous work, we use the clean speech translation training set of 100 hours, including 45K utterances and doubled translations of *Google Translate*. We select the model on the dev set (1,071 utterances) and report results on the test set (2,048 utterances).

MuST-C En-De MuST-C is a multilingual speech translation corpus extracted from the TED talks. We run the experiments on the English-German speech translation dataset of 400 hours speech with 230K utterances. We select the model on the dev set (1,408 utterances) and report results on the tst-COMMON set (2,641 utterances).

Unrestricted Setting We use the additional ASR and MT data for pre-training. The 960 hours LibriSpeech ASR corpus is used for the English ASR model. We extract 10M sentence pairs from the WMT14 English-French and 18M sentence pairs from the Opensubtitle2018³ English-German translation datasets.

Preprocessing Followed the preprocessing recipes of ESPnet (Inaguma et al., 2020), we remove the utterances of more than 3,000 frames and augment speech data by speed perturbation with factors of 0.9, 1.0, and 1.1. The 80-channel log-mel filterbank coefficients with 3-dimensional pitch features are extracted for speech data. We use the lower-cased transcriptions without punctuations. The text is tokenized using the scripts of Moses (Koehn et al., 2007). We learn Byte-Pair Encoding (Sennrich et al., 2016) subword segmentation with 10,000 merge operations based on a shared source and target vocabulary for all datasets.

³<http://opus.nlpl.eu/OpenSubtitles-v2018.php>

5.2 Model Settings

All experiments are implemented based on the ESPnet toolkit⁴. We use the Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.997$ and adopt the default learning schedule in ESPnet. We apply dropout with a rate of 0.1 and label smoothing $\epsilon_{ls} = 0.1$ for regularization.

For reducing the computational cost, the input speech features are processed by two convolutional layers, which have a stride of 2×2 and down-sample the sequence by a factor of 4 (Weiss et al., 2017). The encoder consists of 12 layers for both the ASR and vanilla ST models, and 6 layers for the MT model. The encoder of SATE includes an acoustic encoder of 12 layers and a textual encoder of 6 layers. The decoder consists of 6 layers for all models. The weight of CTC objective α for multitask learning is set to 0.3 for all ASR and ST models. The coefficients β and γ are set to 0.5 in Eq. (14) for the MTKD method.

Under the restricted setting, we employ the Transformer architecture, where each layer comprises 256 hidden units, 4 attention heads, and 2048 feed-forward size. For the unrestricted setting, we use the superior architecture Conformer (Gulati et al., 2020) on the ASR and ST tasks and widen the model by increasing the hidden size to 512 and attention heads to 8. The ASR⁵ and MT models pre-train with the additional data and fine-tune the model parameters with the task-specific data.

During inference, we average the model parameters on the best 5 checkpoints based on the performance of the development set. We use beam search with a beam size of 4 for all models. Different from previous work, we report the case-sensitive SacreBLEU⁶ (Post, 2018) for future standardization comparison across papers.

5.3 Results

Results on MuST-C En-De Table 2 summarizes the experimental results on the MuST-C En-De task. Under the restricted setting, the cascaded ST model translates the output of the ASR model, which degrades the performance compared with the MT model that translates from the reference transcription. The performance of the E2E ST baseline with pre-training is only slightly lower than the cascaded counterpart. SATE outperforms the baseline

⁴<https://github.com/espnet/espnet>

⁵We use the pre-trained ASR model offered by ESPnet.

⁶BLEU+case.mixed+numrefs.1+smooth.exp+tok.13a+version.1.4.14

Method	Restricted	Unrestricted
ESPnet MT*	27.63	-
ESPnet Cascaded*	23.65	-
MT	26.9	31.1
Cascaded ST	23.3	28.1
ESPnet E2E ST*	22.33	-
E2E ST	22.1	23.6
+Pre-training	23.1	25.6
SATE	23.3	23.6
+Pre-training	24.1	27.3
+MTKD	24.7	27.9
+SpecAug	25.2	28.1

Table 2: BLEU scores [%] on the test set of MuST-C En-De corpus. *: results reported in the ESPnet toolkit.

model significantly. This demonstrates the superiority of stacked acoustic and textual encoding for the speech translation task. Incorporating the pre-trained ASR and MT models into SATE releases the encoding burden of the model and achieves a remarkable improvement. The MTKD method provides a strong supervised signal and forces the model to preserve the pre-trained knowledge. Furthermore, we utilize the SpecAugment (Park et al., 2019) which is applied in the input speech features for better generalization and robustness⁷. It yields a remarkable improvement of 1.9 BLEU points over the cascaded baseline and achieves a new state-of-the-art performance.

Under the unrestricted setting, the large-scale ASR and MT data is available, whereas the ST data is scarce. This leads to the cascaded method outperforms the vanilla E2E method with a huge margin of 4.5 BLEU points. The pre-training only slightly closes the gap due to the modeling deficiency and representation inconsistency. SATE incorporates the pre-trained models fully, which achieves a significant improvement of 3.7 BLEU points. With the MTKD and SpecAugment methods, we achieve a comparable performance of 28.1 BLEU points. To our knowledge, we are the first to develop an end-to-end ST system that achieves comparable performance with the cascaded counterpart when large-scale ASR and MT data is available.

Results on LibriSpeech En-Fr Table 3 summarizes the experimental results on the LibriSpeech En-Fr task. Different from the MuST-C corpus,

⁷It is a fair comparison because the ASR model in the cascaded ST system also trains with the SpecAugment.

Method	Restricted	Unrestricted
ESPnet MT*	18.09	-
ESPnet Cascaded*	16.96	-
MT	17.5	21.3
Cascaded ST	16.3	20.6
ESPnet E2E ST*	16.22	-
E2E ST	16.7	17.7
+Pre-training	17.1	20.0
SATE	17.6	18.1
+Pre-training	17.4	20.8
+MTKD	17.7	20.8
+SpecAug	18.3	20.8

Table 3: BLEU scores [%] on the test set of LibriSpeech En-Fr corpus. *: results reported in the ESPnet toolkit.

it is of small magnitude with clean speech data. This results in that the performance of the vanilla E2E baseline is even better than the cascaded counterpart under the restricted setting. Furthermore, pre-training helps the model achieve an improvement of 0.8 BLEU points over the cascaded baseline. More interestingly, SATE without pre-training outperforms the above methods significantly, even achieves a slight improvement than the MT model. A possible reason is that the diverse acoustic representation is fed to the textual encoder, which improves the robustness of the model. This demonstrates the superiority of our method.

Combining our proposed methods yields a substantial improvement of 2.0 BLEU points over the cascaded baseline. It is a new state-of-the-art result of 18.3 BLEU points. Also, we outperform the cascaded counterpart by 0.2 BLEU points on the unrestricted task.

6 Analysis

6.1 Model Performance vs. Speedup

In Table 4, we summarize the performance and inference speedup based on the real time factor (RTF). The vanilla E2E ST model yields an inference speedup of $1.91\times$ than the cascaded counterpart and demonstrates the low latency of the end-to-end methods. We increase the encoder layers for comparison with SATE under the similar model parameters. However, there is a remarkable gap of 0.5 or 0.6 BLEU points, with or without pre-training.

Method	BLEU	RTF/Speedup
Cascaded ST	23.3	0.0286/1.00×
E2E ST	22.1	0.0150/1.91×
+Pre-training	23.1	
E2E ST (Enc 18)	22.8	0.0155/1.85×
+Pre-training	23.5	
SATE	23.3	0.0169/1.69×
+Pre-training	24.1	
+All	25.2	

Table 4: BLEU scores [%] and speedup on the test set (2641 utterances) of the MuST-C En-De corpus under the restricted setting. We evaluate the RTF on the NVIDIA V100 GPU with a batch size of 4 for all models.

Pre-trained Module	MuST-C	LibriSpeech
All	27.3	20.8
-ASR Enc	24.7	19.9
-MT	25.1	19.4
-MT Enc	25.7	20.7
-MT Dec	25.3	19.9

Table 5: Effects of the pre-trained modules on BLEU scores [%] under the unrestricted setting. We only remove one pre-trained module in each experiment.

Our method not only improves the performance of 1.9 BLEU points but also reaches up to 1.69× speedup than the cascaded baseline. This encourages the application of the end-to-end ST model in practical scenarios.

6.2 Effects of Pre-trained Modules

The effects of the pre-trained modules are shown in Table 5. The model performance drops significantly without the pre-trained ASR encoder, especially on the MuST-C corpus that contains noisy speech. The model parameters of pre-trained MT model are updated for adapting the output representation of the random initialized acoustic encoder. This results in the catastrophic forgetting problem (Goodfellow et al., 2015). The effect of the pre-trained MT model is more remarkable on the LibriSpeech corpus due to the modeling burden on the translation. The benefit of the pre-trained MT decoder is larger than the MT encoder. This is contrary to the previous conclusions that the MT encoder helps the performance significantly (Li et al., 2020). A possible reason is that the pre-trained

Design	MuST-C	LibriSpeech
None	25.7	21.7
Soft	25.7	21.9
Mapping	26.0	21.8
Fusion	26.4	21.9

Table 6: BLEU scores [%] of different adaptor setups on the development set under the unrestricted setting.




Figure 4: The localness of the vanilla E2E ST model and SATE model with pre-training.

ASR encoder provides a rich representation and acts as part of the MT encoder, this leads to lower performance degradation when the textual encoder trains from scratch.

Each pre-trained module has a great effect on the final performance. With the complete integration of the pre-trained modules, the model parameters are updated slightly, which preserves the pre-trained knowledge.

6.3 Effects of The Adaptor

We show the effects of the adaptor in Table 6. The straight connection which omits the representation inconsistency issue results in the lower benefit of pre-training. Although the soft representation aims at generating the adaptive representation, there is no obvious improvement on the MuST-C corpus. A possible reason is that the noisy speech inputs produce the misalignment probabilities, which disturbs the textual encoding. The mapping method achieves a slight improvement by transforming the acoustic representation to the textual representation. Fusing the soft and mapping representation enriches the information and avoids the representation inconsistency issue, which achieves the best performances.

6.4 Impact on Localness

We show the encoder localness of the vanilla E2E ST model and SATE model with pre-training in

Figure 4. As mentioned above, the vanilla ST model inherits the preference of ASR, which focuses on short-distance dependencies. SATE initializes with the pre-trained ASR and MT encoders, which stacks acoustic and textual encoding. The complementary behaviors of the pre-trained models benefit the translation, that is, the lower layers act like an ASR encoder while the upper layers capture global representation like an MT encoder.

7 Conclusion

In this paper, we investigate the difficulty of speech translation and shed light on the reasons why pre-training has been challenging in ST. This inspires us to propose a *Stacked Acoustic-and-Textual Encoding* method, which is straightforward to incorporate the pre-trained models into ST. We also introduce an adaptor module and a multi-teacher knowledge distillation method for bridging the gap between pre-training and fine-tuning.

Results on the LibriSpeech and MuST-C corpora demonstrate the superiority of our method. Furthermore, we achieve comparable or even better performance than the cascaded counterpart when large-scale ASR and MT data is available.

8 Acknowledgement

This work was supported in part by the National Science Foundation of China (Nos. 61876035 and 61732005), the National Key R&D Program of China (No. 2019QY1801), and the Ministry of Science and Technology of the PRC (Nos. 2019YFF0303002 and 2020AAA0107900). The authors would like to thank anonymous reviewers for their 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 2020, Online, November 16-20, 2020*, pages 8014–8020. Association for Computational Linguistics.

Antonios Anastasopoulos and David Chiang. 2018. [Tied multitask learning for neural speech translation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Lan-*

guage Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers), pages 82–91. Association for Computational Linguistics.

Parnia Bahar, Tobias Bieschke, and Hermann Ney. 2019. [A comparative study on end-to-end speech to text translation](#). In *IEEE Automatic Speech Recognition and Understanding Workshop, ASRU 2019, Singapore, December 14-18, 2019*, pages 792–799. IEEE.

Sameer Bansal, Herman Kamper, Karen Livescu, Adam Lopez, and Sharon Goldwater. 2019. [Pre-training on high-resource speech recognition improves low-resource speech-to-text translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 58–68. Association for Computational Linguistics.

Alexandre Berard, Laurent Besacier, Ali Can Kocabiyikoglu, and Olivier Pietquin. 2018. [End-to-end automatic speech translation of audio-books](#). In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2018, Calgary, AB, Canada, April 15-20, 2018*, pages 6224–6228. IEEE.

Alexandre Berard, Olivier Pietquin, Christophe Servan, and Laurent Besacier. 2016. [Listen and translate: A proof of concept for end-to-end speech-to-text translation](#). *CoRR*, abs/1612.01744.

Qianqian Dong, Mingxuan Wang, Hao Zhou, Shuang Xu, Bo Xu, and Lei Li. 2021. [Consecutive decoding for speech-to-text translation](#). In *The Thirty-fifth AAAI Conference on Artificial Intelligence, AAAI*.

Long Duong, Antonios Anastasopoulos, David Chiang, Steven Bird, and Trevor Cohn. 2016. [An attentional model for speech translation without transcription](#). In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 949–959. The Association for Computational Linguistics.

- Mattia Antonino Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019. [Must-c: a multilingual speech translation corpus](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 2012–2017. Association for Computational Linguistics.
- Ian J. Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. 2015. [An empirical investigation of catastrophic forgetting in gradient-based neural networks](#).
- Alex Graves, Santiago Fernández, Faustino J. Gomez, and Jürgen Schmidhuber. 2006. [Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks](#). In *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006*, volume 148 of *ACM International Conference Proceeding Series*, pages 369–376. ACM.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. 2020. [Conformer: Convolution-augmented transformer for speech recognition](#). In *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, pages 5036–5040. ISCA.
- Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. [Distilling the knowledge in a neural network](#). *CoRR*, abs/1503.02531.
- 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, ACL 2020, Online, July 5-10, 2020*, pages 302–311. Association for Computational Linguistics.
- Shigeki Karita, Nelson Enrique Yalta Soplin, Shinji Watanabe, Marc Delcroix, Atsunori Ogawa, and Tomohiro Nakatani. 2019. [Improving transformer-based end-to-end speech recognition with connectionist temporal classification and language model integration](#). In *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, pages 1408–1412. ISCA.
- Ali Can Kocabiyikoglu, Laurent Besacier, and Olivier Kraif. 2018. [Augmenting librispeech with french translations: A multimodal corpus for direct speech translation evaluation](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation, LREC 2018, Miyazaki, Japan, May 7-12, 2018*. European Language Resources Association (ELRA).
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *ACL 2007, Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, June 23-30, 2007, Prague, Czech Republic*. The Association for Computational Linguistics.
- Bei Li, Ziyang Wang, Hui Liu, Quan Du, Tong Xiao, Chunliang Zhang, and Jingbo Zhu. 2020. [Learning light-weight translation models from deep transformer](#). *CoRR*, abs/2012.13866.
- 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 *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, pages 1128–1132. ISCA.
- Yuchen Liu, Junnan Zhu, Jiajun Zhang, and Chengqing Zong. 2020. [Bridging the modality gap for speech-to-text translation](#). *CoRR*, abs/2010.14920.
- Lambert Mathias and William Byrne. 2006. [Statistical phrase-based speech translation](#). In *2006 IEEE International Conference on Acoustics Speech and Signal Processing, ICASSP 2006*,

- Toulouse, France, May 14-19, 2006, pages 561–564. IEEE.
- Hermann Ney. 1999. [Speech translation: coupling of recognition and translation](#). In *Proceedings of the 1999 IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP '99, Phoenix, Arizona, USA, March 15-19, 1999*, pages 517–520. IEEE Computer Society.
- Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le. 2019. [SpecAugment: A simple data augmentation method for automatic speech recognition](#). In *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, pages 2613–2617. ISCA.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 2227–2237. Association for Computational Linguistics.
- Juan Pino, Liezl Puzon, Jiatao Gu, Xutai Ma, Arya D McCarthy, and Deepak Gopinath. 2019. [Harnessing indirect training data for end-to-end automatic speech translation: Tricks of the trade](#). *arXiv preprint arXiv:1909.06515*.
- Juan Pino, Qiantong Xu, Xutai Ma, Mohammad Javad Dousti, and Yun Tang. 2020. [Self-training for end-to-end speech translation](#). In *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, pages 1476–1480. ISCA.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, pages 186–191. Association for Computational Linguistics.
- Tanja Schultz, Szu-Chen Jou, Stephan Vogel, and Shirin Saleem. 2004. [Using word lattice information for a tighter coupling in speech translation systems](#). In *Eighth International Conference on Spoken Language Processing*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Matthias Sperber, Graham Neubig, Jan Niehues, and Alex Waibel. 2019. [Attention-passing models for robust and data-efficient end-to-end speech translation](#). *Trans. Assoc. Comput. Linguistics*, 7:313–325.
- Matthias Sperber, Jan Niehues, Graham Neubig, Sebastian Stüker, and Alex Waibel. 2018. [Self-attentional acoustic models](#). In *Interspeech 2018, 19th Annual Conference of the International Speech Communication Association, Hyderabad, India, 2-6 September 2018*, pages 3723–3727. ISCA.
- Mihaela C. Stoian, Sameer Bansal, and Sharon Goldwater. 2020. [Analyzing ASR pretraining for low-resource speech-to-text translation](#). In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, pages 7909–7913. IEEE.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Chengyi Wang, Yu Wu, Shujie Liu, Zhenglu Yang, and Ming Zhou. 2020a. [Bridging the gap between pre-training and fine-tuning for end-to-end speech translation](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI*

2020, New York, NY, USA, February 7-12, 2020, pages 9161–9168. AAAI Press.

Chengyi Wang, Yu Wu, Shujie Liu, Ming Zhou, and Zhenglu Yang. 2020b. [Curriculum pre-training for end-to-end speech translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 3728–3738. Association for Computational Linguistics.

Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplín, Jahn Heymann, Matthew Wiesner, Nanxin Chen, Adithya Renduchintala, and Tsubasa Ochiai. 2018. [Espnet: End-to-end speech processing toolkit](#). In *Interspeech 2018, 19th Annual Conference of the International Speech Communication Association, Hyderabad, India, 2-6 September 2018*, pages 2207–2211. ISCA.

Shinji Watanabe, Takaaki Hori, Suyoun Kim, John R. Hershey, and Tomoki Hayashi. 2017. [Hybrid ctc/attention architecture for end-to-end speech recognition](#). *IEEE J. Sel. Top. Signal Process.*, 11(8):1240–1253.

Ron J. Weiss, Jan Chorowski, Navdeep Jaitly, Yonghui Wu, and Zhifeng Chen. 2017. [Sequence-to-sequence models can directly translate foreign speech](#). In *Interspeech 2017, 18th Annual Conference of the International Speech Communication Association, Stockholm, Sweden, August 20-24, 2017*, pages 2625–2629. ISCA.

Baosong Yang, Zhaopeng Tu, Derek F. Wong, Fandong Meng, Lidia S. Chao, and Tong Zhang. 2018. [Modeling localness for self-attention networks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 4449–4458. Association for Computational Linguistics.

Jiacheng Yang, Mingxuan Wang, Hao Zhou, Chengqi Zhao, Weinan Zhang, Yong Yu, and Lei Li. 2020. [Towards making the most of BERT in neural machine translation](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence*

Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020, pages 9378–9385. AAAI Press.