

Bridging the Gap between Pre-Training and Fine-Tuning for End-to-End Speech Translation

Chengyi Wang,^{1*} Yu Wu,² Shujie Liu,² Zhenglu Yang,¹ Ming Zhou²

¹Nankai University, Tianjin, China

²Microsoft Research Aisa, Beijing, China

cywang@mail.nankai.edu.cn, {Wu.Yu, shujliu, mingzhou}@microsoft.com, yangzl@nankai.edu.cn

Abstract

End-to-end speech translation, a hot topic in recent years, aims to translate a segment of audio into a specific language with an end-to-end model. Conventional approaches employ multi-task learning and pre-training methods for this task, but they suffer from the huge gap between pre-training and fine-tuning. To address these issues, we propose a Tandem Connectionist Encoding Network (TCEN) which bridges the gap by reusing all subnets in fine-tuning, keeping the roles of subnets consistent, and pre-training the attention module. Furthermore, we propose two simple but effective methods to guarantee the speech encoder outputs and the MT encoder inputs are consistent in terms of semantic representation and sequence length. Experimental results show that our model leads to significant improvements in En-De and En-Fr translation irrespective of the backbones.

1 Introduction

Speech-to-Text translation (ST) is essential for a wide range of scenarios: for example in emergency calls, where agents have to respond emergent requests in a foreign language (Munro 2010); or in online courses, where audiences and speakers use different languages (Jan et al. 2018). To tackle this problem, existing approaches can be categorized into cascaded method (Ney 1999; Ma et al. 2019), where a machine translation (MT) model translates outputs of an automatic speech recognition (ASR) system into target language, and end-to-end method (Duong et al. 2016; Weiss et al. 2017), where a single model learns acoustic frames to target word sequence mappings in one step towards the final objective of interest. Although the cascaded model remains the dominant approach due to its better performance, the end-to-end method becomes more and more popular because it has lower latency by avoiding inferences with two models and rectifies the error propagation in theory.

Since it is hard to obtain a large-scale ST dataset, multi-task learning (Weiss et al. 2017; Bérard et al. 2018) and pre-training techniques (Bansal et al. 2019) have been applied to end-to-end ST model to leverage large-scale datasets of ASR

and MT. A common practice is to pre-train two encoder-decoder models for ASR and MT respectively, and then initialize the ST model with the encoder of the ASR model and the decoder of the MT model. Subsequently, the ST model is optimized with the multi-task learning by weighing the losses of ASR, MT, and ST. This approach, however, causes a huge gap between pre-training and fine-tuning, which are summarized into three folds:

- **Subnet Waste:** The ST system just reuses the ASR encoder and the MT decoder, while discards other pre-trained subnets, such as the MT encoder. Consequently, valuable semantic information captured by the MT encoder cannot be inherited by the final ST system.
- **Role Mismatch:** The speech encoder plays different roles in pre-training and fine-tuning. The encoder is a pure acoustic model in pre-training, while it has to extract semantic and linguistic features additionally in fine-tuning, which significantly increases the learning difficulty.
- **Non-pre-trained Attention Module:** Previous work (Bérard et al. 2018) trains attention modules for ASR, MT and ST respectively, hence, the attention module of ST does not benefit from the pre-training.

To address these issues, we propose a Tandem Connectionist Encoding Network (TCEN), which is able to reuse all subnets in pre-training, keep the roles of subnets consistent, and pre-train the attention module. Concretely, the TCEN consists of three components, a speech encoder, a text encoder, and a target text decoder. Different from the previous work that pre-trains an encoder-decoder based ASR model, we only pre-train an ASR encoder by optimizing the Connectionist Temporal Classification (CTC) (Graves et al. 2006) objective function. In this way, the additional decoder of ASR is not required while keeping the ability to read acoustic features into the source language space by the speech encoder. Besides, the text encoder and decoder can be pre-trained on a large MT dataset. After that, we employ common used multi-task learning method to jointly learn ASR, MT and ST tasks.

Compared to prior works, the encoder of TCEN is a concatenation of an ASR encoder and an MT encoder and our model does not have an ASR decoder, so the **subnet waste**

*Works are done during internship at Microsoft
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: An illustration of multi-task learning for speech translation. Networks inherited from pre-trained models are labeled by rectangles.

issue is solved. Furthermore, the two encoders work at tandem, disentangling acoustic feature extraction and linguistic feature extraction, ensuring the **role consistency** between pre-training and fine-tuning. Moreover, we **reuse the pre-trained MT attention module** in ST, so we can leverage the alignment information learned in pre-training.

Since the text encoder consumes word embeddings of plausible texts in MT task but uses speech encoder outputs in ST task, another question is how one guarantees the speech encoder outputs are consistent with the word embeddings. We further modify our model to achieve **semantic consistency** and **length consistency**. Specifically, (1) the projection matrix at the CTC classification layer for ASR is shared with the word embedding matrix, ensuring that they are mapped to the same latent space, and (2) the length of the speech encoder output is proportional to the length of the input frame, so it is much longer than a natural sentence. To bridge the length gap, source sentences in MT are lengthened by adding word repetitions and blank tokens to mimic the CTC output sequences.

We conduct comprehensive experiments on the IWSLT18 speech translation benchmark (Jan et al. 2018), demonstrating the effectiveness of each component. Our model can lead to significant improvements for both LSTM and Transformer backbone.

Our contributions are three-folds: 1) we shed light on why previous ST models cannot sufficiently utilize the knowledge learned from the pre-training process; 2) we propose a new ST model, which alleviates shortcomings in existing methods; and 3) we empirically evaluate the proposed model on a large-scale public dataset.

2 Background

End-to-end speech translation aims to translate a piece of audio into a target-language translation in one step. The raw speech signals are usually converted to sequences of acoustic features. Here, we define the speech feature sequence as $\mathbf{x} = (x_1, \dots, x_{T_x})$.

The transcription and translation sequences are denoted as $\mathbf{y}^s = (y_1^s, \dots, y_{T_s}^s)$, and $\mathbf{y}^t = (y_1^t, \dots, y_{T_t}^t)$ respectively.

Each symbol in \mathbf{y}^s or \mathbf{y}^t is an integer index of the symbol in a vocabulary V_{src} or V_{trg} respectively (e.g. $y_i^s = k, k \in [0, |V_{src}| - 1]$). In this work, we suppose that an ASR dataset, an MT dataset, and a ST dataset are available, denoted as $\mathcal{A} = \{(\mathbf{x}_i, \mathbf{y}_i^s)\}_{i=0}^I$, $\mathcal{M} = \{(\mathbf{y}_j^s, \mathbf{y}_j^t)\}_{j=0}^J$ and $\mathcal{S} = \{(\mathbf{x}_l, \mathbf{y}_l^t)\}_{l=0}^L$ respectively. Given a new piece of audio \mathbf{x} , our goal is to learn an end to end model to generate a translation sentence \mathbf{y}^t without generating an intermediate result \mathbf{y}^s .

2.1 Multi-Task Learning and Pre-training for ST

To leverage large scale ASR and MT data, multi-task learning and pre-training techniques are widely employed to improve the ST system. As shown in Figure 1, there are three popular multi-task strategies for ST, including 1) one-to-many setting, in which a speech encoder is shared between ASR and ST tasks; 2) many-to-one setting in which a decoder is shared between MT and ST tasks; and 3) many-to-many setting where both the encoder and decoder are shared.

A many-to-many multi-task model contains two encoders as well as two decoders. It can be jointly trained on ASR, MT, and ST tasks. As the attention module is task-specific, three attentions are defined.

3 Our method

In this section, we first introduce the architecture of TCEN, which consists of two encoders connected in tandem, and one decoder with an attention module. Then we give the pre-training and fine-tuning strategy for TCEN. Finally, we propose our solutions for semantic and length inconsistency problems, which are caused by multi-task learning.

3.1 Unified formulation for TCEN Architecture

Figure 2 sketches the overall architecture of TCEN, including a speech encoder enc_s , a text encoder enc_t and a decoder dec with an attention module att . The enc_s usually contains two modules: EncPre and EncBody. During training, the enc_s acts like an acoustic model which reads the input \mathbf{x} to

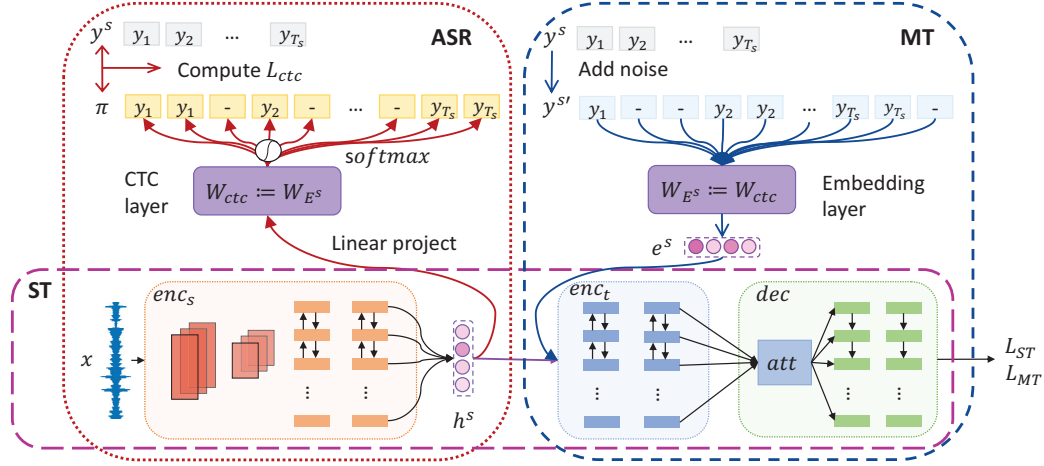


Figure 2: The architecture of our model. The linear projection matrix in ASR is shared with the word embedding matrix in MT.

word or subword representations h^s :

$$\tilde{x} = \text{EncPre}(x) \quad (1)$$

$$h^s = \text{EncBody}(\tilde{x}) \quad (2)$$

Then enc_t learns high-level linguistic knowledge into hidden representations h^t :

$$h^t = enc_t(h^s) \quad (3)$$

Finally, the dec defines a distribution probability over target words through attention mechanism:

$$c_k = \text{att}(z_{k-1}, h^t) \quad (4)$$

$$z_k = \text{dec}(z_{k-1}, y_{k-1}^t, c_k) \quad (5)$$

$$P(y_k^t | y_{<k}^t, x) = \text{softmax}(W \cdot z_k) \quad (6)$$

Here, z_k is the the hidden state of the decoder at k step and c_k is a time-dependent context vector computed by the attention att .

The advantage of our architecture is that two encoders disentangle acoustic feature extraction and linguistic feature extraction, making sure that valuable knowledge learned from ASR and MT tasks can be effectively leveraged for ST training. However, there exists another problem: In ST task, enc_t accepts speech encoder output h^s as input. While in MT, enc_t consumes the word embedding representation e^s derived from y^s , where each element e_i^s is computed by choosing the y_i^s -th vector from the source embedding matrix W_{E^s} . Since h^s and e^s belong to different latent space and have different lengths, there remain semantic and length inconsistency problems. We will provide our solutions in Section 3.3. To verify the generalization of our framework, we test on LSTM based setting and Transformer (Vaswani et al. 2017) based setting.

3.2 Training Procedure

Following previous work, we split the training procedure to pre-training and fine-tuning stages. In pre-training stage, the speech encoder enc_s is trained towards CTC objective using

Transcript y^s	we were not v @en @ge @ful at all
CTC path π_1	-(11) we we -(3) were -(3) not -(4) v @en @en @ge - @ful -(8) at at -(3) all -(10)
CTC path π_2	-(9) we -(3) were were -(4) not not -(3) v v @en @en @en @ge - @ful -(7) at -(3) all all -(10)

Table 1: An example of the comparison between the golden transcript and the predicted CTC paths given the corresponding speech. ‘-’ denotes the blank token and the following number represents repeat times.

dataset \mathcal{A} , while the text encoder enc_t and the decoder dec are trained on MT dataset \mathcal{M} . In fine-tuning stage, we jointly train the model on ASR, MT, and ST tasks.

Pre-training To sufficiently utilize the large dataset \mathcal{A} and \mathcal{M} , the model is pre-trained on CTC-based ASR task and MT task in the pre-training stage.

For ASR task, in order to get rid of the requirement for decoder and enable the enc_s to generate subword representation, we leverage connectionist temporal classification (CTC) (Graves et al. 2006) loss to train the speech encoder.

Given an input x , enc_s emits a sequence of hidden vectors h^s , then a softmax classification layer predicts a CTC path π , where $\pi_t \in V_{src} \cup \{‘-’\}$ is the observing label at particular RNN step t , and ‘-’ is the blank token representing no observed labels:

$$P(\pi|x) \approx \prod_{t=1}^T P(\pi_t|x) = \prod_{t=1}^T \text{softmax}(W_{ctc} \cdot h_t^s) \quad (7)$$

where $W_{ctc} \in \mathbb{R}^{d \times (|V_{src}|+1)}$ is the weight matrix in the classification layer and T is the total length of encoder RNN.

A legal CTC path π is a variation of the source transcription y^s by allowing occurrences of blank tokens and repetitions, as shown in Table 1. For each transcription y , there exist many legal CTC paths in length T . The CTC objective trains the model to maximize the probability of observing the golden sequence y^s , which is calculated by summing

the probabilities of all possible legal paths:

$$P(\mathbf{y}|\mathbf{x}) = \sum_{\pi \in \Phi_T(\mathbf{y})} P(\pi|\mathbf{x}) \quad (8)$$

$$\mathcal{L}_{CTC}(\theta) = - \sum_{(\mathbf{x}, \mathbf{y}^s) \in \mathcal{A}} \log P(\mathbf{y}^s|\mathbf{x}; \theta_{enc_s}, \theta_{W_{ctc}}) \quad (9)$$

where $\Phi_T(\mathbf{y})$ is the set of all legal CTC paths for sequence \mathbf{y} with length T . The loss can be easily computed using forward-backward algorithm. More details about CTC are provided in supplementary material.

For MT task, we use the cross-entropy loss as the training objective. During training, \mathbf{y}^s is converted to embedding vectors \mathbf{e}^s through embedding layer W_{E^s} , then enc_t consumes \mathbf{e}^s and pass the output \mathbf{h}^t to decoder. The objective function is defined as:

$$\mathcal{L}_{MT}(\theta) = - \sum_{(\mathbf{y}^s, \mathbf{y}^t) \in \mathcal{M}} \log P(\mathbf{y}^t|\mathbf{y}^s; \theta_{enc_t}, \theta_{dec}, \theta_{W_{E^s}}) \quad (10)$$

Fine-tune In fine-tune stage, we jointly update the model on ASR, MT, and ST tasks. The training for ASR and MT follows the same process as it was in pre-training stage.

For ST task, the enc_s reads the input \mathbf{x} and generates \mathbf{h}^s , then enc_t learns high-level linguistic knowledge into \mathbf{h}^t . Finally, the dec predicts the target sentence. The ST loss function is defined as:

$$\mathcal{L}_{ST}(\theta) = - \sum_{(\mathbf{x}, \mathbf{y}^t) \in \mathcal{S}} \log P(\mathbf{y}^t|\mathbf{x}; \theta_{enc_s}, \theta_{enc_t}, \theta_{dec}) \quad (11)$$

Following the update strategy proposed by Luong et al. (2016), we allocate a different training ratio α_i for each task. When switching between tasks, we select randomly a new task i with probability $\frac{\alpha_i}{\sum_j \alpha_j}$.

3.3 Subnet-Consistency

Our model keeps role consistency between pre-training and fine-tuning by connecting two encoders for ST task. However, this leads to some new problems: 1) The text encoder consumes \mathbf{e}^s during MT training, while it accepts \mathbf{h}^s during ST training. However, \mathbf{e}^s and \mathbf{h}^s may not follow the same distribution, resulting in the semantic inconsistency. 2) Besides, the length of \mathbf{h}^s is not the same order of magnitude with the length of \mathbf{e}^s , resulting in the length inconsistency.

In response to the above two challenges, we propose two countermeasures: 1) We share weights between CTC classification layer and source-end word embedding layer during training of ASR and MT, encouraging \mathbf{e}^s and \mathbf{h}^s in the same space. 2) We feed the text encoder source sentences in the format of CTC path, which are generated from a seq2seq model, making it more robust toward long inputs.

Semantic Consistency As shown in Figure 2, during multi-task training, two different hidden features will be fed into the text encoder enc_t : the embedding representation \mathbf{e}^s in MT task, and the enc_s output \mathbf{h}^s in ST task. Without any regularization, they may belong to different latent spaces. Due to the space gap, the enc_t has to compromise between two tasks, limiting its performance on individual tasks.

To bridge the space gap, our idea is to pull \mathbf{h}^s into the latent space where \mathbf{e}^s belong. Specifically, we share the




Figure 3: The architecture of seq2seq model. It predicts the next token and its number of repetition at the same time.

weight W_{ctc} in CTC classification layer with the source embedding weights W_{E^s} , which means $W_{ctc} = W_{E^s}$. In this way, when predicting the CTC path π , the probability of observing the particular label $w_i \in V_{src} \cup \{ '- ' \}$ at time step t , $p(\pi_t = w_i|\mathbf{x})$, is computed by normalizing the product of hidden vector \mathbf{h}_t^s and the i -th vector in W_{E^s} :

$$P(\pi_t = w_i|\mathbf{x}) = \frac{\exp(W_{E^s}^{(i)} \cdot \mathbf{h}_t^s)}{\sum_j^{|V_{src}|+1} \exp(W_{E^s}^{(j)} \cdot \mathbf{h}_t^s)} \quad (12)$$

The loss function closes the distance between \mathbf{h}_t^s and golden embedding vector, encouraging \mathbf{h}^s have the same distribution with \mathbf{e}^s .

Length Consistency Another existing problem is length inconsistency. The length of the sequence \mathbf{h}^s is proportional to the length of the input frame \mathbf{x} , which is much longer than the length of \mathbf{e}^s . To solve this problem, we train an RNN-based seq2seq model to transform normal source sentences to noisy sentences in CTC path format, and replace standard MT with denoising MT for multi-tasking.

Specifically, we first train a CTC ASR model based on dataset $\mathcal{A} = \{(\mathbf{x}_i, \mathbf{y}_i^s)\}_{i=0}^I$, and generate a CTC-path π_i for each audio \mathbf{x}_i by greedy decoding. Then we define an operation $S(\cdot)$, which converts a CTC path π to a sequence of the unique tokens \mathbf{u} and a sequence of repetition times for each token l , denoted as $S(\pi) = (\mathbf{u}, \mathbf{l})$. Notably, the operation is reversible, meaning that $S^{-1}(\mathbf{u}, \mathbf{l}) = \pi$. We use the example π_1 in Table 1 and show the corresponding \mathbf{u} and \mathbf{l} in Table 2.

Then we build a dataset $\mathcal{P} = \{(\mathbf{y}_i^s, \mathbf{u}_i, \mathbf{l}_i)\}_{i=0}^I$ by decoding all the audio pieces in \mathcal{A} and transform the resulting path by the operation $S(\cdot)$. After that, we train a seq2seq model, as shown in Figure 3, which takes \mathbf{y}_i^s as input and decodes $\mathbf{u}_i, \mathbf{l}_i$ as outputs. With the seq2seq model, a noisy MT dataset $\mathcal{M}' = \{(\pi_i, \mathbf{y}_i^t)\}_{i=0}^L$ is obtained by converting every source sentence $\mathbf{y}_i^s \in \mathcal{M}$ to π_i , where $\pi_i = S^{-1}(\mathbf{u}_i, \mathbf{l}_i)$. We did not use the standard seq2seq model which takes \mathbf{y}^s as input and generates π directly, since there are too many blank tokens ' - ' in π and the model tends to generate a long sequence with only blank tokens. During MT training, we randomly sample text pairs from \mathcal{M}' and \mathcal{M} according to a hyper-parameter k . After tuning on the validation set, about

CTC path π_1	-(11) we we -(3) were -(3) not -(4) v @en @en @ge - @ful -(8) at at -(3) all -(10)																
u	-	we	-	were	-	not	-	v	@en	@ge	-	@ful	-	at	-	all	-
l	11	2	3	1	3	1	4	1	2	1	1	1	8	2	3	1	10

Table 2: The CTC path π_1 and corresponding unique tokens u and repetition times l , where $S(\pi) = (u, l)$.

30% pairs are sampled from \mathcal{M}' . In this way, the enc_t is more robust toward the longer inputs given by the enc_s .

4 Experiments

4.1 Dataset

We conduct experiments on the Speech Translation TED (ST-TED) En-De corpus (Jan et al. 2018) and the augmented Librispeech En-Fr corpus (Kocabiyikoglu, Besacier, and Kraif 2018).

ST-TED En-De The corpus contains 272 hours of English speech with 171k segments. Each example consists of raw English wave, English transcription, and aligned German translation. Aside from ST-TED, we use TED-LIUM2 corpus (Rousseau, Deléglise, and Esteve 2014) with 207h of speech data for ASR pre-training. For MT model, we use WMT2018 en-de data in pre-training stage and use sentence pairs in the ST-TED corpus as well as WIT3¹ in fine-tune stage. The pre-training data contains 41M sentence pairs and fine-tuning data contains 330k sentence paris in total. We split 2k segments from the ST-TED corpus as dev set and tst2010, tst2013, tst2014, tst2015 are used as test sets.

Librispeech En-Fr This corpus is colleted by aligning e-books in French with English utterances, which contains 236 hours of speech in total. The English speech, English transcription, French text translations from alignment and Google Translate references are provided. Following previous work (Bérard et al. 2018), we only use the 100 hours clean train set and double the training size by concatenating the aligned referenes with Google Translate references. We use the speech-transcription pairs and transcription-translation pairs for ASR and MT pre-training. No additional data is used. The dev set is used as validation set and we report results on the test set.

Data preprocessing Our acoustic features are 80-dimensional log-Mel filterbanks and 3-dimensional pitch features extracted with a step size of 10ms and window size of 25ms and extended with mean subtraction and variance normalization. The utterances with more than 3000 frames are discarded. All the sentences are in lower-case and the punctuation is removed. To increase the amount of training data, we perform speed perturbation on the raw signals with speed factors 0.9 and 1.1.

For the MT pre-training data, sentences longer than 80 words or shorter than 10 words are removed. Besides, we discard pairs whose length ratio between source and target sentence is smaller than 0.5 or larger than 2.0. Word

¹<https://wit3.fbk.eu/mt.php?release=2017-01-trnted>

tokenization is performed using the Moses scripts² and all words are in lower-case.

For ST-TED experiments, we apply both subword-level decoding and character-level decoding. For the subword setting, both English and German vocabularies are generated using sentencepiece³ (Kudo 2018) with a fixed size of 5k tokens. For Librispeech En-Fr experiments, we only apply character-level decoding.

Since there are no human annotated alignments provided in ST-TED test sets, we segment each audio with the LIUM SpkDiarization tool (Meignier and Merlin 2010) and then perform MWER segmentation with RWTH toolkit (Bender et al. 2004). Case-insensitive BLEU is used as evaluation metric.

4.2 Experimental setups

Model architecture For LSTM based models, we follow the model structure in Inaguma et al. (2018). The EncPre corresponds to 2-layers of VGG-like max-pooling, resulting 4-fold downsampling of input feature. The EncBody is five bidirectional LSTM layers with cell size of 1024. The decoder is defined as two unidirectional LSTM layers with an additive attention. The decoder has the same dimension as the encoder RNNs.

For Transformer based models, we use two-layer CNN with 256 channels, stride size 2 and kernel size 3 as EncPre. The other modules are similar as in paper Dong, Xu, and Xu (2018) ($e = 12, d = 6, d_{model} = 256, d_{ff} = 2048$ and $d_{head} = 4$).

Baselines We compare our method with the following baselines.

- Vanilla ST baseline: The vanilla ST has only a speech encoder and a decoder, which is trained from scratch with only the speech-translation data.
- Pre-training baselines: We conduct three pre-training baseline experiments: 1) encoder pre-training, 2) decoder pre-training, and 3) encoder-decoder pre-training. The pre-trained ASR model has the same architecture with vanilla ST model. The MT model has a enc_t and dec with the same architecture of which in TCEN.
- Pre-training + MTL: In this setting, we train a many-to-many multi-task model where the encoders and decoders are derived from pre-trained ASR and MT models.

Implementation All our models are implemented based on ESPnet (Watanabe et al. 2018). For LSTM based models, we use a dropout of 0.3 for embedding and encoders. The

²<https://github.com/moses-smt/mosesdecoder/blob/master/scripts/tokenizer/tokenizer.perl>

³<https://github.com/google/sentencepiece>

	Subword Level Decoder					Char Level Decoder				
	tst2010	tst2013	tst2014	tst2015	Average	tst2010	tst2013	tst2014	tst2015	Average
Vanilla	7.52	7.04	6.77	6.57	6.98	13.77	12.50	11.50	12.68	12.61
+enc pretrain	10.70	10.12	8.82	7.76	9.35	16.00	14.49	12.66	12.20	13.76
+dec pretrain	9.75	9.02	8.34	8.01	8.78	14.44	12.99	11.91	12.87	13.05
+enc dec pretrain	12.14	11.07	9.96	8.77	10.49	15.52	14.62	13.39	13.33	14.22
pretrain+MTL	11.92	11.78	9.89	9.27	10.72	15.70	15.42	13.43	12.66	14.30
Triangle+pretrain	9.89	9.91	7.48	7.22	8.63	11.35	10.73	9.43	9.47	10.25
TCEN-LSTM	15.49	15.50	13.21	13.02	14.31	17.61	17.67	15.73	14.94	16.49

Table 3: Results of LSTM-based models on ST TED. ‘‘Average’’ denotes it averages the results of all test sets. We copy the numbers of vanilla model from <https://github.com/espnet/espnet/blob/master/egs/iwslt18/st1/RESULTS>. Since pre-training data is different, we run ESPnet code to obtain the numbers of pre-training and multi-task learning method, which are slightly higher than numbers in their report.

model is trained using Adadelta with initial learning rate of 1.0.

For Transformer based model, we use a dropout rate of 0.1 and a gradient clip of 5.0. Following Dong, Xu, and Xu(2018), we use Adam optimizer according to the learning rate schedule formula:

$$\text{lrate} = k \cdot d_{\text{model}}^{-0.5} \cdot \min(n^{-0.5}, n \cdot \text{warmup}_n^{-1.5})$$

We set $k = 10$ and $\text{warmup}_n = 25000$ in our experiments. All the models are trained on 4 Tesla P40 GPU for a maximum of 20 epochs.

For training of TCEN, we set $\alpha_{\text{asr}} = 0.2$ and $\alpha_{\text{mt}} = 0.8$ in the pre-training stage, since the MT dataset is much larger than ASR dataset. For fine-tune, we use $\alpha_{\text{st}} = 0.6$, $\alpha_{\text{asr}} = 0.2$ and $\alpha_{\text{mt}} = 0.2$. At inference time, we use a beam size of 10 and a length normalization weight of 0.2.

4.3 Experimental Results

Results on ST TED Table 3 shows the LSTM-based results on four test sets as well as the average performance. In this setting, we also re-implement the triangle multi-task strategy (Anastasopoulos and Chiang 2018) as our baseline, denoted as ‘triangle+pretrain’. They concatenate a ST decoder to an ASR encoder-decoder model.

From the table, we can see that our method significantly outperforms the strong ‘pretrain+MTL’ baseline by 3.6 and 2.2 BLEU scores respectively, indicating the proposed method is very effective. Besides, both pre-training and multi-task learning can improve translation quality. We observe a performance degradation in the ‘triangle+pretrain’ baseline. Compared to our method, where the decoder receives higher-level knowledge extracted from text encoder, their ASR decoder can only provide lower word-level linguistic information. Besides, their model cannot utilize the large-scale MT data in all the training stages.

Table 4 shows the comparison between our best model with the cascaded systems, which combines the ASR model and MT model. In addition to a simple combination system, we also re-segment the ASR outputs before feeding to the MT system, denoted as ‘cascaded+re-seg’. Specifically, we train a seq2seq model (Bahdanau, Cho, and Bengio 2015) on the MT dataset, where the source side is a no punctuation sentence and the target side is a natural sentence. After that, we use the seq2seq model to add sentence boundaries and

	tst2010	tst2013	tst2014	tst2015
cascaded	13.38	15.84	12.94	13.79
cascaded+re-seg	17.12	17.77	14.94	15.01
our model	17.61	17.67	15.73	14.94

Table 4: BLEU comparison of cascaded results and our best end-to-end results. re-seg denotes the ASR outputs are re-segmented before fed into the MT model.

System	tst2013
Vanilla	-
+enc pretrain	13.41
+enc dec pretrain	14.46
pretrain+MTL	14.98
TCEN-Transformer	17.11

Table 5: BLEU of Transformer-based models on tst2013 set. ‘-’: failed training.

punctuation on ASR outputs. It can be seen that our end-to-end model outperforms the simple cascaded model over 2 BLEU scores, and achieves a comparable performance with the ‘cascaded+re-seg’ system.

We list Transformer-based results on tst2013 in Table 5. In this setting, we use character-level decoding strategy due to its better performance. Only in-domain MT data is used during pre-training. It can be seen that our TCEN framework works well on Transformer-based architecture and it outperforms the ‘pretrain+MTL’ baseline by 2.1 BLEU scores.

Results on Librispeech For this dataset, we only perform LSTM-based experiments and report results in Table 6. Even without utilizing large-scale ASR data or MT data, our method can outperform the pre-training baselines and achieve the same performance with Park et al. (2019), which uses a MT model as a teacher model to guide the ST model.

4.4 Discussion

Ablation Study To better understand the contribution of each component, we perform an ablation study on subword-level experiments for ST TED corpus. The results are shown in Table 7. In ‘-MT noise’ setting, we do not add noise to source sentences for MT. In ‘-weight sharing’ setting, we use different parameters in CTC classification layer and

Model		BLEU
MT	Bérard et al.(2018)	19.2
	ESPnet*	18.3
Cascaded	Bérard et al.(2018)	14.6
	ESPnet*	15.8
E2E ST	Bérard et al.(2018)	12.9
	+Pretrain+MTL	13.4
	Liu et al.(2019)	17.02
	ESPnet*	15.71
	+enc pretrain	16.30
	+enc dec pretrain	16.78
	TCEN-LSTM	17.05

Table 6: BLEU results of LSTM-based models on Librispeech En-Fr. *: The ESPnet baseline results are copied from https://github.com/espnet/espnet/blob/master/egs/libri_trans

System	tst2010	tst2013	tst2014	tst2015
TCEN	15.49	15.50	13.21	13.02
-MT noise	15.01	14.95	13.34	12.80
-weight sharing	13.51	14.02	12.25	11.66
-pretrain	8.98	8.42	7.94	8.08

Table 7: Ablation study for subword-level experiments.

source embedding layer. These two experiments prove that both weight sharing and using noisy MT input benefit to the final translation quality. Performance degrades more in ‘-weight sharing’, indicating the semantic consistency contributes more to our model.

In the ‘-pretrain’ experiment, we remove the pre-training stage and directly update the model on three tasks, leading to a dramatic decrease on BLEU score, indicating the pre-training is an indispensable step for end-to-end ST.

Learning Curve It is interesting to investigate why our method is superior to baselines. We find that TCEN achieves a higher final result owing to a better start-point in fine-tuning. Figure 4 provides learning curves of subword accuracy on validation set. The x-axis denotes the fine-tuning training steps. The vanilla model starts at a low accuracy, be-




Figure 4: Model learning curves in fine-tuning.

cause its networks are not pre-trained on the ASR and MT data. The trends of our model and ‘many-to-many+pretrain’ are similar, but our model outperforms it about five points in the whole fine-tuning process. It indicates that the gain comes from bridging the gap between pre-training and fine-tuning rather than a better fine-tuning process.

5 Related Work

Early works conduct ST in a pipeline manner (Ney 1999; Matusov, Kanthak, and Ney 2005), where the ASR output are fed into an MT system to generate target sentences. HMM (Juang and Rabiner 1991), DenseNet (Huang et al. 2017), TDNN (Peddinti, Povey, and Khudanpur 2015) are commonly used ASR systems, while RNN with attention (Bahdanau, Cho, and Bengio 2015) and Transformer (Vaswani et al. 2017) are top choices for MT.

To avoid error propagation and high latency issues, recent works propose translating the acoustic speech into text in target language without yielding the source transcription (Duong et al. 2016; Berard et al. 2016). Since ST data is scarce, pre-training (Bansal et al. 2019), multi-task learning (Duong et al. 2016; Bérard et al. 2018), curriculum learning (Kano, Sakti, and Nakamura 2018), attention-passing (Sperber et al. 2019), and knowledge distillation (Liu et al. 2019; Jia et al. 2019) strategies have been explored to utilize ASR data and MT data. Specifically, Weiss et al. (2017) show improvements of performance by training the ST model jointly with the ASR and the MT model. Bérard et al. (2018) observe faster convergence and better results due to pre-training and multi-task learning on a larger dataset. Bansal et al. (2019) show that pre-training a speech encoder on one language can improve ST quality on a different source language. All of them follow the traditional multi-task training strategies. Kano, Sakti, and Nakamura (2018) propose to use curriculum learning to improve ST performance on syntactically distant language pairs. To effectively leverage transcriptions in ST data, Anastasopoulos and Chiang (2018) augment the multi-task model where the target decoder receives information from the source decoder and they show improvements on low-resource speech translation. Jia et al. (2019) use pre-trained MT and text-to-speech (TTS) synthesis models to convert weakly supervised data into ST pairs and demonstrate that an end-to-end MT model can be trained using only synthesised data.

6 Conclusion

This paper has investigated the end-to-end method for ST. We propose a method to reuse every sub-net and keep the role of sub-net consistent between pre-training and fine-tuning, alleviating the gap between pre-training and fine-tuning in previous methods. Empirical studies have demonstrated that our model significantly outperforms baselines.

7 Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant No.U1636116, 11431006.

References

- Anastasopoulos, A., and Chiang, D. 2018. Tied multitask learning for neural speech translation. In *NAACL-HLT 2018*, 82–91.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR 2015*.
- 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 *NAACL-HLT 2019*, 58–68.
- Bender, O.; Zens, R.; Matusov, E.; and Ney, H. 2004. Alignment templates: the RWTH SMT system. In *IWSLT 2004*, 79–84.
- Berard, A.; Pietquin, O.; Servan, C.; and Besacier, L. 2016. Listen and translate: A proof of concept for end-to-end speech-to-text translation. In *NeurIPS Workshop on End-to-end Learning for Speech and Audio Processing*.
- Bérard, A.; Besacier, L.; Kocabiyikoglu, A. C.; and Pietquin, O. 2018. End-to-end automatic speech translation of audio-books. In *ICASSP 2018*, 6224–6228.
- Dong, L.; Xu, S.; and Xu, B. 2018. Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition. In *ICASSP 2018*, 5884–5888.
- Duong, L.; Anastasopoulos, A.; Chiang, D.; Bird, S.; and Cohn, T. 2016. An attentional model for speech translation without transcription. In *NAACL 2016*, 949–959.
- Graves, A.; Fernández, S.; Gomez, F. J.; and Schmidhuber, J. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *ICML 2006*, 369–376.
- Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *CVPR 2017*, 4700–4708.
- Inaguma, H.; Zhang, X.; Wang, Z.; Renduchintala, A.; Watanabe, S.; and Duh, K. 2018. The jhu/kyotou speech translation system for iwslt 2018. In *IWSLT 2018*.
- Jan, N.; Cattoni, R.; Sebastian, S.; Cettolo, M.; Turchi, M.; and Federico, M. 2018. The iwslt 2018 evaluation campaign. In *IWSLT*, 2–6.
- Jia, Y.; Johnson, M.; Macherey, W.; Weiss, R. J.; Cao, Y.; Chiu, C.; Ari, N.; Laurenzo, S.; and Wu, Y. 2019. Leveraging weakly supervised data to improve end-to-end speech-to-text translation. In *ICASSP 2019*, 7180–7184.
- Juang, B. H., and Rabiner, L. R. 1991. Hidden markov models for speech recognition. *Technometrics* 33(3):251–272.
- Kano, T.; Sakti, S.; and Nakamura, S. 2018. Structured-based curriculum learning for end-to-end english-japanese speech translation. *Interspeech 2017* 2630–2634.
- Kocabiyikoglu, A. C.; Besacier, L.; and Kraif, O. 2018. Augmenting librispeech with french translations: A multi-modal corpus for direct speech translation evaluation. In *LREC 2018*.
- Kudo, T. 2018. Subword regularization: Improving neural network translation models with multiple subword candidates. In *ACL 2018*, 66–75.
- Liu, Y.; Xiong, H.; He, Z.; Zhang, J.; Wu, H.; Wang, H.; and Zong, C. 2019. End-to-end speech translation with knowledge distillation. In *InterSpeech 2019*, volume abs/1904.08075.
- Luong, M.-T.; Le, Q. V.; Sutskever, I.; Vinyals, O.; and Kaiser, L. 2016. Multi-task sequence to sequence learning. *ICLR 2016*.
- Ma, M.; Huang, L.; Xiong, H.; Zheng, R.; Liu, K.; Zheng, B.; Zhang, C.; He, Z.; Liu, H.; Li, X.; Wu, H.; and Wang, H. 2019. STACL: simultaneous translation with implicit anticipation and controllable latency using prefix-to-prefix framework. In *ACL 2019*, 3025–3036.
- Matusov, E.; Kanthak, S.; and Ney, H. 2005. On the integration of speech recognition and statistical machine translation. In *INTERSPEECH 2005*, 3177–3180.
- Meignier, S., and Merlin, T. 2010. Lium spkdiarization: an open source toolkit for diarization. In *CMU SPUD Workshop*.
- Munro, R. 2010. Crowdsourced translation for emergency response in haiti: the global collaboration of local knowledge. In *AMTA Workshop on Collaborative Crowdsourcing for Translation*, 1–4.
- Ney, H. 1999. Speech translation: coupling of recognition and translation. In *ICASSP 1999*, 517–520.
- Park, D. S.; Chan, W.; Zhang, Y.; Chiu, C.; Zoph, B.; Cubuk, E. D.; and Le, Q. V. 2019. Specaugment: A simple data augmentation method for automatic speech recognition. *CoRR* abs/1904.08779.
- Peddinti, V.; Povey, D.; and Khudanpur, S. 2015. A time delay neural network architecture for efficient modeling of long temporal contexts. In *Sixteenth Annual Conference of the International Speech Communication Association*.
- 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 2014*, 3935–3939.
- Sperber, M.; Neubig, G.; Niehues, J.; and Waibel, A. 2019. Attention-passing models for robust and data-efficient end-to-end speech translation. *TACL* 7:313–325.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. In *NeurIPS 2017*, 5998–6008.
- Watanabe, S.; Hori, T.; Karita, S.; Hayashi, T.; Nishitoba, J.; Unno, Y.; Soplin, N. E. Y.; Heymann, J.; Wiesner, M.; Chen, N.; et al. 2018. Espnet: End-to-end speech processing toolkit. *arXiv preprint arXiv:1804.00015*.
- Weiss, R. J.; Chorowski, J.; Jaitly, N.; Wu, Y.; and Chen, Z. 2017. Sequence-to-sequence models can directly translate foreign speech. In *Interspeech 2017*, 2625–2629.