

On Using SpecAugment for End-to-End Speech Translation

Parnia Bahar^{1,2}, Albert Zeyer^{1,2}, Ralf Schlüter¹ and Hermann Ney^{1,2}

¹Human Language Technology and Pattern Recognition Group
Computer Science Department, RWTH Aachen University, 52062 Aachen, Germany

²AppTek, 52062 Aachen, Germany

{bahar, zeyer, schluter, ney}@cs.rwth-aachen.de

Abstract

This work investigates a simple data augmentation technique, SpecAugment, for end-to-end speech translation. SpecAugment is a low-cost implementation method applied directly to the audio input features and it consists of masking blocks of frequency channels, and/or time steps. We apply SpecAugment on end-to-end speech translation tasks and achieve up to +2.2% BLEU on LibriSpeech Audiobooks En→Fr and +1.2% on IWSLT TED-talks En→De by alleviating overfitting to some extent. We also examine the effectiveness of the method in a variety of data scenarios and show that the method also leads to significant improvements in various data conditions irrespective of the amount of training data.

1. Introduction

Traditional speech-to-text translation (ST) systems have been build in a cascaded fashion comprised of an automatic speech recognition (ASR) model trained on paired speech-transcribed data and a machine translation (MT) model trained on bilingual text data. Recent advancements in both ASR [1–6] and MT [7–11] have inspired the end-to-end direct ST models which can be trained using a translation speech corpus [12, 13]. Some appealing advantages of the direct models are: (1) no error accumulation from the recognizer, (2) faster decoding throughput and (3) less computational power in total by training all parameters jointly. In spite of these properties, training such end-to-end ST models requires a moderate amount of paired translated speech-to-text data which is not easy to acquire. Therefore these models tend to overfit easily.

In the absence of an adequate volume of training data, one remedy is generating synthetic data like back-translation (BT) [14] as the most common data augmentation method to leverage monolingual data. The idea is to use a pre-trained model to convert weakly supervised data into speech-to-translation pairs for ST training [15]. One way is to use a pre-trained source-to-target MT model to translate ASR transcription into the target language. Another method is the use of a pre-trained text-to-speech (TTS) model to generate speech data from a monolingual text. However, these methods require some effort to train an additional model, as well as computational power to generate a moderate amount of (noisy) synthetic data, which in some cases can be too ex-

pensive to be obtained.

Another method is data augmentation by which new synthetic training samples are generated by corrupting the initial audio data and conserving the same label as the original training sample. Audio-level speech augmentation can be done in different ways such as noise injection (adding random noise), shifting time (transmitting time series forward/backward with a few seconds), speed perturbation (expanding time series by a speed rate) and changing the frequency pitch randomly. Besides increasing the quantity of training data, data augmentation often make the model invariant to the applied noise and enhance its ability to generalize.

Inspired by the success of augmentation methods in ASR [16, 17], as a remedy to avoid overfitting while using low-resource translated speech data, we study the use of spectrogram augmentation (SpecAugment) for direct ST model. SpecAugment [16] is a simple and low-implementation cost approach. Unlike traditional speech augmentation methods that directly manipulate the input signal, SpecAugment is applied on the audio features, which are usually mel spectrogram of the input signal. We utilize two kinds of deformations of the spectrogram which are time and frequency masking, where we mask a block of consecutive time steps and/or mel frequency channels.

Our main motivation of using SpecAugment is the potential avoidance of overfitting, better generalization beyond low-resource training data and improving robustness of the end-to-end models. In this paper, we aim to shed light on the following questions. First, does SpecAugment strategy help the direct ST model? Second, what is the effect of the approach concerning the different amount of training data? Our first contribution is an extensive empirical investigation of SpecAugment on top-performing ST systems to validate or disprove the above conjectures. Our aim is not to compare with other data augmentation strategies, but the effectiveness of the SpecAugment as a stand-alone method. We hope that this method might overcome the data efficiency issue and therefore, as our second contribution, we explore the effect of that on a various amount of training data. Our experimental results on LibriSpeech Audiobooks En→Fr and IWSLT TED-talks En→De show that the method not only greatly outperforms direct ST model up to +1.7% BLEU on average, but also diminishes the overfitting problem. We also

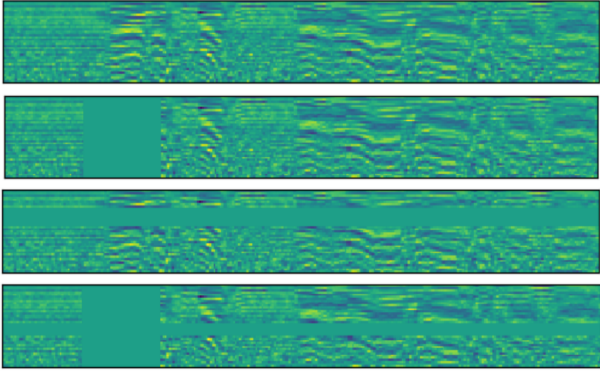


Figure 1: From top to bottom, the figures depict the log mel spectrogram of the input with no augmentation, time masking, frequency masking and both masking applied.

show that our improvements are valid in different data scenarios irrespective of the amount of training data.

2. Spectrogram Augmentation

In spectrogram augmentation (SpecAugment) [16], we randomly apply masking in consecutive frames in the time axis as well as consecutive dimensions in the feature axis. Since the author stated that the time warping is the most expensive and the least influential, we do not explore it here.

2.1. Time Masking

Time masking is masking of τ successive time steps $[t, t + \tau)$, where we set $(x_t, \dots, x_{t+\tau}) := 0$, where τ is the masking window which is selected from a uniform distribution from 0 to the maximum time mask parameter R . (x_1, \dots, x_T) are the input audio features, and T is the length of the input signal. The time position t is picked from another uniform distribution over $[0, T)^1$ such that we never exceed the maximum sequence length T (i.e. if $t + \tau > T$, we set it to T).

We apply the time masking procedure for $m_R \in \mathbb{N}_0$ times. We also ensure that if $m_R > 1$, the same time position t is not selected more than once (i.e. without replacement)².

2.2. Frequency Masking

Frequency masking can be also applied such that ϕ consecutive frequency channels $[f, f + \phi)$ are masked, where ϕ is chosen from a uniform distribution from 0 to the frequency mask parameter F , and f is chosen from $[0, v)^3$. v is the input feature dimension, e.g. the number of mel frequency channels. Similar to time masking, we do not allow for already selected f and check if $f + \phi > v$, we set it to v .

Figure 1 shows examples of the individual augmentations applied to a single input. Multiple frequency and time masks might overlap. $m_F \in \mathbb{N}_0$ refers to the number of times we apply the frequency masks. We note that we standardize the

¹We choose time position differently from the original paper where they select t in the interval of $[0, T - \tau)$ [16].

²It is not clear whether the original paper allows replacement or not.

³Again we note the difference between our implementation and the original paper where the selection interval is $[0, v - \phi)$ [16].




Figure 2: Overview of the direct speech translation model. Shallow grey blocks correspond to pre-trained components, and dark grey blocks are fine-tuned on the ST task.

log mel features to zero mean and variance of one. Therefore, masking to zero is equivalent to setting it to the mean value. In this work, we mainly investigate a series of combinations to find a reliable recipe for direct ST model. We only apply the SpecAugment during training.

3. Network

In ST, given an input observation (audio feature vectors) of variable length T , x_1^T , a sequence of discrete label of unknown length J (source sequence/transcribed words), f_1^J and a sequence of target words of unknown length I , e_1^I , the posterior probability of a target sequence is defined as:

$$p(e_1^I | x_1^T) = \prod_{i=1}^I p(e_i | e_1^{i-1}, x_1^T) \quad (1)$$

where usually $T > I, J$. This posterior can be modeled directly in an end-to-end fashion. Here, we only address the direct end-to-end architecture that is used in our experiments.

The direct model is based on the attention sequence-to-sequence model [8] composed of long short-term memories (LSTMs) [18] similar to [19]. We only focus on LSTM-based models and leave the transformer architecture as our future study [10, 20, 21]. An abstract overview of the network and a summary of the model are shown in Figure 2 and written in Equation 2 respectively. A bidirectional LSTM (BLSTM) scans the input sequence once from left to right and once from right to left. To handle the long speech utterances, we apply max-pooling in the time-dimension at multiple steps inside the speech encoder. For the input sequence x_1^T , we end up with the encoder states $h_1^{T'}$, where $T' = T / red$ with the time reduction factor red . Then, the decoder LSTM generates an output sequence conditioned on the encoder representations. While computing e_i at each time step, an additive attention function is used to produce normalized attention weights $\alpha_{i,t}$. The context vector c_i is then computed as a weighted sum of encoder representations. A transformation followed by a softmax operation predicts e_i . Finally, the decoder state is updated to s_i . Here, L_e and L_d are the number

of encoder and decoder layers respectively. \circ is the concatenation operator of functions.

$$\begin{aligned}
h_1^{T'} &= (\text{BLSTM}_{L_c} \circ \dots \circ \text{max-pool}_1 \circ \text{BLSTM}_1)(x_1^T) \\
\alpha_{i,t} &= \text{softmax}_{\bar{t}}(\text{linear}(\tanh(\text{linear}(s_{i-1}, h_t)))) \\
c_i &= \sum_{t=1}^{T'} \alpha_{i,t} h_t \\
p(e_i | e_1^{i-1}, x_1^T) &= \text{softmax}_{\bar{e}}(\text{linear}(e_{i-1}, s_{i-1}, c_i)) \\
s_i &= \text{LSTM}_{L_d} \circ \dots \circ \text{LSTM}_1(e_i, s_{i-1}, c_i)
\end{aligned} \tag{2}$$

4. Experiments

4.1. Datasets and Metrics

We have conducted our experiments on two ST tasks which are publicly available: the LibriSpeech Audio-books En→Fr [22, 23]⁴ and the IWSLT TED-talks En→De [24, 25]⁵. The training data statistics are listed in Table 1.

Table 1: Training data statistics.

Task	LibriSpeech En→Fr		IWSLT En→De	
	# of seg.	hours	# of seg.	hours
ASR	61.3k	130h	92.9k	207h
ST	94.5k	200h	171.1k	272h
MT	94.5k	-	32M	-

LibriSpeech En→Fr: As suggested by [23], to double the training data size, we concatenate the original translation and the Google Translate reference which have been provided in the dataset package. Hence, we end up to 200h of clean speech corresponding to 94.5k segments for the ST task. We apply 40-dimensional Gammatone features [26] using the RASR feature extractor [27]. For MT training, we utilize no additional data and only use the source-target data from the ST task, i.e. 94.5k. For ASR training, we take both the ASR and ST data resulting in 330h. The dev and test sets contain 2h and 4h of speech, 1071 and 2048 segments respectively. Here, the dev set is used as our cross-validation set as well as checkpoint selection.

IWSLT En→De: Similar to [19, 28], we extract 80-dimensional Mel-frequency cepstral coefficients (MFCC) features. We automatically recompute the provided audio-to-source-sentence alignments to reduce the problem of speech segments without a translation. We use the TED-LIUM corpus including 207h and the IWSLT speech translation TED corpus with 272h of speech data for ASR training. For MT training, we use the TED, and the OpenSubtitles2018 corpora, as well as the data provided by the WMT 2018 evaluation (Europarl, ParaCrawl, CommonCrawl, News Commentary, and Rapid), a total of 65M lines of parallel sentences. We filter these data based on several heuristics resulting in 32M samples. We randomly select a part of our segments as our cross-validation set and choose dev2010 and test2015 as our dev and test sets with 888 and 1080 segments respec-

tively. We select our checkpoints based on dev2010 set.

For both tasks, we remove the punctuation from the transcriptions (i.e. the English text) and keep the punctuation on the target side. After tokenization using Moses toolkit [29]⁶, we apply frequent casing for the IWSLT tasks while lowercase for the LibriSpeech data. Therefore, the evaluation of the IWSLT En→De is case-sensitive, while that of the LibriSpeech is case-insensitive⁷. The translation models are evaluated using the official scripts of WMT campaign, i.e. BLEU [31] computed by `mteval-v13a`⁸ and TER [32] computed by `tercom`⁹. WER is computed by `sclite`¹⁰.

4.2. Models

In our experiments, we build ASR, MT, and ST models all based on the network described in Section 3. The ASR and MT models are used for building the cascade pipeline as well as pre-training the components of the ST model. Thus, the ASR and ST models use the same speech encoder architecture, whilst the MT and ST models use the same text decoder topology, as illustrated in Figure 2.

For both tasks, we apply separate byte pair encoding (BPE) [33] with 20k symbols on both side of the MT data, whereas 10k merge operations on the ASR transcriptions.

ASR model: All tokens are mapped into a 512-dimensional embedding space. The encoder is composed of 6 stacked BLSTM layers with 1024 nodes. The decoder is a 1-layer unidirectional LSTM of size 1024. A single head additive attention with alignment feedback [34, 35] is used as our attention component. Similar to [4, 19], we apply layer-wise pre-training for the encoder, where we start with two encoder layers and a single max-pool in between BLSTM layers. We apply 2 max-pooling layers with pool-size of 3 and 2, i.e. we get a total time reduction factor of 6. We also use CTC auxiliary loss function [36] on top of the speech encoder only during training [37].

MT model: Our MT model follows the ASR model with a 6-layer BLSTM encoder without max-pooling, with a cell size of 1024. The decoder is a 1-layer unidirectional LSTM with cell size 1024, with single head additive attention equipped with alignment feedback.

ST model: The encoder has a similar architecture to the ASR encoder, and the decoder is similar to the MT decoder.

The models are trained end-to-end using the Adam optimizer [38] with a learning rate of 0.0008, and a dropout of 0.3 [39]. We warm-up the learning rate by linearly increasing it for a few training steps. Label smoothing [40] with a ratio of 0.1 is utilized. We employ a learning rate scheduling scheme, where we lower the learning rate with a decay factor of 0.9 if the perplexity on the dev set does not improve for 5 consecutive checkpoints and save the checkpoints every fifth of an epoch. We remove sequences longer than 75 tokens

⁶<http://www.statmt.org/moses/?n=Moses.SupportTools>

⁷We do the case-insensitive evaluation to be comparable with the other works, however, it is not clear which BLEU script they used [23, 30].

⁸<ftp://jaguar.ncsl.nist.gov/mt/resources/mteval-v13a.pl>

⁹<http://www.cs.umd.edu/snoover/tercom/>

¹⁰<http://www1.icsi.berkeley.edu/Speech/docs/sctk-1.2/sclite.htm>

⁴<https://persyval-platform.univ-grenoble-alpes.fr/DS91/detaildataset>

⁵<https://sites.google.com/site/iwslt-evaluation2018/Lectures-task>

Table 2: ASR results measured in WER [%].

Task	WER [\downarrow]	
	dev	test
LibriSpeech En→Fr	6.47	6.47
IWSLT En→De	12.36	13.80

Table 3: MT results measured in BLEU [%] and TER [%] trained using ground truth source text.

Task	BLEU [\uparrow]		TER [\downarrow]	
	dev	test	dev	test
LibriSpeech En→Fr	20.1	18.2	65.3	67.7
IWSLT En→De	30.5	31.5	50.6	-

before batching them together. All batch sizes are specified to be as big as possible to fit in a single GPU memory. A beam size of 12 is used in inference. In order to explore the impact of SpecAugment, we choose to have relatively large models, as explained above. These models include 181M and 192M free parameters for LibriSpeech and IWSLT tasks respectively. The models are built using our in-house NN-toolkit software that relies on TensorFlow [41]. The code is open source and the configurations of the setups are available online.

5. Results

Table 2 and 3 present the results for the ASR and MT models (described in Section 4.2) on LibriSpeech and IWSLT tasks respectively. On the test sets, we achieve 6.47% and 13.80% WER. We note that one might gain better WER using the conventional hybrid hidden Markov model (HMM) - neural network (NN) approach on phoneme level [42], which is out of the scope of this paper.

The MT task on LibriSpeech seems more challenging as both scores are lower. We obtain 18.2% BLEU on the LibriSpeech and respectively 31.5% BLEU on the IWSLT by pure MT. Table 4 shows the traditional cascade pipeline where the output of our ASR model, a sequence of tokens is fed as the input to our MT system. We gain 15.7% BLEU and 70.6% TER on the LibriSpeech and 24.4% BLEU and 62.5% TER on IWSLT test set. As expected, the ST systems are behind the pure MT models (cf. Table 3 and Table 4). In the rest of the paper, we only focus on the results of direct models.

Table 4: ST results using cascaded pipeline of independent ASR and MT models measured in BLEU [%] and TER [%]. We highlight that the cascaded pipeline used more training data compared to the direct model.

Task	BLEU [\uparrow]		TER [\downarrow]	
	dev	test	dev	test
LibriSpeech En→Fr	17.3	15.7	69.1	70.6
IWSLT En→De	24.7	24.4	58.9	62.5

Table 5: SpecAugment results on LibriSpeech En→Fr using various augmentation parameters. “-”: failed training. The bold augmentation parameters are used in the rest of our experiments.

F	m_F	R	m_R	BLEU [%]		TER [%]	
				dev	test	dev	test
0	0	0	0	15.8	15.2	74.1	75.8
2	1	0	0	17.3	16.1	72.5	75.0
4	1	0	0	18.0	16.1	72.2	74.9
4	2	0	0	17.6	16.1	72.3	74.6
4	4	0	0	17.1	16.1	73.2	74.6
5	1	0	0	17.6	16.1	72.4	74.7
8	1	0	0	17.4	15.7	72.8	75.2
8	2	0	0	17.5	15.7	73.0	75.9
20	1	0	0	17.2	15.7	72.2	75.0
20	2	0	0	16.7	15.4	73.0	75.9
35	5	0	0	16.1	15.0	74.1	76.8
0	0	20	1	17.2	15.7	73.3	75.5
0	0	20	2	17.3	16.2	72.9	75.2
0	0	40	1	17.4	15.6	73.2	75.8
0	0	40	2	17.3	16.2	72.5	74.7
0	0	40	4	16.6	15.6	73.1	75.9
0	0	80	1	17.5	16.1	72.5	75.1
0	0	100	1	17.3	15.9	73.8	75.7
0	0	100	2	16.8	15.4	74.1	76.6
0	0	200	2	16.6	15.3	73.8	76.0
0	0	400	5	-	-	-	-
5	1	40	2	17.9	16.1	72.2	74.6
5	1	80	1	17.6	15.6	72.6	75.5
4	1	40	2	17.8	16.1	72.6	75.1
4	2	40	2	16.8	15.7	72.3	74.6

5.1. Using SpecAugment

In this section, we explore different types of masking. In the first set of experiments, we deactivate either time or frequency masking. The results are listed in Table 5 and 6 for LibriSpeech and IWSLT tasks respectively. As it is shown in Table 5, we apply various types of frequency masking with different values of F , between 2 and 35, and m_F while the time masking is disabled. As listed, the optimum value of F is around 4 and 5 with 1.5% in BLEU and 1.4% in TER on average of dev and test sets. A further increase of F until 20 hurts the performance by 0.4% in both BLEU and TER on the test set. To verify the effect of SpecAugment, we also employ a coarse policy where we randomly mask 35 frequency channels out of 40 and apply it 5 times. As expected, the performance drops behind the direct baseline. Interestingly enough, even a small value of frequency masking ($F = 2$) leads to improvements. It is important to highlight that since we randomly select the masking window between zero and maximum value of F , the results are close to each other. The best results of each set of experiments are highlighted in bold.

We also disable the frequency masking by setting F and m_F set to zero and vary the time mask parameter R and the

Table 6: SpecAugment results on IWSLT En→De using various augmentation parameters.

F	m_F	R	m_R	BLEU [%]		TER [%]	
				dev	test	dev	test
0	0	0	0	16.9	16.5	67.3	70.6
4	1	0	0	17.8	17.0	66.3	69.5
4	2	0	0	17.3	17.4	66.5	69.3
5	1	0	0	18.1	17.5	66.3	68.8
8	1	0	0	17.6	17.1	66.2	-
10	1	0	0	16.9	16.9	66.9	70.2
20	1	0	0	17.0	17.1	-	70.3
0	0	20	1	17.7	17.6	-	70.1
0	0	20	2	17.5	17.5	65.8	69.2
0	0	40	1	17.5	17.1	66.1	70.0
0	0	40	2	17.9	17.5	65.8	69.5
0	0	60	1	17.6	17.3	66.6	-
0	0	80	1	16.9	17.1	68.6	69.9
0	0	80	2	16.8	16.6	68.0	72.0
0	0	100	1	17.7	17.0	66.2	-
5	1	40	2	17.8	16.8	66.2	70.9
4	1	40	1	17.5	17.2	66.5	69.4
4	1	40	2	17.7	18.0	66.0	69.2
27	2	100	2	16.6	16.7	67.6	70.4

number of times it has been called. Again, there is a limit of how much data augmentation can be applied. Enlarging the time masking window R to 100 leads to lower BLEU and TER scores. Furthermore, we drastically increase the time masking window R into 400 steps and apply it 5 times which fails due to unstable training. If the initial convergence is stable, in all cases, adding some data augmentation improves the setup, however, at some point, the performance degrades by more augmentation.

Moreover, we observe that in many cases applying small window several times gives slightly more improvements compared to the policy in which a large window applied once (cf. row 4 and 7 in Table 5). After finding the optimum of both time and frequency masking, we have done some combinations of both masking as shown in the table. As expected, the combination of best of two masking gives the largest boost.

We verify the influence of SpecAugment on IWSLT En→De task by an improvement up to 1.2% in BLEU and 1.8% in TER (see Table 6). For IWSLT, we also apply a policy similar to the LD policy of main paper [16] as listed in the last row of Table 6. As shown, this setup is not the optimum case for our task, which leads to the conclusion that SpecAugment might be working better by fine-tuning on a specific task, however, adding some data augmentation improves the setup. Moreover, based on the above results, SpecAugment performs quite well regardless of the features and their dimensions. In our experiments, we use 40-dimensional Gammatone features for LibriSpeech respectively 80-dimensional MFCC features for IWSLT. In both cases, augmentation helps the performance. In the rest of

Table 7: SpecAugment results with with a varying amount of training data on LibriSpeech En→Fr. Δ indicates the absolute difference on average of dev and test sets.

# segments	BLEU [%]			TER [%]		
	dev	test	Δ	dev	test	Δ
23k	11.6	7.0	1.0	90.2	92.3	-1.2
+ SpecAug	13.3	7.3		91.5	93.4	
47k	14.8	10.2	1.8	87.2	89.1	2.8
+ SpecAug	17.3	11.2		83.4	87.4	
71k	15.4	13.2	1.1	78.9	81.0	2.1
+ SpecAug	16.7	14.4		76.7	79.1	
94k	15.8	15.2	1.5	74.1	75.8	1.6
+ SpecAug	17.9	16.1		72.2	74.6	

our experiments, we use the augmentation parameters which is bold in the tables. For IWSLT, we choose $F = 4$ rather $F = 5$, but the rest of the parameters are the same.

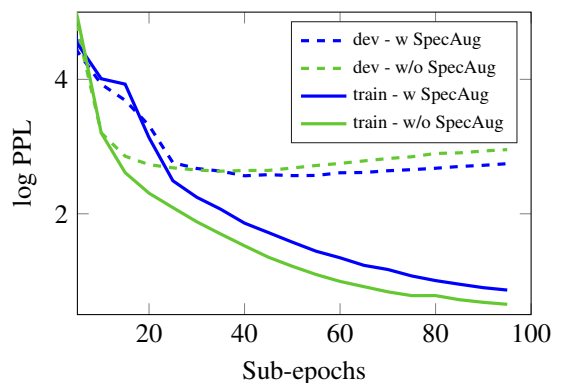


Figure 3: Average log perplexity of training and dev sets across epochs.

5.2. Importance of SpecAugment on overfitting

We also study the effect of SpecAugment on overfitting. Figure 3 shows log-perplexity plots on training and dev set with and without augmentation. It is seen that SpecAugment leads to better generalization, as measured from the difference between the perplexity of training and dev data. The model trained with SpecAugment still has a training data likelihood which is higher than the baseline system. Therefore, we can confirm that the method reduces overfitting up to some degree. In this case, we need to train the SpecAugment system for few more epochs. A corresponding increase in the number of epochs for the baseline system deteriorated the performance.

5.3. Effect on training data size

We go further and show that SpecAugment can be leveraged to improve the performance of a direct ST model, when the amount of training data is limited. To do so, we have conducted studies on a different portion of training data to see how the method performs in different data conditions. Table

Table 8: SpecAugment results with pre-training, which makes use of more training data.

method	LibriSpeech En→Fr				IWSLT En→De			
	BLEU [%]		TER [%]		BLEU [%]		TER [%]	
	dev	test	dev	test	dev	test	dev	test
direct	15.8	15.2	74.1	75.8	16.9	16.5	67.3	70.6
+ pretraining	18.0	15.8	71.3	73.9	21.1	20.7	62.1	65.5
+ SpecAugment	18.5	16.2	71.0	74.5	21.3	20.9	61.9	65.7

7 compares the performance improvement of SpecAugment with a varying amount of training data. It can be seen that it is helpful in all scenarios irrespective of the amount of training data.

The minimal gains, even hurting TER when 23k samples are used. It could be attributed to the fact that 23k segments are not sufficient to train a reliable deep model. We believe that as long as a moderate amount of data is available, SpecAugment helps data efficiency more. The augmentation policy compensates lack of data when we half the training data size to 47k segments. It achieves 1.8% of absolute improvement difference in BLEU and 2.8% in TER compared to the model trained using 94k samples with 1.5% of BLEU and 1.6% of TER.

5.4. Pre-training

We also consider the effect of data augmentation on the top of pre-training. We pre-train the encoder using our pre-trained ASR model, and the decoder using our MT model as described in Section 4.2 which use more training data compared to the direct ST model (see Table 1). After initialization with pre-trained components, the ST model is fine-tuned using the ST data. Here, we add an additional BLSTM layer (adaptor layer) to adapt the output of speech encoder and the input of text decoder without freezing the parameters (see Fig. 2).

As shown in Table 8, SpecAugment slightly outperforms the pre-training. It outperforms the pre-trained models by 0.5% and 0.2% BLEU on LibriSpeech and IWSLT respectively, and no TER improvements. This results confirm that SpecAugment can be used along with pre-training. By comparing Table 5 and 8, one can argue that the SpecAugment might compensate the effect of pre-training strategy by its own for LibriSpeech (compare 18.0% vs. 17.9% in BLEU on the dev set and 15.8% vs. 16.1% on the test set).

We finally compare our model with the other works in the literature in Table 9. On the LibriSpeech test set, our model outperforms both the LSTM-based end-to-end models and the Transformer-based. Contrary to [23] in which character decoder is used, we apply BPE that obtain improvements. Both our direct model and the cascade model outperform the models in [23]. We also beat the Transformer models without augmentation. Our recipe works as good as knowledge distillation method where an MT model is exploited to teach the ST model.

Table 9: Comparison on LibriSpeech En→Fr test set with the literature. In order to be comparable with other works, the results in this table are case-insensitive BLEU computed using `multi-bleu.pl` script [29]. ¹: the evaluation is without punctuation. ²: it correspond to 16.2% BLEU from Table 8.

Method	BLEU [%]
other works	
direct [23]	13.3
multi-task [23]	13.4
cascade pipeline [23]	14.6
unsupervised ¹ [30]	12.2
Transformer [20]	13.8
Transformer+pretraining [21]	14.3
+ knowledge distillation [21]	17.0
this work	
direct+pretraining+SpecAugment ²	17.0

6. Conclusion

We have studied SpecAugment, a simple and low-cost data augmentation for end-to-end direct speech translation. There is a limit of how much data augmentation can be applied. Adding some data augmentation improves the performance in terms of BLEU and TER, however, at some point, the performance degrades by more augmentation. We have also shown that the method avoids overfitting to some extent and it requires longer training. A common criticism of many techniques for low-resource applications is that the improvements go away once we have lots of synthetic parallel data. Therefore, we believe that comparing the SpecAugment approach with generated data using TTS model is a crucial next step. We also aim to explore the effectiveness of the approach on Transformer architecture.

7. Acknowledgements



This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 694537, project “SEQ-CLAS”), the Deutsche Forschungsgemeinschaft (DFG; grant agreement NE 572/8-1, project “CoreTec”) and from a Google Focused Award. The work reflects only the authors’ views and none of the funding parties is responsible for any use that may be made of the information it contains.

8. References

- [1] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, “End-to-end attention-based large vocabulary speech recognition,” in *IEEE Inter. Conf. ICASSP, Shanghai, China, Mar. 20-25, 2016*, pp. 4945–4949.
- [2] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *IEEE Inter. Conf. ICASSP, Shanghai, China, Mar. 20-25, 2016*, pp. 4960–4964.
- [3] J. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, “Attention-based models for speech recognition,” in *Annual Conf. NIPS, Dec. 7-12, Montreal, Quebec, Canada, 2015*, pp. 577–585.
- [4] A. Zeyer, K. Irie, R. Schlüter, and H. Ney, “Improved training of end-to-end attention models for speech recognition,” in *Interspeech 2018, 19th Annual Conference of the International Speech Communication Association, Hyderabad, India, 2-6 September 2018*, pp. 7–11.
- [5] P. Bahar, A. Zeyer, R. Schlüter, and H. Ney, “On using 2d sequence-to-sequence models for speech recognition,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Brighton, UK, May 2019.
- [6] A. Zeyer, P. Bahar, K. Irie, R. Schlüter, and H. Ney, “A comparison of transformer and lstm encoder decoder models for asr,” in *IEEE Automatic Speech Recognition and Understanding Workshop*, Sentosa, Singapore, Dec. 2019.
- [7] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to sequence learning with neural networks,” in *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13, Montreal, Quebec, Canada, 2014*, pp. 3104–3112.
- [8] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *CoRR*, vol. abs/1409.0473, 2015.
- [9] T. Luong, H. Pham, and C. D. Manning, “Effective approaches to attention-based neural machine translation,” in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pp. 1412–1421.
- [10] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December, Long Beach, CA, USA, 2017*, pp. 6000–6010.
- [11] P. Bahar, C. Brix, and H. Ney, “Towards two-dimensional sequence to sequence model in neural machine translation,” in *Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium, Nov. 2018, pp. 3009–3015.
- [12] R. J. Weiss, J. Chorowski, N. Jaitly, Y. Wu, and Z. Chen, “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*, pp. 2625–2629.
- [13] A. Berard, O. Pietquin, C. Servan, and L. Besacier, “Listen and translate: A proof of concept for end-to-end speech-to-text translation,” *CoRR*, vol. abs/1612.01744, 2016.
- [14] R. Sennrich, B. Haddow, and A. Birch, “Improving neural machine translation models with monolingual data,” in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, Berlin, Germany, 2016*.
- [15] Y. Jia, M. Johnson, W. Macherey, R. J. Weiss, Y. Cao, C. Chiu, N. Ari, S. Laurenzo, and Y. Wu, “Leveraging weakly supervised data to improve end-to-end speech-to-text translation,” *CoRR*, vol. abs/1811.02050, 2018.
- [16] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” *arXiv preprint arXiv:1904.08779*, 2019.
- [17] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, “Audio augmentation for speech recognition,” in *INTER-SPEECH 2015, 16th Annual Conference of the International Speech Communication Association, Dresden, Germany, September 6-10, 2015*, pp. 3586–3589.
- [18] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [19] P. Bahar, T. Bieschke, and H. Ney, “A comparative study on end-to-end speech to text translation,” in *IEEE Automatic Speech Recognition and Understanding Workshop*, Sentosa, Singapore, Dec. 2019.
- [20] M. A. D. Gangi, M. Negri, R. Cattoni, R. Dessi, and M. Turchi, “Enhancing transformer for end-to-end speech-to-text translation,” in *Proceedings of Machine Translation Summit XVII Volume 1: Research Track*. Dublin, Ireland: European Association for Machine Translation, 19–23 Aug. 2019, pp. 21–31.
- [21] Y. Liu, H. Xiong, Z. He, J. Zhang, H. Wu, H. Wang, and C. Zong, “End-to-end speech translation with knowledge distillation,” *CoRR*, vol. abs/1904.08075, 2019.
- [22] A. C. Kocabiyikoglu, L. Besacier, and O. Kraif, “Augmenting librispeech with french translations: A multi-modal 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*.
- [23] A. Berard, L. Besacier, A. C. Kocabiyikoglu, and O. Pietquin, “End-to-end automatic speech translation of audiobooks,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2018, Calgary, AB, Canada, April 15-20, 2018*, pp. 6224–6228.
- [24] E. Cho, S. Fünfer, S. Stüker, and A. Waibel, “A cor-

- pus of spontaneous speech in lectures: The KIT lecture corpus for spoken language processing and translation,” in *Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014, Reykjavik, Iceland, May 26-31, 2014*, pp. 1554–1559.
- [25] S. S. M. C. M. T. M. F. J. Niehues, R. Cattoni, “The iwslt 2018 evaluation campaign,” in *15th International Workshop on Spoken Language Translation (IWSLT 2018)*, Bruges, Belgium, October 29 - 30, 2018.
- [26] R. Schlüter, I. Bezrukov, H. Wagner, and H. Ney, “Gammatone features and feature combination for large vocabulary speech recognition,” in *IEEE Inter. Conf. ICASSP, Honolulu, Hawaii, USA, Apr. 15-20, 2007*, pp. 649–652.
- [27] S. Wiesler, A. Richard, P. Golik, R. Schlüter, and H. Ney, “RASR/NN: the RWTH neural network toolkit for speech recognition,” in *IEEE Inter. Conf. ICASSP, Florence, Italy, May 4-9, 2014*, pp. 3281–3285.
- [28] E. Matusov, P. Wilken, P. Bahar, J. Schamper, P. Golik, A. Zeyer, J. A. Silvestre-Cerda, A. Martinez-Villaronga, H. Pesch, and J.-T. Peter, “Neural speech translation at apptek,” in *Proceedings of the 15th International Workshop on Spoken Language Translation*, Oct. 2018, pp. 104–111.
- [29] P. Koehn, H. Hoang, A. Birch, C. Callison-Burch, M. Federico, N. Bertoldi, B. Cowan, W. Shen, C. Moran, R. Zens, *et al.*, “Moses: Open source toolkit for statistical machine translation,” in *Proceedings of the 45th annual meeting of the ACL on interactive poster and demonstration sessions*. Association for Computational Linguistics, 2007, pp. 177–180.
- [30] Y. Chung, W. Weng, S. Tong, and J. Glass, “Towards unsupervised speech-to-text translation,” *CoRR*, vol. abs/1811.01307, 2018.
- [31] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “Bleu: a Method for Automatic Evaluation of Machine Translation,” in *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, Philadelphia, Pennsylvania, USA, July 2002, pp. 311–318.
- [32] M. Snover, B. Dorr, R. Schwartz, L. Micciulla, and J. Makhoul, “A Study of Translation Edit Rate with Targeted Human Annotation,” in *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas*, Cambridge, Massachusetts, USA, August 2006, pp. 223–231.
- [33] R. Sennrich, B. Haddow, and A. Birch, “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, Berlin, Germany, 2016*.
- [34] Z. Tu, Z. Lu, Y. Liu, X. Liu, and H. Li, “Modeling coverage for neural machine translation,” in *Proceedings of the Association for Computational Linguistics (ACL)*, 2016.
- [35] P. Bahar, J. Rosendahl, N. Rossenbach, and H. Ney, “The rwth aachen machine translation systems for iwslt 2017,” in *Int. Workshop on Spoken Language Translation*, 2017, pp. 29–34.
- [36] A. Graves, S. Fernández, F. J. Gomez, and J. Schmidhuber, “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, pp. 369–376.
- [37] T. Hori, S. Watanabe, Y. Zhang, and W. Chan, “Advances in joint CTC-attention based end-to-end speech recognition with a deep CNN encoder and RNN-LM,” in *Interspeech*, 2017.
- [38] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2014.
- [39] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: a simple way to prevent neural networks from overfitting,” *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [40] G. Pereyra, G. Tucker, J. Chorowski, L. Kaiser, and G. E. Hinton, “Regularizing neural networks by penalizing confident output distributions,” *CoRR*, vol. abs/1701.06548, 2017.
- [41] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, and *et. al.*, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015, software available from tensorflow.org.
- [42] H. Bourlard and N. Morgan, *Connectionist speech recognition: a hybrid approach*. Springer, 1994, vol. 247.