

End-to-End Speech-Translation with Knowledge Distillation: FBK@IWSLT2020

Marco Gaido^{1,2}, Mattia Antonino Di Gangi^{1,2}, Matteo Negri¹, Marco Turchi¹

¹Fondazione Bruno Kessler, Trento, Italy

²University of Trento, Italy

{mgaido|digangi|negri|turchi}@fbk.eu

Abstract

This paper describes FBK's participation in the IWSLT 2020 offline speech translation (ST) task. The task evaluates systems' ability to translate English TED talks audio into German texts. The test talks are provided in two versions: one contains the data already segmented with automatic tools and the other is the raw data without any segmentation. Participants can decide whether to work on custom segmentation or not. We used the provided segmentation. Our system is an end-to-end model based on an adaptation of the Transformer for speech data. Its training process is the main focus of this paper and it is based on: *i*) transfer learning (ASR pretraining and knowledge distillation), *ii*) data augmentation (SpecAugment, *time stretch* and synthetic data), *iii*) combining synthetic and real data marked as different domains, and *iv*) multi-task learning using the CTC loss. Finally, after the training with word-level knowledge distillation is complete, our ST models are fine-tuned using label smoothed cross entropy. Our best model scored 29 BLEU on the MuST-C En-De test set, which is an excellent result compared to recent papers, and 23.7 BLEU on the same data segmented with VAD, showing the need for researching solutions addressing this specific data condition.

1 Introduction

The offline speech translation task consists in generating the text translation of speech audio recordings into a different language. In particular, the IWSLT2020 task (Ansari et al., 2020) evaluates German translation of English recordings extracted from TED talks. The test dataset is provided to participants both segmented in a sentence-like format using a Voice Activity Detector (VAD) and in the original unsegmented form. Although custom segmentation of the data can provide drastic

improvements in the final scores, in our work we have not addressed it, participating only with the provided segmentation.

Two main approaches are possible to face the speech translation task. The classic one is the cascade solution, which includes automatic speech recognition (ASR) and machine translation (MT) components. The other option is an end-to-end (E2E) solution, which performs ST with a single sequence-to-sequence model. Both of them are allowed for the IWSLT2020 task, but our submission is based on an E2E model.

E2E ST models gained popularity in the last few years. Their rise is due to the lack of error propagation and the reduced latency in generating the output compared to the traditional cascaded approach. Despite these appealing properties, they failed so far to reach the same results obtained by cascade systems, as shown also by last year's IWSLT campaign (Niehues et al., 2019). One reason for this is the limited amount of parallel corpora compared to those used to separately train ASR and MT components. Moreover, training an E2E ST system is more difficult because the task is more complex, since it deals with understanding the content of the input audio and translating it into a different language directly and without recurring to intermediate representations.

The above-mentioned observations have led researchers to focus on transferring knowledge from MT and ASR systems to improve the ST models. A traditional approach consists in pretraining components: the ST encoder is initialized with the ASR encoder and the ST decoder with the MT decoder. The encoder pretraining has indeed proved to be effective (Bansal et al., 2019), while the decoder pretraining has not demonstrated to be as effective, unless with the addition of adaptation layers (Bahar et al., 2019a). A more promising way to transfer knowledge from an MT model is to use the MT as

a teacher to distill knowledge for the ST training (Liu et al., 2019). This is the approach we explore in the paper.

Despite its demonstrated effectiveness, ASR pretraining has been replaced in some works by multitask learning (Weiss et al., 2017). In this case, the model is jointly trained with two (or more) loss functions and usually the model is composed of 3 components: *i*) a shared encoder, *ii*) a decoder which generates the transcription, and *iii*) a decoder which generates the translation. We adopt the slightly different approach introduced by (Bahar et al., 2019a), which does not introduce an additional decoder but relies on the CTC loss in order to predict the transcriptions (Kim et al., 2017). As this multi-task learning has been proposed for speech recognition and has demonstrated to be useful in that scenario, we also include the CTC loss in ASR pretraining.

Another topic that received considerable attention is data augmentation. Many techniques have been proposed: in this work we focus on SpecAugment (Park et al., 2019), *time stretch* and subsequence sampling (Nguyen et al., 2020). Moreover, we used synthetic data generated by automatically translating the ASR datasets with our MT model. This process can also be considered as a sequence-level knowledge distillation technique, named Sequence KD (Kim and Rush, 2016).

In this paper, we explore different ways to combine synthetic and real data. We also check if the benefits of the techniques mentioned above are orthogonal and joining them leads to better results.

Our experiments show that:

- knowledge distillation, ASR pretraining, multi-task learning and data augmentation are complementary, i.e. they cooperate to produce a better model;
- combining synthetic and real data marking them with different tags (Caswell et al., 2019) leads to a model which generalizes better;
- fine-tuning a model trained with word-level knowledge distillation using the more classical label smoothed cross entropy (Szegedy et al., 2016) significantly improves the results;
- there is a huge performance gap between data segmented in sentences and data segmented with VAD. Indeed, on the same test set, the score on VAD-segmented data is lower by 5.5 BLEU.

To summarize, our submission is characterized by tagged synthetic data, multi-task with CTC loss on the transcriptions, data augmentation and word-level knowledge distillation.

2 Training data

This section describes the data used to build our models. They include: *i*) MT corpora (English-German sentence pairs), for the model used in knowledge distillation; *ii*) ASR corpora (audio and English transcriptions), for generating a pretrained encoder for the ST task; *iii*) ST corpora (audios with corresponding English transcription and German translation), for the training of our ST models. For each task, we used all the relevant datasets allowed by the evaluation campaign¹.

MT. All datasets allowed in WMT 2019 (Barrault et al., 2019) were used for the MT training, with the addition of OpenSubtitles2018 (Lison and Tiedemann, 2016). These datasets contain spurious sentence pairs: some target sentences are in a language different from German (often in English) or are unrelated to the corresponding English source or contain unexpected characters (such as ideograms). As a consequence, an initial training on them caused the model to produce some English sentences, instead of German, in the output. Hence, we cleaned our MT training data using Modern MT (Bertoldi et al., 2017)², in order to remove sentences whose language is not the correct one. We further filtered out sentences containing ideograms with a custom script. Overall, we removed roughly 25% of the data and the final dataset used in the training contains nearly 49 million sentence pairs.

ASR. For this task, we used both pure ASR and ST available corpora. They include TED-LIUM 3 (Hernandez et al., 2018), Librispeech (Panayotov et al., 2015), Mozilla Common Voice³, How2 (Sanabria et al., 2018), the En-De section of MuST-C (Di Gangi et al., 2019a), the Speech-Translation TED corpus provided by the task organizers¹ and the En-De section of Europarl-ST (Iranzo-Sánchez et al., 2020). All data was lowercased and punctuation was removed.

ST. In addition to the allowed ST corpora (MuST-C, Europarl-ST and the Speech-Translation TED

¹http://iwslt.org/doku.php?id=offline_speech_translation

²We run the CleaningPipelineMain class of MMT.

³<https://voice.mozilla.org/>

corpus), we generated synthetic data using Sequence KD (see Section 3.2) for all the ASR datasets missing the German reference. Moreover, we generated synthetic data for the En-Fr section of MuST-C. Overall, the combination of real and generated data resulted in a ST training set of 1.5 million samples.

All texts were preprocessed by tokenizing them, de-escaping special characters and normalizing punctuation with the scripts in the Moses toolkit (Koehn et al., 2007). The words in both languages were segmented using BPE with 8,000 merge rules learned jointly on the two languages of the MT training data (Sennrich et al., 2016). The audio was converted into 40 log Mel-filter banks with speaker normalization using XNMT (Neubig et al., 2018). We discarded samples with more than 2,000 filter-banks in order to prevent memory issues.

3 Models and training

3.1 Architectures

The models we trained are based on Transformer (Vaswani et al., 2017). The MT model is a plain Transformer with 6 layers for both the encoder and the decoder, 16 attention heads, 1,024 features for the attention layers and 4,096 hidden units in feed-forward layers.

2D Self-Attention	Encoder	Decoder	BLEU
2	6	6	16.50
0	8	6	16.90
2	9	6	17.08
2	9	4	17.06
2	12	4	17.31

Table 1: Results on Librispeech with Word KD varying the number of layers.

The ASR and ST models are a revisited version of the S-Transformer introduced by (Di Gangi et al., 2019c). In preliminary experiments on Librispeech (see Table 1), we observed that replacing 2D self-attention layers with additional Transformer encoder layers was beneficial to the final score. Moreover, we noticed that adding more layers in the encoder improves the results, while removing few layers of the decoder does not harm performance. Hence, the models used in this work process the input with two 2D CNNs, whose output is projected into the higher-dimensional space used by the Transformer encoder layers. The projected output is summed with positional embeddings before being fed to the Transformer encoder layers, which

use logarithmic distance penalty.

Both our ASR and ST models have 8 attention heads, 512 features for the attention layers and 2,048 hidden units in FFN layers. The ASR model has 8 encoder layers and 6 decoder layers, while the ST model has 11 encoder layers and 4 decoder layers. The ST encoder is initialized with the ASR encoder (except for the additional 3 layers that are initialized with random values). The decision of having a different number of encoder layers in the two encoders is motivated by the idea of introducing adaptation layers, which (Bahar et al., 2019a) reported to be essential when initializing the decoder with that of a pretrained MT model.

3.2 Data augmentation

One of the main problems for end-to-end ST is the scarcity of parallel corpora. In order to mitigate this issue, we explored the following data augmentation strategies in our participation.

SpecAugment. SpecAugment is a data augmentation technique originally introduced for ASR, whose effectiveness has also been demonstrated for ST (Bahar et al., 2019b). It operates on the input filterbanks and it consists in masking consecutive portions of the input both in the frequency and in the time dimensions. On every input, at each iteration, SpecAugment is applied with probability p . In case of application, it generates *frequency masking num* masks on the frequency axis and *time masking num* masks on the time axis. Each mask has a starting index, which is sampled from a uniform distribution, and a number of consecutive items to mask, which is a random number between 0 and respectively *frequency masking pars* and *time masking pars*. Masked items are set to 0. In our work, we always applied SpecAugment to both the ASR pretraining and the ST training. The configuration we used are: $p = 0.5$, *frequency masking pars* = 13, *time masking pars* = 20, *frequency masking num* = 2 and *time masking num* = 2.

Time stretch. *Time stretch* (Nguyen et al., 2020) is another technique which operates directly on the filterbanks, aiming at generating the same effect of speed perturbation. It divides the input sequence in windows of w features and re-samples each of them by a random factor s drawn by a uniform distribution between 0.8 and 1.25 (in our implementation, the lower bound is set to 1.0 in case of an input sequence with length lower than 10). In this work, we perturb an input sample using *time stretch* with

probability 0.3.

Sub-sequence sampling. As mentioned in the introduction, there is a huge gap in model’s performance when translating data split in well-formed sentences and data split with VAD. In order to reduce this difference, we tried to train the model on sentences which are not always well-formed by using sub-sequence sampling (Nguyen et al., 2020). Sub-sequence sampling requires the alignments between the speech and the target text at word level. As this information is not possible to obtain for the translations, we created the sub-sequences with the alignments between the audio and the transcription, and then we translated the obtained transcription with our MT model to get the target German translation. For every input sentence, we generated three segments: *i*) one starting at the beginning of the sentence and ending at a random word in the second half of the sentence, *ii*) one starting at a random word in the first half of the sentence and ending at the end of the sentence, and *iii*) one starting at a random word in the first quarter of the sentence and ending at a random word in the last quarter of the sentence.

In our experiments, this technique has not provided significant improvements (the gain was less than 0.1 BLEU on the VAD-segmented test set). Hence, it was not included in our final models.

Synthetic data. Finally, we generated synthetic translations for the data in the ASR datasets to create parallel audio-translation pairs to be included in the ST trainings. The missing target sentences were produced by translating the transcript of each audio sample with our MT model, as in (Jia et al., 2019). If the transcription of a dataset was provided with punctuation and correct casing, this was fed to the MT model; otherwise, we had to use the lowercase transcription without punctuation.

Top K	BLEU
4	16.43
8	16.50
64	16.37
1024	16.34

Table 2: Results on Librispeech with different K values, where K is the number of tokens considered for Word KD.

3.3 Knowledge distillation

While the ASR and MT models are optimized on label smoothed cross entropy with smoothing factor 0.1, our ST models are trained with word-level

knowledge distillation (Word KD). In Word KD, the model being trained is named *student* and the goal is to teach it to produce the same output distribution of another - pretrained - model, named *teacher*. This is obtained by computing the KL divergence (Kullback and Leibler, 1951) between the distribution produced by the student and the distribution produced by the teacher. The rationale of knowledge distillation resides in providing additional information to the student, as the output probabilities produced by the teacher reflect its hidden knowledge (the so-called *dark knowledge*), and in the fact that the soft labels produced by the teacher are an easier target to match for the student than cross entropy.

In this work, we follow (Liu et al., 2019), so the teacher model is our MT model and the student is the ST model. Compared to (Liu et al., 2019), we make the training more efficient by extracting only the top 8 tokens from the teacher distribution. In this way, we can precompute and store the MT output instead of computing it at each training iteration, since its size is reduced by three orders of magnitude. Moreover, this approach does not affect negatively the final score, as shown by (Tan et al., 2019) and confirmed for ST by our experiments in Table 2).

Moreover, once the training with Word KD is terminated, we perform a fine-tuning of the ST model using the label smoothed cross entropy. Fine-tuning on a different target is an approach whose effectiveness has been shown by (Kim and Rush, 2016). Nevertheless, they applied a fine-tuning on knowledge distillation after a pretraining with the cross entropy loss, while here we do the opposite. Preliminary experiments on Librispeech showed that there is no difference in the order of the trainings (16.79 vs 16.81 BLEU, compared to 16.5 BLEU before the fine-tuning). In the fine-tuning, we train both on real and synthetic data, but we do not use the other data augmentation techniques.

3.4 Training scheme

A key aspect is the training scheme used to combine the real and synthetic datasets. In this paper, we explore two alternatives:

- **Sequence KD + Finetune:** this is the training scheme suggested in (He et al., 2020). The model is first trained with Sequence KD and Word KD on the synthetic datasets and then it

is fine-tuned on the datasets with ground-truth targets using Word KD.

- **Multi-domain:** similarly to our last year submission (Di Gangi et al., 2019b), the training is executed on all data at once, but we introduce three *tokens* representing the three types of data, namely: *i*) those whose ground-truth translations are provided, *ii*) those generated from true case transcriptions with punctuation, *iii*) those generated from lowercase transcriptions without punctuation. We explore the two most promising approaches according to (Di Gangi et al., 2019d) to integrate the token with the data, i.e. summing the token to all input data and summing the token to all decoder input embeddings.

3.5 Multi-task training

We found that adding the CTC loss (Graves et al., 2006) to the training objective gives better results both in ASR and ST, although it slows down the training by nearly a factor of 2. During the ASR training, we added the CTC loss on the output of the last layer of the encoder. During the ST training, instead, the CTC loss was computed using the output of the last layer pretrained with the ASR encoder, i.e. the 8th layer. In this way, the ST encoder has three additional layers which can transform the representation into features which are more convenient for the ST task, as Bahar et al. (2019a) did introducing an adaptation layer.

4 Experimental settings

For our experiments, we used the described training sets and we picked the best model according to the perplexity on MuST-C En-De validation set. We evaluated our models on three benchmarks: *i*) the MuST-C En-De test set segmented at sentence level, *ii*) the same test set segmented with a VAD (Meignier and Merlin, 2010), and *iii*) the IWSLT 2015 test set (Cettolo et al., 2015).

We trained with Adam (Kingma and Ba, 2015) (betas (0.9, 0.98)). Unless stated otherwise, the learning rate was set to increase linearly from 3e-4 to 5e-4 in the first 5,000 steps and then decay with an inverse square root policy. For fine-tuning, the learning rate was kept fixed at 1e-4. A 0.1 dropout was applied.

Each GPU processed mini-batches containing up to 12K tokens or 8 samples and updates were performed every 8 mini-batches. As we had 8 GPUs,

the actual batch size was about 512. In the case of multi-domain training, a batch for each domain was processed before an update: since we have three domains, the overall batch size was about 1,536. Moreover, the datasets in the different domains had different sizes, so the smaller ones were oversampled to match the size of the largest.

As the truncation of the output values of the teacher model to the top 8 leads to a more peaked distribution, we checked if contrasting this bias is beneficial or not. Hence, we tuned the value of the temperature at generation time in the interval 0.8-1.5. The temperature T is a parameter which is used to divide the *logits* before the *softmax* and determines whether to output a softer (if $T > 1$) or a sharper (if $T < 1$) distribution (Hinton et al., 2015). By default T is 1, returning an unmodified distribution. The generation of the results reported in this paper was performed using $T = 1.3$ for the models trained on Word KD. This usually provided a 0.1-0.5 BLEU increase on our benchmarks compared to $T = 1$, confirming our hypothesis that a compensation of the bias towards a sharper distribution is useful. Instead, the T was set to 1 during the generation with models trained with label smoothed cross entropy, as in this case a higher (or lower) temperature caused performance losses up to 1 BLEU point.

All experiments were executed on a single machine with 8 Tesla K80 with 11GB RAM. Our implementation is built on top of fairseq (Ott et al., 2019), an open source tool based on PyTorch (Paszke et al., 2019).

5 Results

The MT model used as teacher for Sequence KD and Word KD scored 32.09 BLEU on the MuST-C En-De test set. We trained also a smaller MT model to initialize the ST decoder with it. Moreover, we trained two ASR models. One without the multitask CTC loss and one with it. They scored respectively 14.67 and 10.21 WER. All the ST systems having CTC loss were initialized with the latter, while the others were initialized with the former.

Table 3 shows our ST models' results computed on the MuST-C En-De and IWSLT2015 test set.

Model	MuST-C sentence	MuST-C VAD	IWSLT 2015
Seq KD+FT (w/o TS)	25.80	20.94	17.18
+ FT w/o KD	27.55	19.64	16.93
Multi ENC (w/o TS)	25.79	21.37	19.07
+ FT w/o KD	27.24	20.87	19.08
Multi ENC+DEC PT	25.30	20.80	16.76
+ FT w/o KD	27.40	21.90	18.55
Multi ENC+CTC	27.06	21.58	20.23
+ FT w/o KD (1)	27.98	22.51	20.58
Multi ENC+CTC (5e-3)	25.44	20.41	16.36
+ FT w/o KD	29.08	23.70	20.83
+ AVG 5 (2)	28.82	23.66	21.42
Multi DEC+CTC (5e-3)	26.10	19.94	17.92
+ FT w/o KD	28.22	22.61	18.31
Ensemble (1) and (2)	29.18	23.77	21.83

Table 3: Case sensitive BLEU scores for our E2E ST models. Notes: Seq KD: Sequence KD; FT: finetuning on ground-truth datasets; TS: time stretch; Multi ENC: multi-domain model with sum of the language token to the encoder input; Multi DEC: multi-domain model with sum of the language token to the decoder input; DEC PT: pretraining of the decoder with that of an MT model; CTC: multitask training with CTC loss on the 8th encoder layer in addition to the target loss; FT w/o KD: finetuning on all data with label smoothed cross entropy; 5e-3: indicates the learning rate used; AVG 5: average 5 checkpoints around the best.

5.1 Sequence KD + Finetune VS Multi-domain

First, we compare the two training schemes examined. As shown in Table 3, *Sequence KD + Finetune* [Seq KD+FT] has the same performance as *Multi-domain* with language token summed to the input [Multi ENC] (or even slightly better) on the MuST-C test set, but it is significantly worse on the two test set segmented with VAD. This can be explained by the higher generalization capability of the *Multi-domain* model. Indeed, *Sequence KD + Finetune* seems to overfit more the training data; thus, on data coming from a different distribution, as VAD-segmented data are, its performance drops significantly. For this reason, all the following experiments use the *Multi-domain* training scheme.

5.2 Decoder pretraining and *time stretch*

The pretraining of the decoder with that of an MT model does not bring consistent and significant improvements across the test sets [Multi ENC+DEC PT]. Before the fine-tuning with label smoothed cross entropy, indeed, the model performs worse on all test sets. The fine-tuning, though, helps improving performances on all test sets, which was not the case with the previous training. This can

be related to the introduction of *time stretch*, which reduces the overfitting to the training data. Therefore, we decided to discard the MT pretraining and keep *time stretch*.

5.3 CTC loss and learning rate

The multitask training with CTC loss, instead, improves the results consistently. The model trained with it [Multi ENC+CTC] outperforms all the others on all test sets by up to 1.5 BLEU points. During the fine-tuning of these models, we do not perform multitask training with the CTC loss, so the fine-tuning training is exactly the same as for previous models.

Interestingly, increasing the learning rate [Multi ENC+CTC (5e-3)], the performance before the fine-tuning is worse, but the fine-tuning of this models brings an impressive improvement over all test sets. The reason of this behavior is probably related to a better initial exploration of the solution space thanks to the higher learning rate, which, on the other side, prevents to get very close to the local optimum found. In this scenario, the fine-tuning with a lower learning rate helps getting closer to the local optimum, in addition to the usual benefits.

5.4 Token integration strategy

Finally, we tried adding the language token to the embeddings provided to the decoder, instead of the input data [Multi DEC+CTC (5e-3)]. This was motivated by the idea that propagating this information through the decoder may be more difficult due to the CTC loss, which is not dependent on that information so it may hide it to higher layers. The experiments disproved this hypothesis, as after the fine-tuning the results are lower on all benchmarks.

5.5 Submissions

We averaged our best model over 5 checkpoints, centered in the best according to the validation loss. We also created an ensemble with the resulting model and the best among the others. Both operations were not useful on the two variants of the MuST-C test set, but improved the score on the IWSLT2015 test set. We argue this means that they are more robust and generalize better.

Our *primary* submission has been obtained with the ensemble of two models, scoring 20.75 BLEU on the 2020 test set and 19.52 BLEU on the 2019 test set. Our *contrastive* submission has been generated with the 5 checkpoints average of our best

model, scoring 20.25 BLEU on the 2020 test set and 18.92 BLEU on the 2019 test set.

6 Conclusions

We described FBK’s participation in IWSLT2020 offline speech translation evaluation campaign (Ansari et al., 2020). Our work focused on the integration of transfer learning, data augmentation, multi-task training and the training scheme used to combine real and synthetic data. Based on the results of our experiments, our submission is characterized by a multi-domain training scheme, with additional CTC loss on the transcriptions and word-level knowledge distillation, followed by a fine-tuning on label smoothed cross entropy.

Overall, the paper demonstrates that the combination of the above-mentioned techniques can improve the performance of end-to-end ST models so that they can be competitive with cascaded solutions. Moreover, it shows that *i*) tagged synthetic data leads to more robust models than a pretraining on synthetic data followed by a fine-tuning on datasets with ground-truth targets and *ii*) fine-tuning on label smoothed cross entropy after a training with knowledge distillation brings significant improvements. The huge gap (5.5 BLEU) between data segmented in sentences and data segmented with VAD highlights the need of custom solutions for the latter. In light of these considerations, our future research will focus on techniques to improve the results when the audio segmentation is challenging for ST models.

Acknowledgments

This work is part of the “End-to-end Spoken Language Translation in Rich Data Conditions” project,⁴ which is financially supported by an Amazon AWS ML Grant.

References

Ebrahim Ansari, Amittai Axelrod, Nguyen Bach, Ondrej Bojar, Roldano Cattoni, Fahim Dalvi, Nadir Durrani, Marcello Federico, Christian Federmann, Jiatao Gu, Fei Huang, Kevin Knight, Xutai Ma, Ajay Nagesh, Matteo Negri, Jan Niehues, Juan Pino, Elizabeth Salesky, Xing Shi, Sebastian Stüker, Marco Turchi, and Changhan Wang. 2020. Findings of the IWSLT 2020 Evaluation Campaign. In *Proceedings of the 17th International Conference on Spoken Language Translation (IWSLT 2020)*, Seattle, USA.

⁴<https://ict.fbkc.eu/units-hlt-mt-e2eslt/>

Parnia Bahar, Tobias Bieschke, and Hermann Ney. 2019a. A Comparative Study on End-to-end Speech to Text Translation. In *Proceedings of International Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pages 792–799, Sentosa, Singapore.

Parnia Bahar, Albert Zeyer, Ralf Schlüter, and Hermann Ney. 2019b. On Using SpecAugment for End-to-End Speech Translation. In *Proceedings of 16th International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong.

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, Volume 1 (Long and Short Papers)*, pages 58–68, Minneapolis, Minnesota. Association for Computational Linguistics.

Loïc Barrault, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. **Findings of the 2019 Conference on Machine Translation (WMT19)**. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61, Florence, Italy. Association for Computational Linguistics.

Nicola Bertoldi, Roldano Cattoni, Mauro Cettolo, Amin Farajian, Marcello Federico, Davide Caroselli, Luca Mastrostefano, Andrea Rossi, Marco Trombetti, Ulrich Germann, and David Madl. 2017. MMT: New Open Source MT for the Translation Industry. In *Proceedings of the 20th Annual Conference of the European Association for Machine Translation (EAMT)*, pages 86–91, Prague, Czech Republic.

Isaac Caswell, Ciprian Chelba, and David Grangier. 2019. **Tagged Back-Translation**. In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 53–63, Florence, Italy. Association for Computational Linguistics.

Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, Roldano Cattoni, and Marcello Federico. 2015. The IWSLT 2015 Evaluation Campaign. In *Proceedings of 12th International Workshop on Spoken Language Translation (IWSLT)*, Da Nang, Vietnam.

Mattia Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019a. **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)*, page 2012–2017, Minneapolis, Minnesota.

- Mattia Antonino Di Gangi, Matteo Negri, Viet Nhat Nguyen, Amirhossein Tebbifakhr, and Marco Turchi. 2019b. Data Augmentation for End-to-End Speech Translation: FBK@IWSLT '19. In *Proceedings of 16th International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong.
- Mattia Antonino Di Gangi, Matteo Negri, and Marco Turchi. 2019c. [Adapting Transformer to End-to-End Spoken Language Translation](#). In *Proceedings of Interspeech 2019*, pages 1133–1137, Graz, Austria.
- Mattia Antonino Di Gangi, Matteo Negri, and Marco Turchi. 2019d. One-To-Many Multilingual End-to-end Speech Translation. In *Proceedings of the 2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 585–592, Sentosa, Singapore.
- Alex Graves, Santiago Fernández, Faustino J. Gomez, and Jürgen Schmidhuber. 2006. Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks. In *Proceedings of the 23rd international conference on Machine learning (ICML)*, pages 369–376, Pittsburgh, Pennsylvania.
- Junxian He, Jiatao Gu, Jiajun Shen, and Marc’Aurelio Ranzato. 2020. Revisiting Self-Training for Neural Sequence Generation. In *Proceedings of the International Conference on Learning Representations (ICLR)*, Virtual Conference.
- François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia Tomashenko, and Yannick Estève. 2018. [TED-LIUM 3: Twice as Much Data and Corpus Repartition for Experiments on Speaker Adaptation](#). In *Proceedings of the Speech and Computer - 20th International Conference (SPECOM)*, pages 198–208, Leipzig, Germany. Springer International Publishing.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. [Distilling the Knowledge in a Neural Network](#). In *Proceedings of NIPS Deep Learning and Representation Learning Workshop*, Montréal, Canada.
- Javier Iranzo-Sánchez, Joan Albert Silvestre-Cerdà, Javier Jorge, Nahuel Roselló, Giménez. Adrià, Albert Sanchis, Jorge Civera, and Alfons Juan. 2020. [Europarl-ST: A Multilingual Corpus For Speech Translation Of Parliamentary Debates](#). In *Proceedings of 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8229–8233, Barcelona, Spain.
- Ye Jia, Melvin Johnson, Wolfgang Macherey, Ron J. Weiss, Yuan Cao, Chung-Cheng Chiu, Naveen Ari, Stella Laurenzo, and Yonghui Wu. 2019. Leveraging Weakly Supervised Data to Improve End-to-End Speech-to-Text Translation. In *Proceedings of the 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7180–7184, Brighton, UK.
- Suyoun Kim, Takaaki Hori, and Shinji Watanabe. 2017. Joint CTC-Attention based End-to-End Speech Recognition using Multi-task Learning. In *Proceedings of the 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4835–4839, New Orleans, Louisiana.
- Yoon Kim and Alexander M. Rush. 2016. [Sequence-Level Knowledge Distillation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas.
- Diederik Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of 3rd International Conference on Learning Representations (ICLR)*, San Diego, California.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open Source Toolkit for Statistical Machine Translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Solomon Kullback and Richard Arthur Leibler. 1951. [On information and sufficiency](#). *Ann. Math. Statist.*, 22(1):79–86.
- Pierre Lison and Jörg Tiedemann. 2016. OpenSubtitles2016: Extracting Large Parallel Corpora from Movie and TV Subtitles. In *Proceedings of the Tenth International Language Resources and Evaluation Conference (LREC)*, pages 923–929, Portoroz, Slovenia.
- 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 *Proceedings of Interspeech 2019*, pages 1128–1132, Graz, Austria.
- Sylvain Meignier and Teva Merlin. 2010. LIUM SpkDiarization: An Open Source Toolkit For Diarization. In *Proceedings of the CMU SPUD Workshop*, Dallas, Texas.
- Graham Neubig, Matthias Sperber, Xinyi Wang, Matthieu Felix, Austin Matthews, Sarguna Padmanabhan, Ye Qi, Devendra Sachan, Philip Arthur, Pierre Godard, John Hewitt, Rachid Riad, and Liming Wang. 2018. [XNMT: The eXtensible Neural Machine Translation Toolkit](#). In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Papers)*, pages 185–192, Boston, MA. Association for Machine Translation in the Americas.

- Thai-Son Nguyen, Sebastian Stueker, Jan Niehues, and Alex Waibel. 2020. Improving Sequence-to-sequence Speech Recognition Training with On-the-fly Data Augmentation. In *Proceedings of the 2020 International Conference on Acoustics, Speech, and Signal Processing – IEEE-ICASSP-2020*, Barcelona, Spain.
- Jan Niehues, Roldano Cattoni, Sebastian Stueker, Matteo Negri, Marco Turchi, Thanh-Le Ha, Elizabeth Salesky, Ramon Sanabria, Loic Barrault, Lucia Specia, and Marcello Federico. 2019. [The IWSLT 2019 Evaluation Campaign](#). In *Proceedings of 16th International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A Fast, Extensible Toolkit for Sequence Modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. LibriSpeech: An ASR corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210, South Brisbane, Queensland, Australia.
- 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 *Proceedings of Interspeech 2019*, pages 2613–2617, Graz, Austria.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [PyTorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. Alché-Buc, E. Fox, and R. Garnett, editors, *Proceedings of Advances in Neural Information Processing Systems 32 (NIPS)*, pages 8024–8035. Curran Associates, Inc.
- Ramon Sanabria, Ozan Caglayan, Shruti Palaskar, Desmond Elliott, Loïc Barrault, Lucia Specia, and Florian Metze. 2018. [How2: A Large-scale Dataset For Multimodal Language Understanding](#). In *Proceedings of Visually Grounded Interaction and Language (ViGIL)*, Montréal, Canada. Neural Information Processing Society (NeurIPS).
- 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 (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the Inception Architecture for Computer Vision. In *Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, Las Vegas, Nevada, United States.
- Xu Tan, Yi Ren, Di He, Tao Qin, and Tie-Yan Liu. 2019. [Multilingual Neural Machine Translation with Knowledge Distillation](#). In *Proceedings of International Conference on Learning Representations (ICLR)*, New Orleans, Louisiana, United States.
- 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 *Proceedings of Advances in Neural Information Processing Systems 30 (NIPS)*, pages 5998–6008, Long Beach, California.
- Ron J. Weiss, Jan Chorowski, Navdeep Jaitly, Yonghui Wu, and Zhifeng Chen. 2017. [Sequence-to-Sequence Models Can Directly Translate Foreign Speech](#). In *Proceedings of Interspeech 2017*, pages 2625–2629, Stockholm, Sweden.