

# Dealing with training and test segmentation mismatch: FBK@IWSLT2021

Sara Papi<sup>1,2</sup>, Marco Gaido<sup>1,2</sup>, Matteo Negri<sup>1</sup>, Marco Turchi<sup>1</sup>

<sup>1</sup>Fondazione Bruno Kessler, Trento, Italy

<sup>2</sup>University of Trento, Italy

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

## Abstract

This paper describes FBK’s system submission to the IWSLT 2021 Offline Speech Translation task. We participated with a direct model, which is a Transformer-based architecture trained to translate English speech audio data into German texts. The training pipeline is characterized by knowledge distillation and a two-step fine-tuning procedure. Both knowledge distillation and the first fine-tuning step are carried out on manually segmented real and synthetic data, the latter being generated with an MT system trained on the available corpora. Differently, the second fine-tuning step is carried out on a random segmentation of the MuST-C v2 En-De dataset. Its main goal is to reduce the performance drops occurring when a speech translation model trained on manually segmented data (i.e. an ideal, sentence-like segmentation) is evaluated on automatically segmented audio (i.e. actual, more realistic testing conditions). For the same purpose, a custom hybrid segmentation procedure that accounts for both audio content (pauses) and for the length of the produced segments is applied to the test data before passing them to the system. At inference time, we compared this procedure with a baseline segmentation method based on Voice Activity Detection (VAD). Our results indicate the effectiveness of the proposed hybrid approach, shown by a reduction of the gap with manual segmentation from 8.3 to 1.4 BLEU points.

## 1 Introduction

Speech translation (ST) is the task of translating a speech uttered in one language into its textual representation in a different language. Unlike *simultaneous* ST, where the audio is translated as soon as it is produced, in the *offline* setting the audio is entirely available and translated at once. In continuity with the last two rounds of the IWSLT evaluation campaign (Niehues et al., 2019; Ansari

et al., 2020), the IWSLT2021 Offline Speech Translation task (Anastasopoulos et al., 2021) focused on the translation into German of English audio data extracted from TED talks. Participants could approach the task either with a cascade architecture or with a direct end-to-end system. The former represents the traditional pipeline approach (Stentiford and Steer, 1988; Waibel et al., 1991) comprising an automatic speech recognition (ASR) followed by a machine translation (MT) component. The latter (Bérard et al., 2016; Weiss et al., 2017) relies on a single neural network trained to translate the input audio into target language text bypassing any intermediate symbolic representation steps.

The two paradigms have advantages and disadvantages. Cascade architectures have historically guaranteed higher translation quality (Niehues et al., 2018, 2019) thanks to the large corpora available to train their ASR and MT sub-components. However, a well-known drawback of pipelined solutions is represented by error propagation: transcription errors are indeed hard (and sometimes impossible) to recover during the translation step. Direct models, although being penalized by the paucity of training data, have two theoretical competitive advantages, namely: *i*) the absence of error propagation as there are no intermediate processing steps, and *ii*) a less mediated access to the source utterance, which allows them to better exploit speech information (e.g. prosody) without loss of information.

The paucity of parallel (audio, translation) data for direct ST has been previously addressed in different ways, ranging from *model pre-training* to exploit knowledge transfer from ASR and/or MT (Bérard et al., 2018; Bansal et al., 2019; Alinejad and Sarkar, 2020), *knowledge distillation* (Liu et al., 2019; Gaido et al., 2021a), *data augmentation* (Jia et al., 2019; Bahar et al., 2019b; Nguyen et al., 2020), and *multi-task learning* (Weiss et al.,

2017; Anastasopoulos and Chiang, 2018; Bahar et al., 2019a; Gaido et al., 2020b). Thanks to these studies, the gap between the strong cascade models and the new end-to-end ones has gradually reduced during the last few years. As highlighted by the IWSLT 2020 Offline Speech Translation challenge results (Ansari et al., 2020), the rapid evolution of the direct approach has eventually led it to performance scores that are similar to those of cascade architectures. In light of this positive trend, we decided to adopt only the direct approach (described in Section 3) for our participation in the 2021 round of the offline ST task.

Another interesting finding from last year’s campaign concerns the sensitivity of ST models to different segmentations of the input audio. The 2020 winning system (Potapczyk and Przybysz, 2020) shows that, with a custom segmentation of the test data, the same model improved by 3.81 BLEU points the score achieved when using the basic segmentation provided by the task organizers. This noticeable difference is due to a well-known problem in MT, ST and in machine learning at large: any mismatch between training and test data (in terms of domain, text style or a variety of other aspects) can cause unpredictable, often large, performance drops at test time. In ST, this is a critical issue, inherent to the nature of the available resources: while systems are usually trained on corpora that are manually segmented at sentence level, test data come in the form of unsegmented continuous speech.

A possible solution to this problem is to automatically segment the test data with a Voice Activity Detection (VAD) tool (Sohn et al., 1999). This strategy tries to mimic the sentence-based segmentation observed in the training data using pauses as an indirect (hence known to be sub-optimal) cue for sentence boundaries. Custom segmentation strategies, which are allowed to IWSLT participants, typically go in this direction with the aim to reduce the data mismatch by working on evaluation data. An opposite way to look at the problem is to work on the training data. In this case, the goal is to “robustify” the ST model to noisy inputs (i.e. sub-optimal segmentations) at training time, by exposing it to perturbed data where sentence-like boundaries are not guaranteed. Our participation in the offline ST task exploits both solutions (see Section 4): at training time, by fine-tuning the model with a random segmentation of the available in-domain data;

at test time, by feeding it with a custom hybrid segmentation of the evaluation data.

In a nutshell, our participation can be summarized as follows. After a preliminary model selection phase that was carried out in order to select the best architecture, we adopted a pipeline consisting of: *i*) ASR pre-training, *ii*) ST training with knowledge distillation with an MT teacher, and *iii*) two-step fine-tuning by varying the type and the amount of data between the two steps. The second fine-tuning step, which was carried out on artificially perturbed data to increase model robustness, represents the main aspect characterizing our participation to this year’s round of the offline ST task together with our custom automatic segmentation of the test set (see Section 4). Our experimental results proved the effectiveness of our solutions: compared to a standard ST model and a baseline VAD-based method, on the MuST-C v2 English-German test set (Cattoni et al., 2021), the gap with optimal manual segmentation is reduced from 8.3 to 1.4 BLEU.

## 2 Training data

To build our models, we used most of the training data allowed for participation.<sup>1</sup> They include: MT corpora (English-German text pairs), ASR corpora (English audios and their corresponding transcripts) and ST corpora (English audios with corresponding English transcripts and German translations).

**MT.** Among all the available datasets, we selected those allowed for WMT 2019 (Barrault et al., 2019) and OpenSubtitles2018 (Lison and Tiedemann, 2016). Some pre-processing was required to isolate and remove different types of potentially harmful noise present in the data. These include non-unicode characters, both on the source and target side of the parallel sentence pairs, which would have led to an increased dictionary size hindering model training, and whole non-German target sentences (mostly in English). The cleaning of this two types of noise, which was respectively performed using a custom script and Modern MT (Bertoldi et al., 2017), resulted in the removal of roughly 25% of the data, with a final dataset of  $\sim 49$  million sentence pairs.

**ASR.** ASR corpora, together with the ST ones described below, were collected for the ASR training. In detail, the allowed native ASR datasets are:

<sup>1</sup><https://iwslt.org/2021/offline>

LibriSpeech (Panayotov et al., 2015), TEDLIUM v3 (Hernandez et al., 2018) and Mozilla Common Voice.<sup>2</sup> In all of them, English texts were lower-cased and punctuation was removed.

**ST.** The ST benchmarks we used are essentially three: *i*) Europarl-ST (obtained from European Parliament debates – Iranzo-Sánchez et al. 2020), *ii*) MuST-C v2 (built from TED talks – Cattoni et al. 2021), and *iii*) CoVoST 2 (containing the translations of a portion of the Mozilla Common Voice dataset – Wang et al. 2020a). To cope with the scarcity of ST data, we complemented these native ST corpora with synthetic data. To this aim, we used the MT system trained on the available MT data to translate into German the English transcripts of the aforementioned ASR datasets. The resulting texts were used as reference material during the ST model training. The combination of native and generated data resulted in a total of about 1.26 million samples. The transcription-translation pairs were tokenized using, respectively, source/target-language SentencePiece (Sennrich et al., 2016) unigram models trained on the MT corpora with a vocabulary size of 32k tokens. Similar to our last year’s IWSLT submission (Gaido et al., 2020b), the entire dataset was used for training in a multi-domain fashion, where the two domains were *native* (original ST data) and *generated* (synthetic data).

Prior to the extraction of the speech features, the audio was pre-processed with the SpecAugment (Park et al., 2019) data augmentation technique, which masks consecutive portions of the input both in frequency and in time dimensions. From all the audio files, 80 log Mel-filter banks features were extracted using PyKaldi (Can et al., 2018), filtering out those samples containing more than 3,000 frames. Finally, we applied utterance level Cepstral Mean and Variance Normalization both during ASR pre-training and ST training phases. The configuration parameters used are the default ones as set in (Wang et al., 2020b).

### 3 Model and training

In order to select the best performing architecture, we trained several Transformer-based models (Vaswani et al., 2017), which consist of 12 encoder layers, 6 decoder layers, 8 attention heads, 512 features for the attention layers and 2,048 hidden

<sup>2</sup><https://commonvoice.mozilla.org/en/datasets>

units in the feed-forward layers. The ASR and ST models are based on a custom version of the model by (Wang et al., 2020b), which is a Transformer whose encoder has two initial 1D convolutional layers with *gelu* activation functions (Hendrycks and Gimpel, 2020). Also, the encoder self-attentions were biased using a logarithmic distance penalty in favor of the local context as per (Di Gangi et al., 2019). A Connectionist Temporal Classification (CTC) scoring function was applied as described in (Gaido et al., 2020b). This was done by adding a linear layer to either the 6th, 8th or 10th encoder layer to map the encoder states to the vocabulary size and compute the CTC loss. The choice of the final architecture, depending on where the CTC loss is applied, was made based on sacreBLEU score (Post, 2018) after training the models on MuST-C v1 En-De (Cattoni et al., 2021). ST results computed on the test set are reported on Table 1. As it can be seen from the table, two models obtained the highest, identical BLEU score (21.21): they both use logarithmic distance penalty but apply CTC loss to the 6th or the 8th encoder layer.

#### 3.1 Training pipeline

In the following, we describe the pipeline used to build our ST models, as anticipated in Section 1. In details, the ASR model is trained and its encoder used as starting point for the ST model, which is first trained via knowledge distillation and then fine-tuned on native and synthetic data. Then, a second fine-tuning step is performed on a perturbed version of a subset of the native data, focused on reducing the model performance drop over different segmentations. For the initial ST training, we optimized KL divergence (Kullback and Leibler, 1951) and CTC losses. For the first fine-tuning step, we optimized label smoothed cross entropy (LSCE) or CTC+LSCE while, for the second fine-tuning step, the models were refined using LSCE only, with a lower learning rate in order not to override the knowledge acquired during the previous phases.

**ASR pre-training.** Due to the identical BLEU score obtained by applying the CTC loss to the 6th and 8th layer during the ST model selection phase, we opted for training the ASR system using both these architectures, and selected the final model by looking at the Word Error Rate (WER) achieved by averaging 7 checkpoints around the best one. As shown in Table 2, the best overall performing architecture is the one where the CTC is applied to

architecture	CTC encoder layer	distance penalty	BLEU
2d convolutional	6	no	19.04
1d convolutional	6	no	21.16
1d convolutional	6	log	21.21
1d convolutional	8	log	21.21
1d convolutional	10	log	21.08

Table 1: Results of 1d convolutional architectures trained computing CTC loss at different layers and with/without distance penalty. Also the result of a 2d convolutional architecture is reported where the structure is exactly the same except for the use of a different type of convolution.

model	dev	test
CTC on 6th encoder layer	8.67	12.19
<b>CTC on 8th encoder layer</b>	<b>7.52</b>	<b>10.70</b>

Table 2: Results of ASR pre-training in terms of WER. The dev and test sets used are, respectively, dev and tst-COMMON of MuST-C v1 En-De.

the 8th encoder layer. Accordingly, we used this architecture to perform all the successive training phases.

**Training with knowledge distillation.** Two ST models, one with 12 and one with 15 encoder layers, were trained by loading the pre-trained ASR encoder weights and applying word-level Knowledge Distillation (KD) as in (Kim and Rush, 2016). In KD, a *student* model is trained with the goal of learning how to produce the same output distribution as a *teacher* model, and this is obtained by computing the KL divergence between the two output distributions. In our setting, the student and the teacher are respectively the ST system and an MT system that we trained on the MT data described in Section 2. It consists in a plain Transformer model 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 the feed-forward layers. Evaluated on the MuST-C v2 En-De test set, it achieved a BLEU score of 33.3. For ST training with KD, we extracted only the top 8 tokens from the teacher distribution. According to (Tan et al., 2019), this choice results in a significant reduction of the memory required, with no loss in final performance. At the end of this phase, we decided to keep the model with 15 encoder layers as it performs better than the one with 12 encoder layers by 1 BLEU point.

**Fine-tuning step #1: using native and synthetic data.** Once the KD training phase was concluded, we performed a multi-domain fine-tuning where

the ST model was jointly trained on native and synthetic data optimizing LSCE or its combination with the CTC loss.

#### 4 Coping with training/test data mismatch

As mentioned in Section 1, the segmentation of audio files is a crucial aspect in ST. In fact, mismatches between the manual segmentation of the training data and the automatic one required when processing the unsegmented test set can produce significant performance drops. To mitigate this risk, we worked on two complementary fronts: at training and inference time. At training time, we tried to robustify our model by fine-tuning it on a randomly segmented subset of the training data. At inference time, we applied an automatic segmentation procedure to the test set in order to feed the model with input resembling, as much as possible, the gold manual segmentation. These two solutions, which characterize our final submission, are explained in the following.

**Fine-tuning step #2: using randomly segmented data.** For the second fine-tuning step, we re-segmented the MuST-C v2 En-De training set following the procedure described in (Gaido et al., 2020a). The method consists in choosing a random word in the transcript of each sample, and using it as sentence boundary instead of the linguistically-motivated (sentence-level) splits provided in the original data. The corresponding audio segments are then obtained by means of audio-text alignments performed with Gentle.<sup>3</sup> Similarly, the German translation of each re-segmented transcript is extracted with cross-lingual alignments generated by a fast\_align (Dyer et al., 2013) model trained on all the MT data available for the task and on MuST-C v2. In case either of the alignments is

<sup>3</sup><https://github.com/lowerquality/gentle/>

model	MuST-C2	MuST-C2	MuST-C2	IWSLT2015	IWSLT2015
	manual	VAD (WebRTC)	hybrid	VAD (LIUM)	hybrid
1-FT LSCE	27.6	20.8	24.8	16.1	21.9
2-FT LSCE	-	23.4 (+2.6)	<b>26.4</b> (+1.6)	20.7 (+4.6)	22.7 (+0.8)
1-FT LSCE+CTC	27.7	19.9	25.3	14.0	21.7
2-FT LSCE+CTC	-	<b>23.7</b> (+3.8)	26.3 (+1.0)	<b>20.9</b> (+6.9)	<b>23.1</b> (+1.4)

Table 3: Results of the best architectures deriving from KD training after one or two fine-tuning steps. 1-FT stands for one-step fine-tuning and 2-FT stands for two-step fine-tuning (see Section 3). MuST-C v2 results on manual segmentation have been not computed for the 2-step fine-tuned models as we were interested in the evaluation of the improvement on automatically segmented data.

not possible (because `fast_align` is not able to align enough words or Gentle does not recognize the position of the word in the audio), the sentence is discarded. The resulting material, which contains  $\sim 5\%$  less segments than the original MuST-C release, was then used for our second (and final) fine-tuning step. As already stated, we used only the LSCE loss for this stage.

**Automatic segmentation of the test data.** At inference time, the test set was segmented with an hybrid approach that considers both the audio content and the length of the resulting segment (Gaido et al., 2021b). Specifically, every segment is ensured to be at least 17s and at most 20s long, but the exact splitting position is determined by the longest pause detected within this interval. Pauses are identified with the WebRTC VAD tool (Johnston and Burnett, 2012), using 20ms as *frame duration* and 2 as *aggressivity* level.

## 5 Experimental settings

Our implementation is built on top of fairseq Pytorch library (Ott et al., 2019). All our models were trained using the Adam optimizer (Kingma and Ba, 2015) with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.98$ . During training, the learning rate was set to increase linearly from 0 to  $2e-3$  for the first 10,000 warm-up steps and then to decay with an inverse square root policy. Differently, the learning rate was kept constant for model fine-tuning, with a value of  $1e-3$  for the first fine-tuning step and  $1e-4$  for the second one.

All the trainings were performed on 2 Tesla V100 GPUs with 32GB RAM. We set the maximum number of tokens to 10k per batch and 8 as update frequency. For generation, the maximum number of tokens was increased to 50k, using a single Tesla V100 GPU and by applying a standard 5-beam search strategy.

## 6 Results

For the evaluation of the fine-tuned models we considered three different test sets: MuST-C v2 En-De tst-COMMON, IWSLT 2015 and 2019 test sets (available on the Offline ST task Evaluation Campaign web page<sup>4</sup>). While for MuST-C v2 we originally had a manual segmentation of the audio files, for the IWSLT 2015 and 2019 test sets the organizers provided only automatic segmentations obtained by the LIUM VAD tool (Meignier and Merlin, 2010). Furthermore, we segmented MuST-C v2 tst-COMMON using the WebRTC VAD tool to have a comparable framework. Table 3 reports the results before and after the second fine-tuning step, which clearly show that performing the additional training on randomly segmented data highly improves the performance in the non-manual segmentation case, by up to 6 BLEU points. We also created an ensemble with the best two models reported in Table 3, whose KD training also used CTC loss. Results are not reported here since ensembling did not bring any improvement in terms of BLEU score compared to the two separate models. A possible motivation is that our two-step fine-tuning process is already sufficient to build a robust model, which is capable of generalizing without the need of combining two or more model outputs.

For our *primary* submission, we chose the two-step fine-tuned model that uses the LSCE+CTC losses for the first fine-tuning step (2-FT LSCE+CTC) since it achieved the highest BLEU on automatically segmented data. In order to measure the contribution of fine-tuning on randomly segmented data also on the official evaluation set, we selected the same model before the second fine-tuning step (1-FT LSCE+CTC) as our *contrastive* submission.

<sup>4</sup><https://iwslt.org/2021/offline>

Our primary submission scored 30.6 BLEU on the tst2021 test set considering both references while our contrastive scored 29.3 BLEU, showing the effectiveness of our fine-tuning step. In addition, our primary submission scored 24.7 BLEU on the tst2020 test set.

## 7 Conclusions

We described FBK’s participation in the IWSLT2021 Offline Speech Translation task (Anastasopoulos et al., 2021). Our work focused on a multi-step training pipeline involving data augmentation (SpecAugment and MT-based synthetic data), multi-domain transfer learning (KD training first and then fine-tuning on synthetic and native data) and ad-hoc fine-tuning on randomly segmented data. Based on the experimental results, our submission was characterized by the use of the CTC loss on transcripts during word-level knowledge distillation training, followed by a two-stage fine-tuning aimed to fill the gap between the performance of models when tested on manual and automatically segmented data. This huge gap was pointed out in our last year submission (Gaido et al., 2020b), where we highlighted that some strategies should have been adopted in order to mitigate the problem. This paper demonstrates that, following the above-mentioned pipeline, together with some data-driven techniques, we can obtain significant improvements in the performance of end-to-end ST systems. Research in this direction will help us to build models that are not only competitive with cascaded solutions, but also able to handle different segmentation strategies which are going to be more frequently used in the future.

## References

- Ashkan Alinejad and Anoop Sarkar. 2020. [Effectively pretraining a speech translation decoder with Machine Translation data](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8014–8020, Online. Association for Computational Linguistics.
- Antonios Anastasopoulos, Ondřej Bojar, Jacob Bremerman, Roldano Cattoni, Maha Elbayad and Marcello Federico, Xutai Ma, Satoshi Nakamura, Matteo Negri, Jan Niehues, Juan Pino, Elizabeth Salesky, Sebastian Stüker, Katsuhito Sudoh, Marco Turchi, Alex Waibel, Changhan Wang, and Matthew Wiesner. 2021. [Findings of the IWSLT 2021 Evaluation Campaign](#). In *Proceedings of the 18th International Conference on Spoken Language Translation (IWSLT 2021)*, Online.
- Antonios Anastasopoulos and David Chiang. 2018. [Tied Multitask Learning for Neural Speech Translation](#). In *Proc. of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 82–91, New Orleans, Louisiana.
- Ebrahim Ansari, Amittai Axelrod, Nguyen Bach, Ondřej 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, Alexander Waibel, and Changhan Wang. 2020. [FINDINGS OF THE IWSLT 2020 EVALUATION CAMPAIGN](#). In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 1–34, Online. Association for Computational Linguistics.
- Parnia Bahar, Tobias Bieschke, and Hermann Ney. 2019a. [A Comparative Study on End-to-End Speech to Text Translation](#). In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 792–799.
- Parnia Bahar, Albert Zeyer, Ralf Schlüter, and Hermann Ney. 2019b. [On Using SpecAugment for End-to-End Speech Translation](#). In *Proc. of the International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong, China.
- 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 *Proc. of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 58–68, Minneapolis, Minnesota.
- 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.
- Alexandre Bérard, Laurent Besacier, Ali Can Kocabiyikoglu, and Olivier Pietquin. 2018. [End-to-End Automatic Speech Translation of Audiobooks](#). In *Proc. of ICASSP 2018*, pages 6224–6228, Calgary, Alberta, Canada.
- 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 *The 20th Annual Conference of the European Association for Machine Translation (EAMT)*. 20th Annual Conference of the European Association for

- Machine Translation, EAMT 2017 ; Conference date: 29-05-2017 Through 31-05-2017.
- Alexandre Bérard, Olivier Pietquin, Christophe Servan, and Laurent Besacier. 2016. Listen and Translate: A Proof of Concept for End-to-End Speech-to-Text Translation. In *NIPS Workshop on end-to-end learning for speech and audio processing*, Barcelona, Spain.
- Dogan Can, Victor R. Martinez, Pavlos Papadopoulos, and Shrikanth S. Narayanan. 2018. Pykaldi: A python wrapper for kaldi. In *Acoustics, Speech and Signal Processing (ICASSP), 2018 IEEE International Conference on*. IEEE.
- Roldano Cattoni, Mattia A. Di Gangi, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2021. [MuST-C: A multilingual corpus for end-to-end speech translation](#). *Computer Speech & Language*, 66:101155.
- Mattia A. Di Gangi, Matteo Negri, and Marco Turchi. 2019. [Adapting Transformer to End-to-End Spoken Language Translation](#). In *Proc. Interspeech 2019*, pages 1133–1137.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. A Simple, Fast, and Effective Reparameterization of IBM Model 2. In *Proc. of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 644–648, Atlanta, Georgia.
- Marco Gaido, Mattia A. Di Gangi, Matteo Negri, Mauro Cettolo, and Marco Turchi. 2020a. [Contextualized Translation of Automatically Segmented Speech](#). In *Proc. Interspeech 2020*, pages 1471–1475.
- Marco Gaido, Mattia A. Di Gangi, Matteo Negri, and Marco Turchi. 2020b. [End-to-end speech-translation with knowledge distillation: FBK@IWSLT2020](#). In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 80–88, Online. Association for Computational Linguistics.
- Marco Gaido, Mattia A. Di Gangi, Matteo Negri, and Marco Turchi. 2021a. [On Knowledge Distillation for Direct Speech Translation](#). In *Proceedings of CLiC-IT 2020*, Online.
- Marco Gaido, Matteo Negri, Mauro Cettolo, and Marco Turchi. 2021b. [Beyond voice activity detection: Hybrid audio segmentation for direct speech translation](#).
- Dan Hendrycks and Kevin Gimpel. 2020. [Gaussian Error Linear Units \(GELUs\)](#).
- François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia A. Tomashenko, and Yannick Estève. 2018. [TED-LIUM 3: Twice as much data and corpus reparation for experiments on speaker adaptation](#). In *Speech and Computer - 20th International Conference, SPECOM 2018, Leipzig, Germany, September 18-22, 2018, Proceedings*, volume 11096 of *Lecture Notes in Computer Science*, pages 198–208. Springer.
- Javier Iranzo-Sánchez, Joan Albert Silvestre-Cerdà, Javier Jorge, Nahuel Roselló, Adrià Giménez, Albert Sanchis, Jorge Civera, and Alfons Juan. 2020. [EuroParl-ST: A Multilingual Corpus for Speech Translation of Parliamentary Debates](#). In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8229–8233.
- 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 *Proc. of ICASSP 2019*, pages 7180–7184, Brighton, UK.
- Alan B. Johnston and Daniel C. Burnett. 2012. *WebRTC: APIs and RTCWEB Protocols of the HTML5 Real-Time Web*. Digital Codex LLC, St. Louis, MO, USA.
- 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. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- 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 Conference on Language Resources and Evaluation (LREC'16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).
- 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 *Proc. of Interspeech 2019*, pages 1128–1132.
- Sylvain Meignier and Teva Merlin. 2010. [LIUM SP-KDIARIZATION: AN OPEN SOURCE TOOLKIT FOR DIARIZATION](#). In *CMU SPUD Workshop*, Dallas, United States.
- 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 *Proc. of the 2020 Interna-*

- tional Conference on Acoustics, Speech, and Signal Processing – IEEE-ICASSP-2020*, Barcelona, Spain.
- Jan Niehues, Roldano Cattoni, Sebastian Stüker, Mauro Cettolo, Marco Turchi, and Marcello Federico. 2018. The IWSLT 2018 Evaluation Campaign. In *Proceedings of the 15th International Workshop on Spoken Language Translation*, Bruges, Belgium.
- Jan Niehues, Roldano Cattoni, Sebastian Stüker, Matteo Negri, Marco Turchi, Elizabeth Salesky, Ramon Sanabria, Loïc Barrault, Lucia Specia, Marcello Federico, and et al. 2019. [The IWSLT 2019 Evaluation Campaign](#). In *Proceedings of the 16th International Workshop on Spoken Language Translation*.
- 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 NAACL-HLT 2019: Demonstrations*.
- 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.
- 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 *Proc. Interspeech 2019*, pages 2613–2617.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Tomasz Potapczyk and Pawel Przybysz. 2020. [SR-POL’s system for the IWSLT 2020 end-to-end speech translation task](#). In *Proc. of IWSLT*, Online.
- 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.
- Jongseo Sohn, Nam Soo Kim, and Wonyong Sung. 1999. A statistical model-based voice activity detection. *IEEE Signal Processing Letters*, 6(1).
- Frederick W. M. Stentiford and Martin G. Steer. 1988. Machine Translation of Speech. *British Telecom Technology Journal*, 6(2):116–122.
- Xu Tan, Yi Ren, Di He, Tao Qin, and Tie-Yan Liu. 2019. [Multilingual neural machine translation with knowledge distillation](#). In *International Conference on Learning Representations*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Alex Waibel, Ajay N. Jain, Arthur E. McNair, Hiroaki Saito, Alexander G. Hauptmann, and Joe Tebelskis. 1991. JANUS: A Speech-to-Speech Translation System Using Connectionist and Symbolic Processing Strategies. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing, ICASSP 1991*, pages 793–796, Toronto, Canada.
- Changhan Wang, Juan Pino, Anne Wu, and Jiatao Gu. 2020a. [CoVoST: A diverse multilingual speech-to-text translation corpus](#). In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 4197–4203, Marseille, France. European Language Resources Association.
- Changhan Wang, Yun Tang, Xutai Ma, Anne Wu, Dmytro Okhonko, and Juan Pino. 2020b. [fairseq s2t: Fast speech-to-text modeling with fairseq](#). In *Proceedings of the 2020 Conference of the Asian Chapter of the Association for Computational Linguistics (ACL): System Demonstrations*.
- 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.