

# The NiuTrans End-to-End Speech Translation System for IWSLT 2021 Offline Task

Chen Xu<sup>1</sup>, Xiaoqian Liu<sup>1</sup>, Xiaowen Liu<sup>1</sup>, Laohu Wang<sup>1</sup>, Canan Huang<sup>1</sup>,  
Tong Xiao<sup>1,2</sup>, Jingbo Zhu<sup>1,2</sup>

<sup>1</sup>NLP Lab, School of Computer Science and Engineering

Northeastern University, Shenyang, China

<sup>2</sup>NiuTrans Research, Shenyang, China

{xuchenneu, liuxiaoqianneu, liuxiaowenneu}@outlook.com,  
{tigerneu, huangcananneu}@outlook.com,  
{xiaotong, zhujingbo}@mail.neu.edu.cn

## Abstract

This paper describes the submission of the NiuTrans end-to-end speech translation system for the IWSLT 2021 offline task, which translates from the English audio to German text directly without intermediate transcription. We use the Transformer-based model architecture and enhance it by Conformer, relative position encoding, and stacked acoustic and textual encoding. To augment the training data, the English transcriptions are translated to German translations. Finally, we employ ensemble decoding to integrate the predictions from several models trained with the different datasets. Combining these techniques, we achieve 33.84 BLEU points on the MuST-C En-De test set, which shows the enormous potential of the end-to-end model.

## 1 Introduction

Speech translation (ST) aims to learn models that can predict, given some speech in the source language, the translation into the target language. End-to-end (E2E) approaches have become popular recently for its ability to free designers from cascading different systems and shorten the pipeline of translation (Duong et al., 2016; Berard et al., 2016; Weiss et al., 2017). This paper describes the submission of the NiuTrans E2E ST system for the IWSLT 2021 (Anastasopoulos et al., 2021) offline task, which translates from the English audio to the German text translation directly without intermediate transcription.

Our baseline model is based on the DLCL Transformer (Vaswani et al., 2017; Wang et al., 2019) model with Connectionist Temporal Classification (CTC) (Graves et al., 2006) loss on the encoders (Bahar et al., 2019). We enhance it with the superior model architecture Conformer (Gulati et al.,

2020), relative position encoding (RPE) (Shaw et al., 2018), and stacked acoustic and textual encoding (SATE) (Xu et al., 2021). To augment the training data, the English transcriptions of the automatic speech recognition (ASR) and speech translation corpora are translated to the German translation. Finally, we employ the ensemble decoding method to integrate the predictions from multiple models (Wang et al., 2018) trained with the different datasets.

This paper is structured as follows. The training data is summarized in Section 2, then we describe the model architecture in Section 3 and data augmentation in Section 4. We present the ensemble decoding method in Section 5. The experimental settings and final results are shown in Section 6.

## 2 Training Data

Our system is built under the constraint condition. The training data can be divided into three categories: ASR, MT, and ST corpora<sup>1</sup>.

**ASR corpora.** ASR corpora are used to generate synthetic speech translation data. We only use the Common Voice (Ardila et al., 2020) and LibriSpeech (Panayotov et al., 2015) corpora. Furthermore, we filter the noisy training data in the Common Voice corpus by force decoding and keep 1 million utterances.

**MT corpora.** Machine translation (MT) corpora are used to translate the English transcription. We use the allowed English-German translation data from WMT 2020 (Barrault et al., 2020) and OpenSubtitles2018 (Lison and Tiedemann, 2016). We filter the training bilingual data followed Li et al. (2019), which includes length ratio, language detection, and so on.

<sup>1</sup>We only described the training data used in our system.

**ST corpora.** The ST corpora we used include MuST-C (Gangi et al., 2019) English-German<sup>2</sup>, CoVoST (Wang et al., 2020), Speech-Translation TED corpus<sup>3</sup>, and Europarl-ST (Iranzo-Sánchez et al., 2020).

The statistics of the final training data are shown in Table 1. We augment the quantity of the ST training data by translating the English transcription (the details are unveiled in Section 4).

Task	Corpora	Size	Time
ASR	LibriSpeech	281241	960h
	Common Voice	1000000	1387h
	Total	1281241	2347h
MT	CommonCrawl	2014304	-
	Europarl	1802849	-
	ParaCrawl	31528317	-
	Wiki	5714363	-
	OpenSubtitles	14449099	-
Total	55508932	-	
ST	MuST-C	249462	435h
	CoVoST	289411	329h
	ST TED	170133	254h
	Europarl	69537	155h
Total	778543	1173h	

Table 1: Data statistics of the ASR, MT, and ST corpora.

### 3 Model Architecture

In this section, we describe the baseline model and the architecture improvements. Then, the experimental results are shown to demonstrate the effectiveness.

#### 3.1 Baseline Model

Our system is based on deep Transformer (Vaswani et al., 2017) implemented on the fairseq toolkit (Ott et al., 2019). Furthermore, dynamic linear combination of layers (DLCL) (Wang et al., 2019) method is employed to train the deep model effectively (Li et al., 2020a,b).

To reduce the computational cost, the input speech features are processed by two convolutional layers, which have a stride of 2. This downsamples

<sup>2</sup>We use the latest MusST-C v2 dataset released by IWSLT 2021.

<sup>3</sup><http://i13pc106.ira.uka.de/~mmueller/iwslt-corpus.zip>

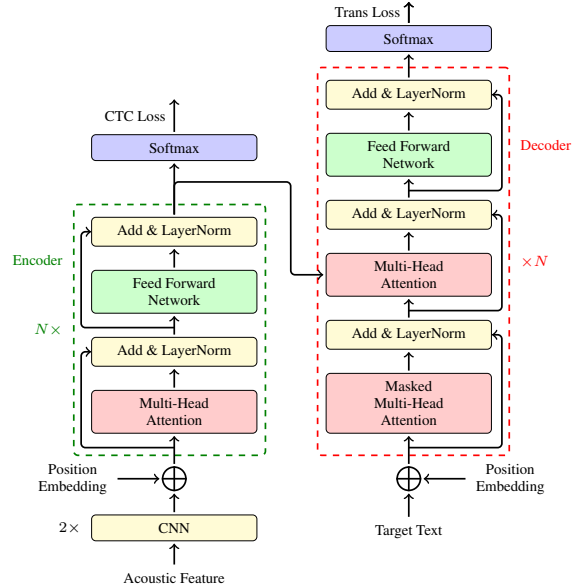


Figure 1: The baseline model architecture.

the sequence by a factor of 4 (Weiss et al., 2017). For strong systems, we use Connectionist Temporal Classification (CTC) (Graves et al., 2006) as the auxiliary loss on the encoders (Watanabe et al., 2017; Karita et al., 2019; Bahar et al., 2019). The weight of CTC objective  $\alpha$  is set to 0.3 for all ASR and ST models. The model architecture is shown in Figure 1<sup>4</sup>.

#### 3.2 Conformer

Conformer (Gulati et al., 2020) models both local and global dependencies by combining the Convolutional Neural Network and Transformers. It has shown superiority and achieved promising results in ASR tasks.

We replace the Transformer blocks in the encoder by the conformer blocks, which include two macaron-like feed-forward networks, multi-head self attention modules, and convolution modules. Note that we use the RPE proposed in Shaw et al. (2018) rather than Transformer-XL (Dai et al., 2019).

#### 3.3 Relative Position Encoding

Due to the non-sequential modeling of the original self attention modules, the vanilla Transformer employs the position embedding by a deterministic sinusoidal function to indicate the absolute position of each input element (Vaswani et al., 2017). However, this scheme is far from ideal for acoustic modeling (Pham et al., 2020).

<sup>4</sup><https://github.com/NiuTrans/MTBook>

Model	tst-COMMON
Baseline	23.98
+ Conformer	24.43
+ RPE	24.69
+ SATE	25.35

Table 2: Effects of the architecture improvements. We report SacreBLEU scores [%] on the MuST-C tst-COMMON set.

The latest work (Pham et al., 2020; Gulati et al., 2020) points out that the relative position encoding enables the model to generalize better for the unseen sequence lengths. It yields a significant improvement on the acoustic modeling tasks. We reimplement the relative position encoding scheme (Shaw et al., 2018). The maximum relative position is set to 100 for the encoder and 20 for the decoder. We use both absolute and relative positional representations simultaneously.

### 3.4 Stacked Acoustic and Textual Encoding

The previous work (Bahar et al., 2019) employs the CTC loss on the top layer of the encoder, which forces the encoders to learn soft alignments between speech and transcription. However, the CTC loss demonstrates strong preference for locally attentive models, which is inconsistent with the ST model (Xu et al., 2021).

In our systems, we use the stacked acoustic-and-textual encoding (SATE) (Xu et al., 2021) method to encode the speech features. It calculates the CTC loss based on the hidden states of the intermediate layer rather than the top layer. The layers below CTC also extract the acoustic representation like an ASR encoder, while the upper layers with no CTC encode the global representation for translation. An adaptor layer is introduced to bridge the acoustic and textual encoding.

### 3.5 Experimental Results

We use the architecture described in Section 3.1 as the baseline model. The encoder consists of 12 layers and the decoder consists of 6 layers. Each layer comprises 256 hidden units, 4 attention heads, and 2048 feed-forward size. The encoder of SATE includes an acoustic encoder of 8 layers and a textual encoder of 4 layers. The model is trained with MuST-C English-German dataset and we test the results on the tst-COMMON set based on the SacreBLEU. The other experimental details are shown

Data	Corpora	Size	Time
Synthetic	LibriSpeech	281241	960h
	Common Voice	1000000	1387h
	MuST-C	249462	435h
	ST TED	170133	254h
	Total	1700836	3036h
Real	Total	778543	1173h
	Total	2479379	4209h

Table 3: All available ST corpora.

in Section 6.

We report the experimental results after applying each architecture improvement in Table 2. Benefitting the power of the deep Transformer, our baseline model achieves 23.98 BLEU points. The Conformer and RPE methods strengthen the encoding and achieve an improvement of 0.45 and 0.26 BLEU points. SATE achieves a remarkable improvement by encoding the acoustic representation and textual representation respectively. We will explore better architecture designs in the future.

## 4 Data Augmentation

A large amount of the training data is necessary for a strong neural model. However, unlike the ASR and MT tasks, annotated speech-to-translation data is scarce, which prevents well-trained ST models. This is the main reason why cascaded systems are the dominant approach in the industrial scenarios. In this section, we describe our data augmentation method.

We train a deep DLCL Transformer (Wang et al., 2019) with the 25 encoder layers on all available MT data. To keep the domain consistency with the original ST data, we finetune the MT model on the MuST-C dataset. The model achieves the SacreBLEU of 35.89 of the MuST-C tst-COMMON test set. For the case-insensitive LibriSpeech dataset, we train a similar MT model except for lower-casing the source text without punctuation during training.

Then, we generate the German translation from English transcription in the LibriSpeech and Common Voice ASR datasets. Furthermore, sequence-level knowledge distillation (Kim and Rush, 2016) is applied to augment the training data. We generate the translation of the MuST-C and Speech-Translation TED ST datasets which are more re-

lated to the target domain.

Corrupting the acoustic feature is another data augmentation method, including SpecAugment, speed perturbation, and so on. SpecAugment (Park et al., 2019) is a simple data augmentation applied on the input acoustic features. The time masking and the frequency masking are applied in our systems. Speed perturbation transforms the audio by a speed rate, which changes the duration of the audio signal. Limited by the size of GPU resources, we do not use this method. Compared with the perturbed data, we think the synthetic samples improve the robustness more effectively. All available ST corpora are shown in Table 3.

## 5 Ensemble Decoding

Ensemble decoding is an effective method to improve performance by integrating the predictions from multiple models. It has been proved in the WMT competitions (Wang et al., 2018; Li et al., 2019). In our systems, we train multiple ST models with different training data for diverse ensemble decoding. The models are chosen based on the performance of the development set. This leads to a significant improvement over a single model.

## 6 Experiments

### 6.1 Preprocessing

We remove the utterances with more than 3000 frames or less than 5 frames. The 80-channel log-mel filterbank features are extracted from the audio file by torchaudio<sup>5</sup> library. We use the lower-cased transcriptions without punctuations for CTC loss computation. We learn SentencePiece<sup>6</sup> subword segmentation with a size of 10,000 based on a shared source and target vocabulary for all datasets.

### 6.2 Model Settings

All experiments are implemented based on the fairseq toolkit<sup>7</sup>. We use Adam optimizer and adopt the default learning schedule in fairseq. We apply dropout with a rate of 0.1 and label smoothing  $\epsilon_{ls} = 0.1$  for regularization. We also set the activate function dropout to 0.1 and attention dropout to 0.1, which improves the regularization and overcomes the overfitting.

We use the best model architecture that combines all the improvements described in Section

<sup>5</sup><https://github.com/pytorch/audio>

<sup>6</sup><https://github.com/google/sentencepiece>

<sup>7</sup><https://github.com/pytorch/fairseq>

3. The encoder includes an acoustic encoder of 12 conformer layers and a textual encoder of 6 transformer layers. The decoder consists of 6 Transformer layers. Each layer comprises 512 hidden units, 8 attention heads, and 2048 feed-forward size. Pre-norm is applied for training a deep model. The weight of CTC objective  $\alpha$  for multitask learning is set to 0.3 for all models. All the models are trained for 50 epochs on one machine with 8 NVIDIA 2080Ti GPUs.

During inference, we average the model parameters on the final 10 checkpoints. We use beam search with a beam size of 5 for all models. The coefficient of length normalization is tuned on the development set. We report the case-sensitive SacreBLEU (Post, 2018) on the MuST-C tst-COMMON set, IWSLT tst2019 and tst2020 test set.

The organizers provide the segmentation of the test sets and allow the participants to use the own segmentation. We simply use the segmentation provided by the WerRTCvad<sup>8</sup> toolkit.

### 6.3 Experimental Results

Firstly, We train the model on all training corpora, including real and synthetic speech-to-translation paired data. As shown in Table 4, we achieve a high BLEU on the tst-COMMON test set, but a low performance on the tst2019 test set compared with the previous work (Gaido et al., 2020). A possible reason is that the data distribution between IWSLT test sets and the synthetic data is different.

tst-COMMON	tst2019
32.65	14.16

Table 4: Performance of the model trained on all training corpora.

To verify this assumption, we pick some subsets from the available datasets for training, including MuST-C and ST TED from the real corpora and MuST-C and LibriSpeech from the synthetic corpora. We present the results in Table 5. Although the performance on the tst-COMMON test set drops by 0.8 BLEU points, the model achieves a reasonable performance on the tst2019 test set. Furthermore, we finetune the model on the MuST-C dataset with a small learning rate. This yields a slight improvement.

<sup>8</sup><https://github.com/wiseman/py-webrtcvad>

Model	tst-COMMON	tst2019
Base	31.85	20.64
+ finetune	32.31	20.73

Table 5: Performance of the model trained with the subsets of all available corpora.

Test sets	Given	Own
tst-COMMON	33.84	-
tst2019	22.68	23.76
tst2020	21.8	22.8
tst2021†	19.0	19.6
tst2021‡	20.7	20.6
tst2021★	30.7	30.3

Table 6: Final results with ensemble decoding. We report the results with given and own segmentation. There are two references on the tst2021 test set: TED reference (†) and IWSLT reference (‡). The final results are based on both references (★) together.

We train multiple models with different training data for diverse ensemble decoding. We select a part of the synthetic corpora randomly, then mix them with the whole real training data. Finally, we use the ensemble decoding with 6 models for the final results and achieve a substantial improvement over a single model. As shown in Table 6, we achieve an excellent performance of 33.84 BLEU points on the MuST-C En-De tst-COMMON set.

The best end-to-end system of last year achieves 20.1 BLEU points on the tst2019 test set and 21.49 BLEU points on the tst2020 test set with the given segmentation. We achieve remarkable improvements of 2.58 and 0.31 BLEU points, which demonstrates the superiority of our systems.

There are two references available for tst2021 test set. The TED reference is the original one from the TED website. Since new regulations for the official regulation lead to translations that are much shorter, they created a second reference translation, called the IWSLT reference. The final results are based on both references. We achieve better performance with the own segmentation on the TED reference, which is consistent with the results on the previous test sets. However, the results with the own segmentation are worse on the IWSLT reference. A possible reason is that we do not optimize the segmentation tool for IWSLT test sets. We will explore better segmentation methods in the future.

## 7 Conclusion

This paper describes the submission of the Niu-Trans E2E ST systems for the IWSLT 2021 offline task, which translates the English audio to German translation directly without intermediate transcription. We build our final submissions considering two mainstreams:

- Model architecture improvements for the speech translation task.
- Data augmentation by translating the English transcription to German translation.

We also find that the distribution of the training data has a great impact on the performance and alleviate it by ensemble decoding. Using the given segmentation, we achieve remarkable improvements over the best end-to-end system of last year.

## 8 Acknowledgement

This work was supported in part by the National Science Foundation of China (Nos. 61876035 and 61732005), the National Key R&D Program of China (No. 2019QY1801), and the Ministry of Science and Technology of the PRC (Nos. 2019YFF0303002 and 2020AAA0107900). The authors would like to thank anonymous reviewers for their comments.

## References

- Antonios Anastasopoulos, Ondřej Bojar, Jacob Bremerman, Roldano Cattoni, Maha Elbayad, 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.
- Rosana Ardila, Megan Branson, Kelly Davis, Michael Kohler, Josh Meyer, Michael Henretty, Reuben Morais, Lindsay Saunders, Francis M. Tyers, and Gregor Weber. 2020. [Common voice: A massively-multilingual speech corpus](#). In *Proceedings of The 12th Language Resources and Evaluation Conference, LREC 2020, Marseille, France, May 11-16, 2020*, pages 4218–4222. European Language Resources Association.

- Parnia Bahar, Tobias Bieschke, and Hermann Ney. 2019. [A comparative study on end-to-end speech to text translation](#). In *IEEE Automatic Speech Recognition and Understanding Workshop, ASRU 2019, Singapore, December 14-18, 2019*, pages 792–799. IEEE.
- Loïc Barrault, Magdalena Biesialska, Ondrej Bojar, Marta R. Costa-jussà, Christian Federmann, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Matthias Huck, Eric Joanis, Tom Kocmi, Philipp Koehn, Chi-kiu Lo, Nikola Ljubesic, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Santanu Pal, Matt Post, and Marcos Zampieri. 2020. [Findings of the 2020 conference on machine translation \(WMT20\)](#). In *Proceedings of the Fifth Conference on Machine Translation, WMT@EMNLP 2020, Online, November 19-20, 2020*, pages 1–55. Association for Computational Linguistics.
- Alexandre Berard, Olivier Pietquin, Christophe Servan, and Laurent Besacier. 2016. [Listen and translate: A proof of concept for end-to-end speech-to-text translation](#). *CoRR*, abs/1612.01744.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc Viet Le, and Ruslan Salakhutdinov. 2019. [Transformer-xl: Attentive language models beyond a fixed-length context](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 2978–2988. Association for Computational Linguistics.
- Long Duong, Antonios Anastasopoulos, David Chiang, Steven Bird, and Trevor Cohn. 2016. [An attentional model for speech translation without transcription](#). In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 949–959. The Association for Computational Linguistics.
- Marco Gaido, Mattia Antonino Di Gangi, Matteo Negri, and Marco Turchi. 2020. [End-to-end speech-translation with knowledge distillation: Fbk@iwslt2020](#). In *Proceedings of the 17th International Conference on Spoken Language Translation, IWSLT 2020, Online, July 9 - 10, 2020*, pages 80–88. Association for Computational Linguistics.
- Mattia Antonino Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019. [Must-c: a multilingual speech translation corpus](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 2012–2017. Association for Computational Linguistics.
- Alex Graves, Santiago Fernández, Faustino J. Gomez, and Jürgen Schmidhuber. 2006. [Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks](#). In *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006*, volume 148 of *ACM International Conference Proceeding Series*, pages 369–376. ACM.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. 2020. [Conformer: Convolution-augmented transformer for speech recognition](#). In *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, pages 5036–5040. ISCA.
- Javier Iranzo-Sánchez, Joan Albert Silvestre-Cerdà, Javier Jorge, Nahuel Roselló, Adrià Giménez, Albert Sanchís, Jorge Civera, and Alfons Juan. 2020. [Europarl-st: A multilingual corpus for speech translation of parliamentary debates](#). In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, pages 8229–8233. IEEE.
- Shigeki Karita, Nelson Enrique Yalta Soplín, Shinji Watanabe, Marc Delcroix, Atsunori Ogawa, and Tomohiro Nakatani. 2019. [Improving transformer-based end-to-end speech recognition with connectionist temporal classification and language model integration](#). In *Interspeech*

- 2019, *20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, pages 1408–1412. ISCA.
- Yoon Kim and Alexander M. Rush. 2016. [Sequence-level knowledge distillation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 1317–1327. The Association for Computational Linguistics.
- Bei Li, Yinqiao Li, Chen Xu, Ye Lin, Jiqiang Liu, Hui Liu, Ziyang Wang, Yuhao Zhang, Nuo Xu, Zeyang Wang, Kai Feng, Hexuan Chen, Tengbo Liu, Yanyang Li, Qiang Wang, Tong Xiao, and Jingbo Zhu. 2019. [The niutrans machine translation systems for WMT19](#). In *Proceedings of the Fourth Conference on Machine Translation, WMT 2019, Florence, Italy, August 1-2, 2019 - Volume 2: Shared Task Papers, Day 1*, pages 257–266. Association for Computational Linguistics.
- Bei Li, Ziyang Wang, Hui Liu, Quan Du, Tong Xiao, Chunliang Zhang, and Jingbo Zhu. 2020a. [Learning light-weight translation models from deep transformer](#). *CoRR*, abs/2012.13866.
- Bei Li, Ziyang Wang, Hui Liu, Yufan Jiang, Quan Du, Tong Xiao, Huizhen Wang, and Jingbo Zhu. 2020b. [Shallow-to-deep training for neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 995–1005. Association for Computational Linguistics.
- Pierre Lison and Jörg Tiedemann. 2016. [Open-subtitles2016: Extracting large parallel corpora from movie and TV subtitles](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC 2016, Portorož, Slovenia, May 23-28, 2016*. European Language Resources Association (ELRA).
- 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: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Demonstrations*, pages 48–53. 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 2015, South Brisbane, Queensland, Australia, April 19-24, 2015*, pages 5206–5210. IEEE.
- Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le. 2019. [SpecAugment: A simple data augmentation method for automatic speech recognition](#). In *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, pages 2613–2617. ISCA.
- Ngoc-Quan Pham, Thanh-Le Ha, Tuan-Nam Nguyen, Thai-Son Nguyen, Elizabeth Salesky, Sebastian Stüker, Jan Niehues, and Alex Waibel. 2020. [Relative positional encoding for speech recognition and direct translation](#). In *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, pages 31–35. ISCA.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, pages 186–191. Association for Computational Linguistics.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-attention with relative position representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 2 (Short Papers)*, pages 464–468. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing*

*Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.

Changhan Wang, Anne Wu, and Juan Pino. 2020. [Covost 2: A massively multilingual speech-to-text translation corpus](#). *CoRR*, abs/2007.10310.

Qiang Wang, Bei Li, Jiqiang Liu, Bojian Jiang, Zheyang Zhang, Yinqiao Li, Ye Lin, Tong Xiao, and Jingbo Zhu. 2018. [The niutrans machine translation system for WMT18](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, pages 528–534. Association for Computational Linguistics.

Qiang Wang, Bei Li, Tong Xiao, Jingbo Zhu, Changliang Li, Derek F. Wong, and Lidia S. Chao. 2019. [Learning deep transformer models for machine translation](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 1810–1822. Association for Computational Linguistics.

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

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

Chen Xu, Bojie Hu, Yanyang Li, Yuhao Zhang, Shen Huang, Qi Ju, Tong Xiao, and Jingbo Zhu. 2021. [Stacked acoustic-and-textual encoding: Integrating the pre-trained models into speech translation encoders](#). *CoRR*, abs/2105.05752.