

CTC-based Non-autoregressive Textless Speech-to-Speech Translation

Qingkai Fang^{1,3}, Zhengrui Ma^{1,3}, Yan Zhou^{1,3}, Min Zhang⁴, Yang Feng^{1,2,3*}

¹Key Laboratory of Intelligent Information Processing

Institute of Computing Technology, Chinese Academy of Sciences (ICT/CAS)

²Key Laboratory of AI Safety, Chinese Academy of Sciences

³University of Chinese Academy of Sciences, Beijing, China

⁴School of Future Science and Engineering, Soochow University

{fangqingkai21b, fengyang}@ict.ac.cn zhangminmt@hotmail.com

Abstract

Direct speech-to-speech translation (S2ST) has achieved impressive translation quality, but it often faces the challenge of slow decoding due to the considerable length of speech sequences. Recently, some research has turned to non-autoregressive (NAR) models to expedite decoding, yet the translation quality typically lags behind autoregressive (AR) models significantly. In this paper, we investigate the performance of CTC-based NAR models in S2ST, as these models have shown impressive results in machine translation. Experimental results demonstrate that by combining pretraining, knowledge distillation, and advanced NAR training techniques such as glancing training and non-monotonic latent alignments, CTC-based NAR models achieve translation quality comparable to the AR model, while preserving up to $26.81\times$ decoding speedup.¹

1 Introduction

Direct speech-to-speech translation (S2ST) refers to the process of generating target speech directly from the source speech, without the need for generating intermediate source or target text. Distinguished from traditional cascaded S2ST (Lavie et al., 1997; Nakamura et al., 2006), which involves cascading automatic speech recognition (ASR), machine translation (MT), and text-to-speech (TTS) models, direct S2ST can mitigate error accumulation and may be applicable to low-resource languages without written forms (Chen et al., 2022).

Early direct S2ST models often require text supervision during training (Jia et al., 2019), which is not feasible for numerous unwritten languages. To address this issue, researchers have proposed textless S2ST based on discrete units (Lee et al., 2022a,b), which does not rely on any text supervision during training. Specifically, they first extract

the discrete representation (i.e., discrete units) of the target speech based on the speech pretraining model HuBERT (Hsu et al., 2021). Then, they train a speech-to-unit translation (S2UT) model to generate target discrete units from the source speech and subsequently employ a unit-based vocoder to synthesize the target speech. However, due to the considerable length of discrete unit sequences, these models often exhibit higher decoding latency. To address this problem, Huang et al. (2023) propose the use of non-autoregressive (NAR) model to generate discrete units in parallel. Despite its faster decoding speed, it suffers from a significant decrease in translation quality compared to S2UT limited by the expressive ability of NAR models.

In this paper, we aim to explore the performance of more powerful NAR models in S2ST. We opt for models based on connectionist temporal classification (CTC; Graves et al., 2006), which improves the model’s expressive ability by expanding the output space. Previous work has demonstrated that CTC-based NAR models exhibit excellent performance in machine translation (Libovický and Helcl, 2018; Yan et al., 2023) and speech-to-text translation (Xu et al., 2023). In this paper, we build upon CTC and incorporate multiple techniques including pretraining, knowledge distillation, and advanced NAR training techniques such as glancing training and non-monotonic latent alignments. Experimental results show that CTC-based models can attain translation quality comparable to the AR model, while maintaining up to $26.81\times$ decoding speedup.

2 Background

2.1 Speech-to-unit Translation

Speech-to-unit translation (S2UT; Lee et al., 2022a) introduces *discrete units* as training targets and has been proven to be one of the most effective approaches in textless S2ST. Specifically, discrete units are obtained by applying K-means clustering

*Corresponding author: Yang Feng.

¹Code: <https://github.com/ictnlp/CTC-S2UT>.

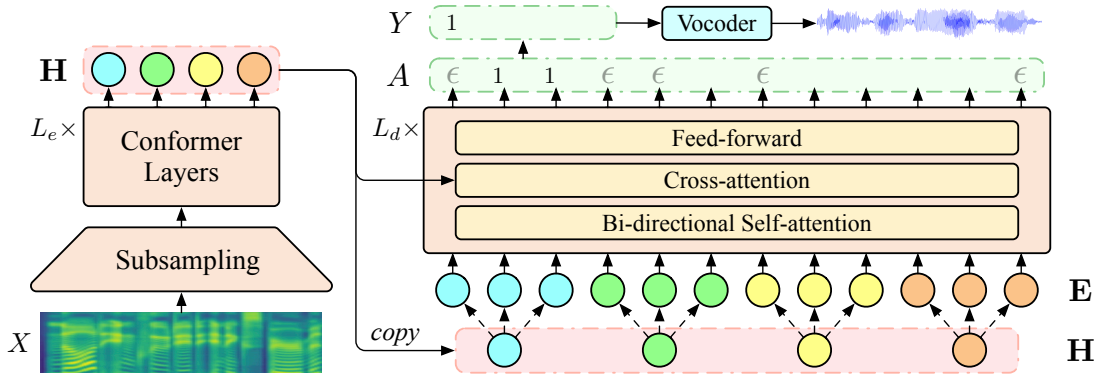


Figure 1: Model architecture of CTC-S2UT.

to the 50Hz continuous representation generated by the speech pretraining model HuBERT (Hsu et al., 2021). The cluster indices serve as the discrete units, ensuring that each discrete unit is an integer in the range of 0 to $K - 1$, where K represents the number of clusters.

S2UT trains a sequence-to-sequence model to directly generate the discrete unit sequence corresponding to the target speech from the source speech. Finally, a pre-trained unit-based HiFi-GAN (Kong et al., 2020) vocoder synthesizes the target speech waveform from the sequence of discrete units, completing the S2ST task.

2.2 Connectionist Temporal Classification

Connectionist temporal classification (CTC; Graves et al., 2006) is a sequence modeling algorithm used to model *variable-length* mappings between sequences, originally applied to speech recognition (Graves and Jaitly, 2014) and more recently extended to machine translation (Libovický and Helcl, 2018). Formally, given an input sequence $X = (x_1, \dots, x_N)$ and an output sequence $Y = (y_1, \dots, y_M)$, CTC introduces a special blank token ϵ into the output space $\mathcal{Y}^* = \mathcal{Y} \cup \{\epsilon\}$. It first generates a sequence $A = (a_1, \dots, a_T)$ based on the input sequence X , where T is a pre-determined length, and A is referred to as the *alignment*. Furthermore, CTC defines a collapsing function $\beta(A)$, which first merges consecutive identical tokens in A into one and then removes all blank tokens to obtain the final output sequence Y . During the training process, CTC marginalizes over all possible alignments for optimization:

$$P(Y|X) = \sum_{A \in \beta^{-1}(Y)} P(A|X), \quad (1)$$

where $\beta^{-1}(Y)$ contains all possible alignments of length T that can be collapsed to Y . The probability $P(A|X)$ is modeled in a non-autoregressive way:

$$P(A|X) = \prod_{t=1}^T p(a_t|X). \quad (2)$$

3 Method

In this section, we investigate the performance of CTC-based non-autoregressive translation models in S2ST. We develop our model based on S2UT, named CTC-S2UT, where X represents the source speech, and Y represents the sequence of discrete units corresponding to the target speech. Section 3.1 will introduce the model architecture, and Section 3.2 will introduce the training approach.

3.1 Model Architecture

CTC-S2UT comprises a speech encoder and a non-autoregressive unit decoder. The speech encoder encodes the source speech, while the unit decoder generates the target discrete units in parallel.

Speech Encoder The speech encoder comprises a subsampling module and L_e Conformer layers (Gulati et al., 2020). The subsampling module consists of two 1D convolutional layers, responsible for a $4 \times$ subsampling of the input speech features. Conformer layers employ both attention mechanisms and convolutional layers to concurrently model global and local features of the speech. Relative positional encoding (Dai et al., 2019) is employed in the multi-head attention. The output representations of the speech encoder are denoted as $\mathbf{H} = (\mathbf{h}_1, \dots, \mathbf{h}_{N'})$, where $N' = \lfloor N/4 \rfloor$.

Unit Decoder The unit decoder consists of L_d non-autoregressive Transformer decoder layers (Gu

et al., 2018), each comprising a bidirectional self-attention layer, a cross-attention layer, and a feed-forward layer. The input to the decoder is obtained by uniformly upsampling the encoder output representations by λ times: $\mathbf{E} = (\mathbf{e}_1, \dots, \mathbf{e}_T)$, where $T = \lambda \cdot N'$ and $\mathbf{e}_i = \mathbf{h}_{\lfloor i/\lambda \rfloor}$. We add learnable positional encodings to \mathbf{E} and feed it into the decoder. The decoder generates the alignment A in parallel, and finally obtains the sequence of target discrete units $Y = \beta(A)$.

3.2 Training

In theory, we can calculate the conditional probability $P(Y|X)$ using Eq. (1) and train the model by minimizing $-\log P(Y|X)$. However, due to the complexity of the speech-to-unit translation task, training from scratch using a simple CTC objective often faces challenges. Therefore, we adopt several techniques to facilitate training, including encoder pretraining, knowledge distillation, glancing training and non-monotonic latent alignments.

Encoder Pretraining We initialize the encoder of CTC-S2UT using the encoder from a pretrained autoregressive S2UT model to provide the encoder with a good initial state.

Knowledge Distillation We perform sequence-level knowledge distillation (KD; Kim and Rush, 2016) with the autoregressive S2UT model, which is a widely used approach in non-autoregressive translation to reduce data multimodality (Zhou et al., 2020; Shao et al., 2022).

Glancing Training Glancing training (GLAT; Qian et al., 2021) has been proven to significantly enhance the translation quality of NAR translation models. It first selects the alignment with the highest posterior probability: $A^* = \arg \max_{A \in \beta^{-1}(Y)} P(A|X)$, and then randomly chooses some positions in the decoder input \mathbf{E} to replace with the corresponding token embeddings from the alignment A^* .

Non-monotonic Latent Alignments Shao and Feng (2022) introduces non-monotonic latent alignments (NMLA) to enhance CTC-based NAR models. It maximizes the F1 score of expected bigram matching between alignments and the target. More details can be found in the original paper.

Direction	Source Hours	Target Hours	#Samples
Fr→En	264h	174h	207K
En→Fr	174h	264h	207K
En→Es	70h	113h	79K

Table 1: Statistics of datasets.

4 Experiments

4.1 Experimental Setups

Dataset We use the CVSS-C dataset (Jia et al., 2022b) for experiments, which includes speech-to-speech translation pairs between 21 languages and English. Following Huang et al. (2023), we conduct experiments on three language pairs, including French→English (Fr→En), English→French (En→Fr), and English→Spanish (En→Es). Table 1 lists the detailed statistics of the dataset. We compute the 80-dimensional mel-filterbank features for the source speech, and extract the discrete unit sequences for the target speech using the publicly available pretrained mHuBERT model and K-means quantizer².

Model Configurations and Training For both autoregressive S2UT and CTC-S2UT, the models consist of 12 Conformer encoder layers and 6 Transformer decoder layers. The dropout is set to 0.3. We initialize the encoder of CTC-S2UT with the encoder from S2UT and train the model in two stages. In the first stage, we employ conventional CTC loss for training. Each batch contains 320k source audio frames, and the learning rate warms up to 1e-3 over 10k steps. The glancing ratio linearly decays from 0.5 to 0.3 within 100k steps. In the second stage, we fine-tune the model using NMLA loss for 6k steps. Each batch contains 1,280k source audio frames, and the learning rate warms up to 3e-4 over 500 steps. The glancing ratio is fixed at 0.3. The upsample factor λ is set to 2, 6, and 6 for Fr→En, En→Fr, and En→Es, respectively. We use Adam (Kingma and Ba, 2015) optimizer in all training stages. All models are trained on 4 RTX 3090 GPUs.

Evaluation and Baseline Systems We use the ASR-BLEU³ toolkit to evaluate the translation quality. It transcribes the translated speech into

²https://github.com/facebookresearch/fairseq/blob/main/examples/speech_to_speech/docs/textless_s2st_real_data.md

³https://github.com/facebookresearch/fairseq/tree/ust/examples/speech_to_speech/asr_bleu

Model	#Iter	#Param	Textless	ASR-BLEU				Speedup
				Fr→En	En→Fr	En→Es	Avg.	
S2UT	M	58M	✓	24.80	21.38	21.71	22.63	1.00×
TranSpeech	5	67M	✓	17.24	16.30	11.79	15.11	11.04×
TranSpeech	15	67M	✓	18.03	16.97	12.62	15.87	5.34×
TranSpeech (b=15)	15	67M	✓	18.10	17.05	12.70	15.95	2.75×
TranSpeech (b=15 + NPD)	15	67M	✓	18.39	17.50	12.77	16.22	2.53×
DASpeech (Lookahead)	1+1	93M	×	24.71	N/A	N/A	N/A	18.53×
DASpeech (Joint-Viterbi)	1+1	93M	×	25.03	N/A	N/A	N/A	16.29×
CTC-S2UT	1	59M	✓	25.16	20.85	21.60	22.54	26.81×

Table 2: Results on CVSS-C test sets. Results of TranSpeech (Huang et al., 2023) and DASpeech (Fang et al., 2023) are quoted from the original paper. NPD: noisy parallel decoding; b: length beam in NAR decoding. The best results among NAR models are marked in **bold**.

Pretrain	GLAT	KD	NMLA	ASR-BLEU
×	✓	✓	×	0.00
✓	×	✓	×	11.44
✓	✓	×	×	23.19
✓	✓	✓	×	23.85
✓	✓	✓	✓	25.16

Table 3: Results on CVSS-C Fr→En test set with different combinations of training techniques.

text using a pretrained ASR model, and calculates the BLEU score (Papineni et al., 2002). We measure the decoding speedup on the CVSS-C Fr→En test set using 1 RTX 3090 GPU with a batch size of 1. Besides autoregressive S2UT, we include TranSpeech (Huang et al., 2023) and DASpeech (Fang et al., 2023) as NAR baseline systems for comparison.

4.2 Main Results

Table 2 presents the experimental results on the CVSS-C test set, where CTC-S2UT achieves translation quality comparable to autoregressive S2UT while preserving a significant decoding speedup of 26.81 times. Compared to the previous NAR S2UT model TranSpeech, which relies on CMLM (Ghazvininejad et al., 2019), CTC-S2UT attains higher translation quality without the need for iterative decoding, demonstrating the importance of employing more powerful NAR models. Compared to the previous two-pass NAR model DASpeech, CTC-S2UT achieves slightly higher translation quality. Additionally, since CTC-S2UT does not require decoding the target text, it exhibits faster decoding speed and can be applied to target languages without written form.

#Frames	[0, 300)	[300, 600)	[600, +∞)
Speedup	11.67×	24.95×	34.28×

Table 4: Speedup of CTC-S2UT compared with S2UT across different source speech lengths (#Frames) on CVSS-C Fr→En test set.

4.3 Ablation Study

Table 3 shows the results with different combinations of training techniques. We first explore the influence of encoder pretraining, GLAT and KD on training. The results indicate that encoder pretraining is indispensable, as the model fails to converge without it. GLAT also significantly impacts performance, while KD slightly enhances model performance. Building upon these techniques, fine-tuning the model using NMLA further improves performance, confirming the effectiveness of four training techniques we employed.

4.4 Speedup Analysis

We further investigate the decoding speedup ratio of CTC-S2UT compared to autoregressive S2UT for different lengths of speech. As shown in Table 4, for longer source speech, the speedup ratio of CTC-S2UT becomes more significant, reaching up to 34.28 times, indicating the speed advantage of NAR models in translating long speech segments.

5 Related Work

Speech-to-speech translation extends speech-to-text translation (S2TT) (Fang et al., 2022; Fang and Feng, 2023a,b; Zhou et al., 2023) that further synthesizes the target speech. Jia et al. (2019) first introduces direct S2ST. Lee et al. (2022a,b); Chen et al. (2022) use discrete units as training

targets and successfully apply S2UT to unwritten languages. Jia et al. (2022a); Inaguma et al. (2023) propose two-pass S2ST models to improve the translation quality, which generate target text and target speech successively. Zhang et al. (2024) introduces a unified framework for simultaneous speech recognition, translation and synthesis. To improve the decoding efficiency, Inaguma et al. (2021); Xu et al. (2023) propose performing the S2TT task using non-autoregressive models. Huang et al. (2023) first proposes NAR S2UT, achieving higher decoding speed but lower translation quality. Wu (2023); Zhu et al. (2023) adopt diffusion models to iteratively generate discrete units. Fang et al. (2023) proposes the first two-pass NAR S2ST model, which achieves both good translation quality and fast decoding speed, but requires target text for training. Fang et al. (2024) introduces a composite S2ST model which can combine existing speech-to-text translation and TTS models. Ma et al. (2024) proposes a CTC-based NAR model for simultaneous S2ST. Compared with previous models, CTC-S2UT achieves better translation quality and faster decoding speed, and does not require any text supervision during training.

6 Conclusion

In this paper, we propose CTC-S2UT, an NAR textless S2ST model based on CTC. By incorporating several training techniques including encoder pre-training, knowledge distillation, glancing training, and non-monotonic latent alignment, CTC-S2UT achieves comparable translation quality to the autoregressive S2UT model, while achieving $26.81\times$ decoding speedup. Further ablation studies validate the effectiveness of each training technique.

Limitations

Although CTC-S2UT achieves good performance in translation quality and decoding speed, it still has some limitations: (1) Training CTC-S2UT requires pretraining an autoregressive S2UT model, as its success relies on S2UT for encoder pretraining and knowledge distillation. This makes the training process more complex compared to S2UT; (2) Due to the limitations of existing publicly available datasets, we did not conduct experiments on real-world unwritten languages but instead simulated textless scenarios on language pairs such as Fr \rightarrow En, En \rightarrow Fr, and En \rightarrow Es.

Acknowledgement

We thank all the anonymous reviewers for their insightful and valuable comments. This paper is supported by National Natural Science Foundation of China (Grant No.62376260).

References

- Peng-Jen Chen, Kevin Tran, Yilin Yang, Jingfei Du, Justine Kao, Yu-An Chung, Paden Tomasello, Paul-Ambroise Duquenne, Holger Schwenk, Hongyu Gong, Hirofumi Inaguma, Sravya Popuri, Changhan Wang, Juan Miguel Pino, Wei-Ning Hsu, and Ann Lee. 2022. *Speech-to-speech translation for A real-world unwritten language*. *CoRR*, abs/2211.06474.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. *Transformer-XL: Attentive language models beyond a fixed-length context*. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.
- Qingkai Fang and Yang Feng. 2023a. Back translation for speech-to-text translation without transcripts. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*.
- Qingkai Fang and Yang Feng. 2023b. Understanding and bridging the modality gap for speech translation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*.
- Qingkai Fang, Rong Ye, Lei Li, Yang Feng, and Mingxuan Wang. 2022. Stemm: Self-learning with speech-text manifold mixup for speech translation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*.
- Qingkai Fang, Shaolei Zhang, Zhengrui Ma, Min Zhang, and Yang Feng. 2024. Can we achieve high-quality direct speech-to-speech translation without parallel speech data? In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*.
- Qingkai Fang, Yan Zhou, and Yang Feng. 2023. DASpeech: Directed acyclic transformer for fast and high-quality speech-to-speech translation. In *Advances in Neural Information Processing Systems*.
- Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. 2019. *Mask-predict: Parallel decoding of conditional masked language models*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6112–6121, Hong Kong, China. Association for Computational Linguistics.

- Alex Graves, Santiago Fernández, Faustino 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 '06*, page 369–376, New York, NY, USA. Association for Computing Machinery.
- Alex Graves and Navdeep Jaitly. 2014. [Towards end-to-end speech recognition with recurrent neural networks](#). In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1764–1772, Beijing, China. PMLR.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor O.K. Li, and Richard Socher. 2018. [Non-autoregressive neural machine translation](#). In *International Conference on Learning Representations*.
- 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 *Proc. Interspeech 2020*, pages 5036–5040.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. [Hubert: Self-supervised speech representation learning by masked prediction of hidden units](#). *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, 29:3451–3460.
- Rongjie Huang, Jinglin Liu, Huadai Liu, Yi Ren, Lichao Zhang, Jinzheng He, and Zhou Zhao. 2023. [Transpeech: Speech-to-speech translation with bilateral perturbation](#). In *The Eleventh International Conference on Learning Representations*.
- Hirofumi Inaguma, Yosuke Higuchi, Kevin Duh, Tatsuya Kawahara, and Shinji Watanabe. 2021. Non-autoregressive end-to-end speech translation with parallel autoregressive rescoring. *arXiv preprint arXiv:2109.04411*.
- Hirofumi Inaguma, Sravya Popuri, Iliia Kulikov, Peng-Jen Chen, Changhan Wang, Yu-An Chung, Yun Tang, Ann Lee, Shinji Watanabe, and Juan Pino. 2023. [UnitY: Two-pass direct speech-to-speech translation with discrete units](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15655–15680, Toronto, Canada. Association for Computational Linguistics.
- Ye Jia, Michelle Tadmor Ramanovich, Tal Remez, and Roi Pomerantz. 2022a. [Translatotron 2: High-quality direct speech-to-speech translation with voice preservation](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 10120–10134. PMLR.
- Ye Jia, Michelle Tadmor Ramanovich, Quan Wang, and Heiga Zen. 2022b. CVSS corpus and massively multilingual speech-to-speech translation. In *Proceedings of Language Resources and Evaluation Conference (LREC)*, pages 6691–6703.
- Ye Jia, Ron J. Weiss, Fadi Biadsy, Wolfgang Macherey, Melvin Johnson, Zhifeng Chen, and Yonghui Wu. 2019. [Direct speech-to-speech translation with a sequence-to-sequence model](#). In *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, pages 1123–1127. 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*, 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*.
- Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. 2020. [Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis](#). In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS'20*, Red Hook, NY, USA. Curran Associates Inc.
- A. Lavie, A. Waibel, L. Levin, M. Finke, D. Gates, M. Gavalda, T. Zeppenfeld, and Puming Zhan. 1997. [Janus-iii: speech-to-speech translation in multiple languages](#). In *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 1, pages 99–102 vol.1.
- Ann Lee, Peng-Jen Chen, Changhan Wang, Jiatao Gu, Sravya Popuri, Xutai Ma, Adam Polyak, Yossi Adi, Qing He, Yun Tang, Juan Pino, and Wei-Ning Hsu. 2022a. [Direct speech-to-speech translation with discrete units](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3327–3339, Dublin, Ireland. Association for Computational Linguistics.
- Ann Lee, Hongyu Gong, Paul-Ambroise Duquenne, Holger Schwenk, Peng-Jen Chen, Changhan Wang, Sravya Popuri, Yossi Adi, Juan Pino, Jiatao Gu, and Wei-Ning Hsu. 2022b. [Textless speech-to-speech translation on real data](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 860–872, Seattle, United States. Association for Computational Linguistics.
- Jindřich Libovický and Jindřich Helcl. 2018. [End-to-end non-autoregressive neural machine translation with connectionist temporal classification](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3016–

- 3021, Brussels, Belgium. Association for Computational Linguistics.
- Zhengru Ma, Qingkai Fang, Shaolei Zhang, Shoutao Guo, Yang Feng, and Min Zhang. 2024. A non-autoregressive generation framework for end-to-end simultaneous speech-to-any translation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*.
- S. Nakamura, K. Markov, H. Nakaiwa, G. Kikui, H. Kawai, T. Jitsuhiro, J.-S. Zhang, H. Yamamoto, E. Sumita, and S. Yamamoto. 2006. [The atr multi-lingual speech-to-speech translation system](#). *IEEE Transactions on Audio, Speech, and Language Processing*, 14(2):365–376.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Lihua Qian, Hao Zhou, Yu Bao, Mingxuan Wang, Lin Qiu, Weinan Zhang, Yong Yu, and Lei Li. 2021. [Glancing transformer for non-autoregressive neural machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1993–2003, Online. Association for Computational Linguistics.
- Chenze Shao and Yang Feng. 2022. Non-monotonic latent alignments for ctc-based non-autoregressive machine translation. In *Proceedings of NeurIPS 2022*.
- Chenze Shao, Xuanfu Wu, and Yang Feng. 2022. One reference is not enough: Diverse distillation with reference selection for non-autoregressive translation. In *Proceedings of NAACL 2022*.
- Xianchao Wu. 2023. [Duplex diffusion models improve speech-to-speech translation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8035–8047, Toronto, Canada. Association for Computational Linguistics.
- Chen Xu, Xiaoqian Liu, Xiaowen Liu, Qingxuan Sun, Yuhao Zhang, Murun Yang, Qianqian Dong, Tom Ko, Mingxuan Wang, Tong Xiao, Anxiang Ma, and Jingbo Zhu. 2023. [CTC-based non-autoregressive speech translation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13321–13339, Toronto, Canada. Association for Computational Linguistics.
- Brian Yan, Siddharth Dalmia, Yosuke Higuchi, Graham Neubig, Florian Metze, Alan W Black, and Shinji Watanabe. 2023. [CTC alignments improve autoregressive translation](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1623–1639, Dubrovnik, Croatia. Association for Computational Linguistics.
- Shaolei Zhang, Qingkai Fang, Shoutao Guo, Zhengru Ma, Min Zhang, and Yang Feng. 2024. [Stream-Speech: Simultaneous speech-to-speech translation with multi-task learning](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*.
- Chunting Zhou, Jiatao Gu, and Graham Neubig. 2020. [Understanding knowledge distillation in non-autoregressive machine translation](#). In *International Conference on Learning Representations*.
- Yan Zhou, Qingkai Fang, and Yang Feng. 2023. [CMOT: Cross-modal mixup via optimal transport for speech translation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7873–7887, Toronto, Canada. Association for Computational Linguistics.
- Yongxin Zhu, Zhujin Gao, Xinyuan Zhou, Ye Zhongyi, and Linli Xu. 2023. [DiffS2UT: A semantic preserving diffusion model for textless direct speech-to-speech translation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11573–11583, Singapore. Association for Computational Linguistics.