



# Voice Transformer Network: Sequence-to-Sequence Voice Conversion Using Transformer with Text-to-Speech Pretraining

Wen-Chin Huang<sup>1</sup>, Tomoki Hayashi<sup>1</sup>, Yi-Chiao Wu<sup>1</sup>, Hirokazu Kameoka<sup>2</sup>, Tomoki Toda<sup>1</sup>

<sup>1</sup>Nagoya University, Japan

<sup>2</sup>NTT Communication Science Laboratories, Japan

wen.chinhuang@g.sp.m.is.nagoya-u.ac.jp

## Abstract

We introduce a novel sequence-to-sequence (seq2seq) voice conversion (VC) model based on the Transformer architecture with text-to-speech (TTS) pretraining. Seq2seq VC models are attractive owing to their ability to convert prosody. While seq2seq models based on recurrent neural networks (RNNs) and convolutional neural networks (CNNs) have been successfully applied to VC, the use of the Transformer network, which has shown promising results in various speech processing tasks, has not yet been investigated. Nonetheless, their data-hungry property and the mispronunciation of converted speech make seq2seq models far from practical. To this end, we propose a simple yet effective pretraining technique to transfer knowledge from learned TTS models, which benefit from large-scale, easily accessible TTS corpora. VC models initialized with such pretrained model parameters are able to generate effective hidden representations for high-fidelity, highly intelligible converted speech. Experimental results show that such a pretraining scheme can facilitate data-efficient training and outperform an RNN-based seq2seq VC model in terms of intelligibility, naturalness, and similarity.

**Index Terms:** Voice Conversion, Sequence-to-Sequence Learning, Transformer, pretraining

## 1. Introduction

Voice conversion (VC) aims to convert the speech from a source to that of a target without changing the linguistic content [1]. Conventional VC systems follow an *analysis—conversion—synthesis* paradigm [2]. First, a high quality vocoder such as WORLD [3] or STRAIGHT [4] is utilized to extract different acoustic features, such as spectral features and fundamental frequency (F0). These features are converted separately, and a waveform synthesizer finally generates the converted waveform using the converted features. Past VC studies have focused on the conversion of spectral features while only applying a simple linear transformation to F0. In addition, the conversion is usually performed frame-by-frame, i.e., the converted speech and the source speech are always of the same length. To summarize, the conversion of prosody, including F0 and duration, is overly simplified in the current VC literature.

This is where sequence-to-sequence (seq2seq) models [5] can play a role. Modern seq2seq models, often equipped with an attention mechanism [6, 7] to implicitly learn the alignment between the source and output sequences, can generate outputs of various lengths. This ability makes the seq2seq model a natural choice to convert duration in VC. In addition, the F0 contour can also be converted by considering F0 implicitly using mel spectrograms as the feature representation.

Existing seq2seq VC models are based on either recurrent neural networks (RNNs) [8, 9, 10, 11, 12, 13, 14, 15] or convolutional neural networks (CNNs) [16]. In recent years, the

Transformer architecture [17] has been shown to perform efficiently [18] in various speech processing tasks such as automatic speech recognition (ASR) [19], speech translation (ST) [20, 21], and text-to-speech (TTS) [22]. On the basis of attention mechanism solely, the Transformer enables parallel training by avoiding the use of recurrent layers, and provides a receptive field that spans the entire input by using multi-head self-attention rather than convolutional layers. Nonetheless, the above-mentioned speech applications that have successfully utilized the Transformer architecture all attempted to find a mapping between text and acoustic feature sequences. VC, in contrast, attempts to map between acoustic frames, whose high time resolution introduces challenges regarding computational memory cost and accurate attention learning.

Despite the promising results, seq2seq VC models suffer from two major problems. First, seq2seq models usually require a large amount of training data, although a large-scale parallel corpus, i.e., pairs of speech samples with identical linguistic contents uttered by both source and target speakers, is impractical to collect. Second, as pointed out in [11], the converted speech often suffers from mispronunciations and other instability problems such as phonemes and skipped phonemes. Several techniques have been proposed to address these issues. Some used external resources such as text labels [11, 12], extra data [15] or pretrained modules such as ASR [10, 13], while others designed extra regularization objectives [9, 16].

In this work, we propose a novel yet simple pretraining technique to transfer knowledge from learned TTS models to seq2seq VC training. We pretrain the encoder and decoder in separate steps to transfer the core abilities, which are to generate and utilize fine representations. First, the decoder is pretrained with a large-scale TTS corpus and a conventional TTS objective. This ensures a well-trained decoder that can generate high-quality speech with the correct hidden representations. Then, we train the encoder in an autoencoder style with the pretrained decoder fixed, so that the encoder is forced to encode input speech into hidden representations that can be recognized by the decoder. We demonstrate that the VC model initialized with the pretrained model parameters can generate high-quality, highly intelligible speech even with very limited training data.

Our contributions in this work are as follows:

- We apply the Transformer network to VC. To our knowledge, this is the first work to investigate this combination.
- We propose a TTS pretraining technique for VC. The pretraining process provides a prior for fast, sample-efficient VC model learning, thus reducing the data size requirement and training time. In this work, we verify the effectiveness of this scheme by transferring knowledge from Transformer-based TTS models to a Transformer-based VC model.

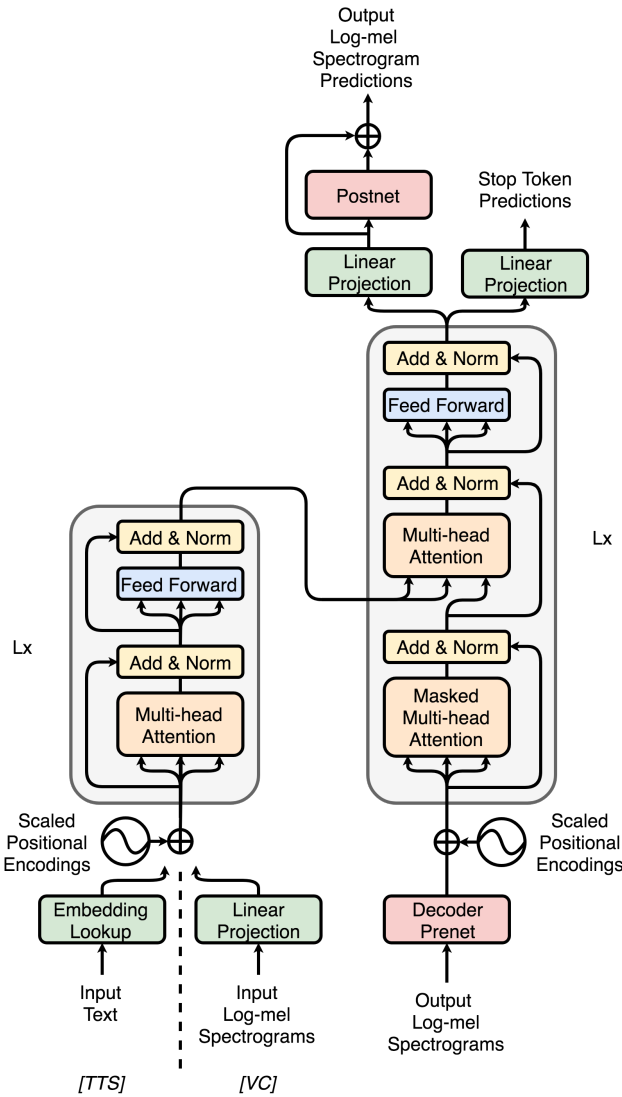


Figure 1: Model architecture of Transformer-TTS and VC.

## 2. Background

### 2.1. Sequence-to-sequence speech synthesis

Seq2seq models are used to find a mapping between a source feature sequence  $\mathbf{x}_{1:n} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$  and a target feature sequence  $\mathbf{y}_{1:m} = (\mathbf{y}_1, \dots, \mathbf{y}_m)$  which do not necessarily have to be of the same length, i.e.,  $n \neq m$ . Most seq2seq models have an encoder-decoder structure [5], where advanced ones are equipped with an attention mechanism [6, 7]. First, an encoder (Enc) maps  $\mathbf{x}_{1:n}$  into a sequence of hidden representations  $\mathbf{h}_{1:n} = (\mathbf{h}_1, \dots, \mathbf{h}_n) = \text{Enc}(\mathbf{x}_{1:n})$ . The decoder (Dec) is autoregressive, which means that at time step  $t$ , the previously generated features  $\mathbf{y}_{1:t-1}$  are considered when decoding the current output  $\mathbf{y}_t$ . A context vector  $\mathbf{c}_t$ , which is a weighted sum of  $\mathbf{h}_{1:n}$  calculated by an attention probability vector  $\mathbf{a}_t = (a_t^{(1)}, \dots, a_t^{(n)})$ , is also used by the decoder. Each attention probability  $a_t^{(k)}$  can be thought of as the importance of the hidden representation  $\mathbf{h}_k$  at the  $t$ th time step. The above-mentioned procedure can be formulated as:  $\mathbf{y}_t = \text{Dec}(\mathbf{y}_{1:t-1}, \mathbf{c}_t)$ . During training, the model is optimized via backpropagation using an L1 or L2 loss.

### 2.2. Transformer-based text-to-speech synthesis

In this subsection we describe the Transformer-based TTS system proposed in [22], which we will refer to as Transformer-TTS. Transformer-TTS is a combination of the Transformer [17] architecture and the Tacotron 2 [23] TTS system.

We first briefly introduce the Transformer model [17]. The Transformer relies solely on a so-called multi-head self-attention module that learns sequential dependences by jointly attending to information from different representation subspaces. The main body of Transformer-TTS resembles the original Transformer architecture, which consists of an encoder stack and a decoder stack that are composed of  $L$  encoder layers and  $L$  decoder layers, respectively. An encoder layer contains a multi-head self-attention sublayer followed by a positionwise fully connected feedforward network. A decoder layer, in addition to the two sub-layers in the encoder layer, contains a third sub-layer, which performs multi-head attention over the output of the encoder stack. Each layer is equipped with residual connections and layer normalization. Finally, since no recurrent relation is employed, sinusoidal positional encoding [24] is added to the inputs of the encoder and decoder so that the model can be aware of information about the relative or absolute position of each element.

The model architecture of Transformer-TTS is depicted in Figure 1. Since the Transformer architecture was originally designed for machine translation, several changes have been made to the architecture in [22] to make it compatible in the TTS task. First, as in Tacotron 2, prenets are added to the encoder and decoder sides. Since the text space and the acoustic feature space are different, the positional embeddings are employed with corresponding trainable weights to adapt to the scale of each space. In addition to the linear projection to predict the output acoustic feature, an extra linear layer is added to predict the stop token [23]. A weighted binary cross-entropy loss is used so that the model can learn when to stop decoding. As a common practice in recent TTS models, a five-layer CNN postnet predicts a residual to refine the final prediction.

In this work, our implementation is based on the open-source ESPnet-TTS [25, 26], where the encoder prenet is discarded and the guided attention loss is applied [27] to partial heads in partial decoder layers [18].

## 3. Voice Transformer Network

In this section we describe the combination of Transformer and seq2seq VC. Our proposed model, called the Voice Transformer Network (VTN), is largely based on Transformer-TTS introduced in Section 2.2. Our model consumes the source log-mel spectrogram and outputs the converted log-mel spectrogram. Since both TTS and VC respectively encode text and acoustic features to decode acoustic features, we make a very simple modification to the TTS model, which is to replace the embedding lookup layer in the encoder with a linear projection layer, as shown in Figure 1. Although more complicated networks can be employed, we found that this simple design is sufficient to generate satisfying results. The rest of the model architecture as well as the training process remains the same as that for Transformer-TTS.

An important trick we found to be useful here is to use a reduction factor in both the encoder and the decoder for accurate attention learning. In seq2seq TTS, since the time resolution of acoustic features is usually much larger than that of the text input, a reduction factor  $r_d$  is commonly used on the decoder

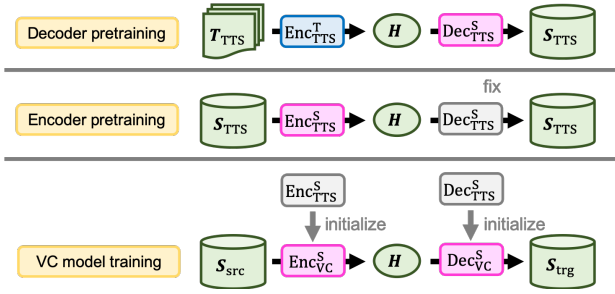


Figure 2: Illustration of proposed TTS pretraining technique for VC.

side [28], where multiple stacked frames are decoded at each time step. On the other hand, although the input and output of VC are both acoustic features, the high time resolution (about 100 frames per second) not only makes attention learning difficult but also increases the training memory footprint. While pyramid RNNs were used to reduce the time resolution in [10], here we simply introduce an encoder reduction factor  $r_e$ , where adjacent frames are stacked to reduce the time axis. We found that this not only leads to better attention alignment but also reduces the training memory footprint by half and subsequently the number of required gradient accumulation steps [26].

## 4. Proposed training strategy with text-to-speech pretraining

We present a text-to-speech pretraining technique that enables fast, sample-efficient training without introducing additional modification or loss to the original model structure or training loss. Assume that, in addition to a small, parallel VC dataset  $D_{VC} = \{S_{src}, S_{trg}\}$ , access to a large single-speaker TTS corpus  $D_{TTS} = \{T_{TTS}, S_{TTS}\}$  is also available.  $S_{src}, S_{trg}$  denote the source, target speech respectively, and  $T_{TTS}, S_{TTS}$  denote the text and speech of the TTS speaker respectively. Our setup is highly flexible in that we do not require any of the speakers to be the same, nor any of the sentences between the VC and TTS corpus to be parallel. We employ a two-stage training procedure, where in the first stage we use  $D_{TTS}$  to learn the initial parameters as a prior, and then use  $D_{VC}$  to adapt to the VC model in the second stage. As argued in Section 1, the ability to generate fine-grained hidden representations  $H$  is the key to a good VC model, so our goal is to find a set of prior model parameters to train the final encoder  $Enc_{VC}^S$  and decoder  $Dec_{VC}^S$ . The overall procedure is depicted in Figure 2.

### 4.1. Decoder pretraining

The decoder pretraining is as simple as training a conventional TTS model using  $D_{TTS}$ . Since text itself contains pure linguistic information, the text encoder  $Enc_{TTS}^T$  here is ensured to learn to encode an effective hidden representation that can be consumed by the decoder  $Dec_{TTS}^S$ . Furthermore, by leveraging large-scale corpora, the decoder is expected to be more robust by capturing various speech features, such as articulation and prosody.

### 4.2. Encoder pretraining

A well pretrained encoder should be capable of encoding acoustic features into hidden representations that are recognizable by the pretrained decoder. With this goal in mind, we train an autoencoder whose decoder is the one pretrained in Section 4.1 and kept fixed during training. The desired pretrained encoder

$Enc_{TTS}^S$  can be obtained by minimizing the reconstruction loss of  $S_{TTS}$ . As the decoder pretraining process described in Section 4.1 takes a hidden representation encoded from text as the input, fixing it in the encoder pretraining process guarantees the encoder to behave similarly to  $Enc_{TTS}^T$ , which is to extract fine-grained, linguistic-information-rich representations.

### 4.3. VC model training

Finally, using  $D_{VC}$ , we train the desired VC models, with the encoder and decoder initialized with  $Enc_{TTS}^S$  and  $Dec_{TTS}^S$  pretrained in Section 4.2 and Section 4.1, respectively. The pretrained model parameters serve as a very good prior to adapt to the relatively scarce VC data, as we will show later. Also, compared with training from scratch, the model takes less than half the training time to converge with the pretraining scheme, enabling extremely efficient training.

## 5. Experimental evaluation

### 5.1. Experimental settings

We conducted our experiments on the CMU ARCTIC database [29], which contains parallel recordings of professional US English speakers sampled at 16 kHz. One female (*slt*) was chosen as the target speaker and one male (*bdl*) and one female (*clb*) were chosen as sources. The number of train/valid/test utterances were 932/100/100. For the TTS corpus, we chose a US female English speaker (*judy bieber*) from the M-AILABS speech dataset [30], which contained 15,200 utterances (32 hours).

The entire implementation was carried out on the open-source ESPnet toolkit [26, 25], including feature extraction, training and benchmarking. We extracted 80-dimensional mel spectrograms with 1024 FFT points and a 256 point frame shift. The base settings for the TTS model and training follow the *Transformer.v1* configuration in [26], and we made minimal modifications to it for VC. The reduction factors  $r_e, r_d$  are both 2 in all VC models. For the waveform synthesis module, we used Parallel WaveGAN (PWG) [31], which is a non-autoregressive variant of the WaveNet vocoder [32, 33]<sup>1</sup>. We trained a speaker-dependent PWG by conditioning on natural mel spectrograms using the full training data of *slt*.

We carried out two types of objective evaluations between the converted speech and the ground truth: the mel cepstrum distortion (MCD), a commonly used measure of spectral distortion in VC, and the character error rate (CER) as well as the word error rate (WER), which estimate the intelligibility. The ASR engine is based on the Transformer [19] and is trained using the LibriSpeech dataset [34]. The CER and WER for the ground-truth test set of *slt* were 0.9% and 3.8%, respectively.

### 5.2. Effectiveness of TTS pretraining

To evaluate the importance and the effectiveness of each pretraining scheme we proposed, we conducted a systematic comparison between different training processes and different sizes of training data. The objective results are in Table 1. First, when the network was trained from scratch without any pretraining, the performance was not satisfactory even with the full training set. With decoder pretraining, a performance boost in MCD was obtained, whereas the ASR results were similar. Nonetheless, as we reduced the training size, the performance dropped dramatically. Finally, by incorporating encoder pretraining, the

<sup>1</sup><https://github.com/kan-bayashi/ParallelWaveGAN>

Table 1: Validation-set objective results of adapted TTS, baseline (ATTS2S), and variants of the proposed VTN.

Speaker	Description	pretraining		932 training utterances			250 training utterances			80 training utterances		
		Encoder	Decoder	MCD	CER	WER	MCD	CER	WER	MCD	CER	WER
clb-slt	VTN	-	-	7.16	15.1	23.2	-	-	-	-	-	-
		-	✓	6.84	9.9	18.1	8.20	52.3	71.7	-	-	-
		✓	✓	<b>6.44</b>	<b>2.4</b>	<b>6.3</b>	<b>6.82</b>	<b>2.9</b>	<b>7.0</b>	<b>7.15</b>	<b>10.6</b>	<b>16.1</b>
	ATTS2S	-	-	7.18	7.5	14.1	7.83	19.9	30.4	8.46	40.4	58.5
bdl-slt	VTN	-	-	7.32	20.9	32.6	-	-	-	-	-	-
		-	✓	7.07	20.3	31.7	8.52	58.9	80.7	-	-	-
		✓	✓	<b>6.56</b>	<b>6.1</b>	<b>12.0</b>	<b>7.14</b>	<b>9.6</b>	<b>16.2</b>	<b>7.34</b>	<b>15.0</b>	<b>23.7</b>
	ATTS2S	-	-	7.36	13.0	22.7	8.01	28.4	43.9	8.76	51.2	72.7
slt	TTS adaptation	-	-	-	3.3	6.2	-	4.5	7.2	-	4.8	7.8

Table 2: Test set subjective evaluation results with 95% confidence intervals of the ground truth, the TTS model, the baseline ATTS2S and the proposed model, VTN. The numbers in the parenthesis indicate the number of training utterances.

Description	Naturalness	Similarity
Ground truth	4.65 ± 0.19	-
TTS adaptation (932)	3.80 ± 0.17	-
ATTS2S (932)	2.67 ± 0.23	50% ± 16%
VTN (932)	<b>4.02 ± 0.17</b>	<b>83% ± 9%</b>
VTN (80)	<b>3.75 ± 0.17</b>	<b>72% ± 12%</b>

model exhibited a significant improvement in all objective measures, where the effectiveness was robust against the reduction in the size of training data.

### 5.3. Comparison with baseline method

Next, we compared our VTN model with an RNN-based seq2seq VC model called ATTS2S [9]. This model is based on the Tacotron model [28] with the help of context preservation loss and guided attention loss to stabilize training and maintain linguistic consistency after conversion. We followed the configurations in [9] but used mel spectrograms instead of WORLD features.

The objective evaluation results of the baseline are reported in Table 1. For the different sizes of training data, our system not only consistently outperformed the baseline method but also remained robust, whereas the performance of the baseline method dropped dramatically as the size of training data was reduced. This proves that our proposed method can improve data efficiency as well as pronunciation. We also observed that when trained from scratch, our VTN model had a similar MCD and inferior ASR performance compared with the baseline. As the ATTS2S employed an extra mechanism to stabilize training, this result may indicate the superiority of using the Transformer architecture over RNNs. We will investigate this in the future.

Systemwise subjective tests on naturalness and conversion similarity were also conducted to evaluate the perceptual performance. For naturalness, participants were asked to evaluate the naturalness of the speech by the mean opinion score (MOS) test on a five-point scale. For conversion similarity, each listener was presented a natural speech of the target speaker and a converted speech, and asked to judge whether they were produced by the same speaker with the confidence of the decision,

i.e., sure or not sure. Ten listeners were recruited.<sup>2</sup>

Table 2 shows the subjective results on the test set. First, with the full training set, our proposed VTN model significantly outperformed the baseline ATTS2S by over one point for naturalness and 30% for similarity. Moreover, when trained with 80 utterances, our proposed method showed only a slight drop in performance, and was still superior to the baseline method. This result justifies the effectiveness of our method and also showed that the pretraining technique can greatly increase data efficiency without severe performance degradation.

## 6. Conclusion

In this work, we successfully applied the Transformer structure to seq2seq VC. Also, to address the problems of data efficiency and mispronunciation in seq2seq VC, we proposed the transfer of knowledge from easily accessible, large-scale TTS corpora by initializing the VC models with pretrained TTS models. A two-stage training strategy that pretrains the decoder and the encoder subsequently ensures that fine-grained intermediate representations are generated and fully utilized. Objective and subjective evaluations showed that our pretraining scheme can greatly improve speech intelligibility, and it significantly outperformed an RNN-based seq2seq VC baseline. Even with limited training data, our system can be successfully trained without significant performance degradation. In the future, we plan to more systematically examine the effectiveness of the Transformer architecture compared with RNN-based models. Extension of our pretraining methods to more flexible training conditions, such as nonparallel training [12], is also an important future task.

## 7. Acknowledgements

This work was supported in part by JST PRESTO Grant Number JPMJPR1657 and JST CREST Grant Number JPMJCR19A3, Japan.

## 8. References

- [1] Y. Stylianou, O. Cappe, and E. Moulines, “Continuous probabilistic transform for voice conversion,” *IEEE Transactions on Speech and Audio Processing*, vol. 6, no. 2, pp. 131–142, 1998.

<sup>2</sup>Audio samples: <https://unilight.github.io/Publication-Demos/publications/transformer-vc/>

- [2] T. Toda, A. W. Black, and K. Tokuda, "Voice Conversion Based on Maximum-Likelihood Estimation of Spectral Parameter Trajectory," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 8, pp. 2222–2235, 2007.
- [3] M. Morise, F. Yokomori, and K. Ozawa, "WORLD: A Vocoder-Based High-Quality Speech Synthesis System for Real-Time Applications," *IEICE Transactions on Information and Systems*, vol. 99, pp. 1877–1884, 2016.
- [4] H. Kawahara, I. Masuda-Katsuse, and A. de Cheveign, "Restructuring speech representations using a pitch-adaptive time-frequency smoothing and an instantaneous-frequency-based F0 extraction: Possible role of a repetitive structure in sounds," *Speech Communication*, vol. 27, no. 3, pp. 187–207, 1999.
- [5] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to Sequence Learning with Neural Networks," in *Advances in Neural Information Processing Systems*, 2014, pp. 3104–3112.
- [6] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [7] T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Lisbon, Portugal, Sep. 2015, pp. 1412–1421.
- [8] H. Miyoshi, Y. Saito, S. Takamichi, and H. Saruwatari, "Voice Conversion Using Sequence-to-Sequence Learning of Context Posterior Probabilities," in *Proc. Interspeech*, 2017, pp. 1268–1272.
- [9] K. Tanaka, H. Kameoka, T. Kaneko, and N. Hojo, "ATTS2S-VC: Sequence-to-sequence Voice Conversion with Attention and Context Preservation Mechanisms," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 6805–6809.
- [10] J. Zhang, Z. Ling, L. Liu, Y. Jiang, and L. Dai, "Sequence-to-Sequence Acoustic Modeling for Voice Conversion," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 3, pp. 631–644, 2019.
- [11] J.-X. Zhang, Z.-H. Ling, Y. Jiang, L.-J. Liu, C. Liang, and L.-R. Dai, "Improving Sequence-to-sequence Voice Conversion by Adding Text-supervision," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6785–6789, 2019.
- [12] J.-X. Zhang, Z.-H. Ling, and L.-R. Dai, "Non-Parallel Sequence-to-Sequence Voice Conversion with Disentangled Linguistic and Speaker Representations," *arXiv preprint arXiv:1906.10508*, 2019.
- [13] G. Zhao, S. Ding, and R. Gutierrez-Osuna, "Foreign Accent Conversion by Synthesizing Speech from Phonetic Posteriorgrams," in *Proc. Interspeech*, 2019, pp. 2843–2847.
- [14] P. Narayanan, P. Chakravarty, F. Charette, and G. Puskorius, "Hierarchical sequence to sequence voice conversion with limited data," *arXiv preprint arXiv:1907.07769*, 2019.
- [15] F. Biadsy, R. J. Weiss, P. J. Moreno, D. Kanvesky, and Y. Jia, "Parrottron: An End-to-End Speech-to-Speech Conversion Model and its Applications to Hearing-Impaired Speech and Speech Separation," in *Proc. Interspeech*, 2019, pp. 4115–4119.
- [16] H. Kameoka, K. Tanaka, T. Kaneko, and N. Hojo, "ConvS2S-VC: Fully convolutional sequence-to-sequence voice conversion," *CoRR*, vol. abs/1811.01609, 2018.
- [17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is All you Need," in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [18] S. Karita, N. Chen, T. Hayashi, T. Hori, H. Inaguma, Z. Jiang, M. Someki, N. E. Y. Soplin, R. Yamamoto, X. Wang *et al.*, "A comparative study on transformer vs RNN in speech applications," *arXiv preprint arXiv:1909.06317*, 2019.
- [19] L. Dong, S. Xu, and B. Xu, "Speech-Transformer: A No-Recurrence Sequence-to-Sequence Model for Speech Recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5884–5888.
- [20] L. Cross Vila, C. Escolano, J. A. R. Fonollosa, and M. R. Costa-Juss, "End-to-End Speech Translation with the Transformer," in *Proc. IberSPEECH*, 2018, pp. 60–63.
- [21] M. A. D. Gangi, M. Negri, R. Cattoni, R. Dessi, and M. Turchi, "Enhancing Transformer for End-to-end Speech-to-Text Translation," in *Proceedings of Machine Translation Summit XVII Volume 1: Research Track*, 2019, pp. 21–31.
- [22] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, "Neural Speech Synthesis with Transformer Network," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 6706–6713.
- [23] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan, R. A. Saurous, Y. Ajiomyriannakis, and Y. Wu, "Natural TTS Synthesis by Conditioning WaveNet on MEL Spectrogram Predictions," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4779–4783.
- [24] J. Gehring, M. Auli, D. Grangier, D. Yarats, and Y. N. Dauphin, "Convolutional Sequence to Sequence Learning," in *Proceedings of the 34th International Conference on Machine Learning*, vol. 70, 2017, pp. 1243–1252.
- [25] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, "ESPnet: End-to-End Speech Processing Toolkit," in *Proc. Interspeech*, 2018, pp. 2207–2211.
- [26] T. Hayashi, R. Yamamoto, K. Inoue, T. Yoshimura, S. Watanabe, T. Toda, K. Takeda, Y. Zhang, and X. Tan, "ESPnet-TTS: Unified, Reproducible, and Integratable Open Source End-to-End Text-to-Speech Toolkit," *arXiv preprint arXiv:1910.10909*, 2019.
- [27] H. Tachibana, K. Uenoyama, and S. Aihara, "Efficiently Trainable Text-to-Speech System Based on Deep Convolutional Networks with Guided Attention," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4784–4788.
- [28] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. Le, Y. Ajiomyriannakis, R. Clark, and R. A. Saurous, "Tacotron: Towards end-to-end speech synthesis," in *Proc. Interspeech*, 2017, pp. 4006–4010.
- [29] J. Kominek and A. W. Black, "The CMU ARCTIC speech databases," in *Fifth ISCA Workshop on Speech Synthesis*, 2004.
- [30] Munich Artificial Intelligence Laboratories GmbH, "The M-AILABS speech dataset," 2019, accessed 30 November 2019. [Online]. Available: <https://www.caito.de/2019/01/the-m-ailabs-speech-dataset/>
- [31] R. Yamamoto, E. Song, and J.-M. Kim, "Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram," *arXiv preprint arXiv:1910.11480*, 2019.
- [32] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," *arXiv preprint arXiv:1609.03499*, 2016.
- [33] A. Tamamori, T. Hayashi, K. Kobayashi, K. Takeda, and T. Toda, "Speaker-dependent wavenet vocoder," in *Proc. Interspeech*, 2017, pp. 1118–1122.
- [34] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "LibriSpeech: An ASR corpus based on public domain audio books," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.