





Pretraining Techniques for Sequence-to-Sequence Voice Conversion

Wen-Chin Huang , *Student Member, IEEE*, Tomoki Hayashi , Yi-Chiao Wu ,
Hirokazu Kameoka , *Senior Member, IEEE*, and Tomoki Toda, *Senior Member, IEEE*

Abstract—Sequence-to-sequence (seq2seq) voice conversion (VC) models are attractive owing to their ability to convert prosody. Nonetheless, without sufficient data, seq2seq VC models can suffer from unstable training and mispronunciation problems in the converted speech, thus far from practical. To tackle these shortcomings, we propose to transfer knowledge from other speech processing tasks where large-scale corpora are easily available, typically text-to-speech (TTS) and automatic speech recognition (ASR). We argue that VC models initialized with such pretrained ASR or TTS model parameters can generate effective hidden representations for high-fidelity, highly intelligible converted speech. In this work, we examine our proposed method in a parallel, one-to-one setting. We employed recurrent neural network (RNN)-based and Transformer based models, and through systematical experiments, we demonstrate the effectiveness of the pretraining scheme and the superiority of Transformer based models over RNN-based models in terms of intelligibility, naturalness, and similarity.

Index Terms—Voice conversion, sequence-to-sequence, pretraining, transformer.

I. INTRODUCTION

VOICE conversion (VC) aims to convert the speech from a source to that of a target without changing the linguistic content [1]. In conventional VC systems, with the assumption that a parallel corpus, i.e., pairs of speech samples with identical linguistic contents uttered by both source and target speakers, exists, an *analysis—conversion—synthesis* paradigm is often adopted [2]. First, a high-quality vocoder such as WORLD [3] or STRAIGHT [4] is utilized to extract different acoustic features,

Manuscript received August 6, 2020; revised November 22, 2020 and December 16, 2020; accepted December 19, 2020. Date of publication January 5, 2021; date of current version February 1, 2021. This work was supported in part by JST, CREST under Grant JPMJCR19A3, and in part by JSPS KAKENHI under Grant 17H06101. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Yu Tsao. (*Corresponding author: Wen-Chin Huang.*)

Wen-Chin Huang and Yi-Chiao Wu are with the Graduate School of Informatics, Nagoya University, Nagoya, Aichi 464-8601, Japan (e-mail: wen.chinhuang@g.sp.m.is.nagoya-u.ac.jp; yichiao.wu@g.sp.m.is.nagoya-u.ac.jp).

Tomoki Hayashi is with the Human Dataware Laboratory Company Ltd., and the Graduate School of Information Science, Nagoya University, Nagoya, Aichi 464-8601, Japan (e-mail: hayashi.tomoki@g.sp.m.is.nagoya-u.ac.jp).

Hirokazu Kameoka is with the NTT Communication Science Laboratories, Nippon Telegraph and Telephone Corporation, Atsugi, Kanagawa 243-0198, Japan (e-mail: kameoka.hirokazu@lab.ntt.co.jp).

Tomoki Toda is with the Information Technology Center, Nagoya University, Nagoya, Aichi 464-8601, Japan (e-mail: tomoki@icts.nagoya-u.ac.jp).

Digital Object Identifier 10.1109/TASLP.2021.3049336

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. This restricts the modeling of the speaking rate. To summarize, the conversion of prosody, including F0 and duration, is overly simplified in the past 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 and capture long-term dependencies. This ability makes the seq2seq model a natural choice to convert prosody in VC. It has been shown that seq2seq VC models can outperform conventional frame-wise VC systems, especially in terms of conversion similarity [8]–[10]. This is owing to the fact that the suprasegmental characteristics of F0 and duration patterns well handled in seq2seq VC models are closely correlated with the speaker identity.

Despite the promising results, seq2seq VC models suffer from two major problems. First, seq2seq models usually require a large amount of training data to generalize well, although parallel VC corpora are often hard to collect. Second, as pointed out in [11], the converted speech in seq2seq VC systems still suffers from mispronunciations and other linguistic-inconsistency problems, such as inserted, repeated and skipped phonemes. This is mainly due to the failure in attention alignment learning, which can also be seen as a consequence of insufficient data. In this study, we focus on improving seq2seq modeling in a parallel, one-to-one setting.

As one popular solution to limited training data, pretraining has been regaining attention in recent years, where knowledge from massive, out-of-domain data is transferred to aid learning in the target domain. This concept is usually realized by learning universal, high-level feature representations. In the field of computer vision, *supervised* pretraining (e.g., ImageNet classification [12]–[14]) followed by fine-tuning on tasks with less training data (e.g., object detection [15]–[17], segmentation [18], [19] and style transfer [20], [21]) often leads to state-of-the-art results. On the other hand, many natural

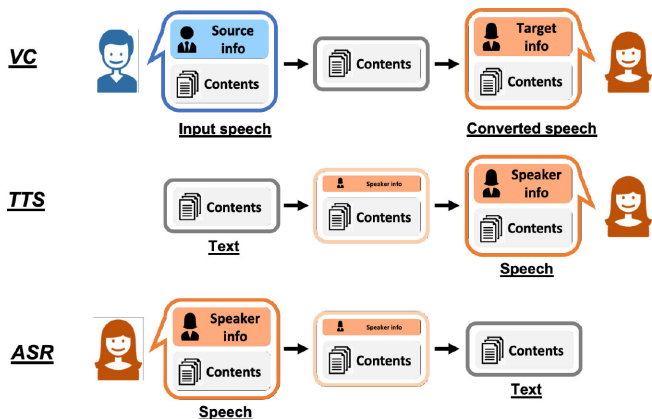


Fig. 1. Illustration of the relationship of VC, TTS and ASR in an information perspective.

language understanding (NLU) tasks learn rich representation through an *self-supervised* language model objective [22]–[26], which have also been shown to boost performance. Efforts have also been dedicated to speech representation learning for tasks including automatic speech recognition (ASR) [27] and speaker identification/verification [28].

In our prior work [10], we proposed a text-to-speech (TTS) pretraining strategy for seq2seq VC. Instead of learning *feature representations*, we focused on finding *model parameters* that can easily learn effective representations in an *supervised* manner. We initially adopted the text-to-speech (TTS) task based on the facts that (1) in a broader definition, both VC and TTS can be categorized into the speech synthesis task [29], so the respective models should share some characteristics. (2) Vast studio-level large-scale TTS corpora have been made easily accessible by the community. It was demonstrated that the VC model initialized with TTS pretrained model parameters can generate high-quality, highly intelligible speech even with very limited training data. However, we compared with a baseline model [8] in a system-wise manner, and we found the need of a complete study on the proposed method.

In this work, we pursue a comprehensive understanding of the pretraining technique for seq2seq VC. In addition to TTS, we consider another task for knowledge transfer: ASR. We refer to them as *TTS-oriented pretraining* and *ASR-oriented pretraining*, respectively. Here, we provide a unified, intuitive explanation on the choice of these two source of transfer from an information perspective, as depicted in Fig. 1. Roughly speaking, speech consists of the linguistic contents and the speaker identity. The goal of VC is to remove the source speaker information from the source speech, and then inject the identity of the target speaker. Thus, a speaker-free intermediate feature space would be essential for a successful VC model, which is hard to facilitate given only a parallel corpus. On the other hand, TTS and ASR both aim to find a mapping between text and speech, as the former tries to add speaker information to the source while the latter tries to remove. We therefore suspect that the intermediate hidden representation spaces of these two tasks contain somewhat little speaker information, and serve as a suitable fit for VC.

Our contributions in this work are as follows:

- For the first time, we propose ASR-oriented pretraining for seq2seq VC.
- We examine, through systematical objective and subjective evaluations, the TTS and ASR tasks for pretraining in seq2seq VC. Results show that they are both effective with sufficient data while only TTS pretraining maintains robust against the reduction of data.
- We visualize the hidden representation spaces of the learned models using different pretraining tasks and how they relate to the performance.
- We compare two different model architectures for seq2seq VC: recurrent neural networks (RNNs) and Transformers, and we show the supremacy of the latter over the former, which is consistent with the finding in most speech processing tasks [30].

II. RELATED WORK

A. Data Deficiency in Sequence-to-Sequence Voice Conversion

In VC, it is common to limit the size of training data to around 5 or 10 minutes [31], [32], wherein existing seq2seq VC literature, approximately 1 h of data is usually used. Even with such amount of data, it is still required to resort to certain techniques to successfully train seq2seq VC models. These techniques can be categorized into the followings:

Extra module: Many have utilized an external ASR module pretrained on a large dataset during training and runtime. For example, the phonetic posteriorgram (PPG) extracted from ASR is a commonly used linguistic clue in VC [33], and can be used as the only input [34], [35] or as an additional clue [9], [11] in seq2seq VC. On the other hand, Parrotron [36] used an external TTS system to generate artificial data from a large hand-transcribed corpus for training an any-to-one (normalization) VC model. The disadvantage of using external modules is that the performance depends on the extra module. The accuracy of the PPG and the quality of the TTS system can bound the performance of the final VC system.

Text label: Text labels provide strong supervision to ensure linguistic consistency. Methods utilizing such labels include multitask learning meaningful hidden representation [11], [36], [37], data augmentation [11] or representations disentanglement [37]. Yet, labeling errors and failed force alignments might cause potential performance degradation.

Regularization: As multitask learning and feature disentanglement can be seen as regularizations, some have also proposed to impose constraints on the model without any external resources. [8], [38] proposed the context preservation loss and the guided attention loss, and [34] proposed to use local attention to stabilize training. Nonetheless, such regularization often requires rigorous weight tuning.

Our proposed pretraining method is closest to the use of extra modules, except that the external resource is only used to obtain a prior for effective finetuning, instead of generating PPGs or artificial data. Also, our method needs neither text label of the VC data nor carefully designed regularization methods, yet can still achieve great data efficiency.

B. Pretraining in Speech Processing

Early applications of pretraining deep neural networks for speech processing mainly lied in ASR, with the main goal of speeding up optimization and reducing generalization error [39], [40]. In recent years, inspired by the breakthrough in NLU, unsupervised or self-supervised speech representation learning utilizing massive, untranscribed speech data has become a popular research topic. As language modeling objectives have been widely employed for pretraining in NLU, finding a universally effective objective is still an active research area. Various objectives have been proposed, such as autoencoding [41]–[43] sometimes with an autoregressive model [44], [45] or contrastive learning [27], [28], [46], [47]. Nonetheless, different pretraining objectives lead to different representations, and an effective objective for VC is still unclear.

Different from the above mentioned unsupervised approaches, our method relies on *supervised* pretraining with well-defined speech processing objectives. As we adopt popular speech processing tasks, large scale datasets can be assumed easily accessible thanks to the vastly growing community. We also expect that performance of pretraining would benefit from the rapid development of state-of-the-art models, thus improving the quality of the downstream VC task.

C. Transfer Learning From ASR and TTS for VC

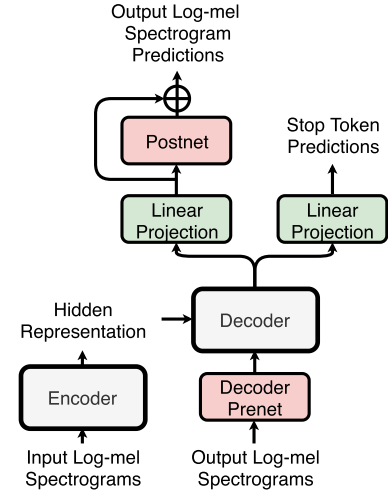
In a recent VC review paper [48], it was shown that a sufficient amount of efforts has been dedicated to transferring knowledge from ASR and TTS to improving various aspects of VC, regardless of using a seq2seq model or not. The PPG-based methods [49]–[53] and the Parratron system described in Section II-A facilitated nonparallel, any-to-one VC by utilizing ASR and TTS modules, respectively. Another line of work show that training an integrated system capable of performing either TTS or VC can boost individual performances [29], [54]. Bootstrapping from a speaker-adaptive or multispeaker TTS model is another active direction [55], [56]. However, few of the abovementioned methods were designed to tackle the data deficiency problem for seq2seq based parallel, one-to-one VC, which is the main scope of this paper.

III. MODEL ARCHITECTURES


Most existing seq2seq VC models are based on RNNs [8], [9], [11], [34]–[37], [57], [58]. In our prior work [10], we successfully applied the Transformer architecture [59] to seq2seq VC. We first provide a unified model structure of seq2seq VC and describe the shared components.

A. Unified seq2seq VC Model Structure

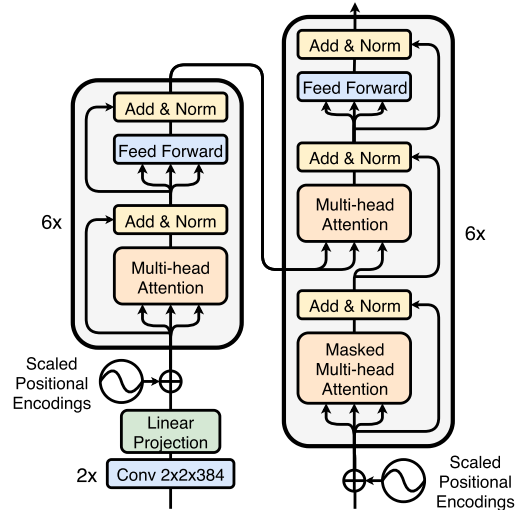
Seq2seq models learn a mapping between a source feature sequence $X = \mathbf{x}_{1:n} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ and a target feature sequence $Y = \mathbf{y}_{1:m} = (\mathbf{y}_1, \dots, \mathbf{y}_m)$ which are often of different length, i.e. $n \neq m$. As with most seq2seq models, our seq2seq VC model also has an encoder–decoder structure [5], as depicted in Fig. 2(a). The encoder (Enc) first maps the input acoustic feature



(a) Illustration of the unified seq2seq VC model architecture and the shared components.



(b) RNN-based encoder and decoder.



(c) Transformer-based encoder and decoder.

Fig. 2. Model architectures that are used in this work.

sequence $\mathbf{x}_{1:n}$ into a sequence of hidden representations:

$$H = \mathbf{h}_{1:n} = \text{Enc}(\mathbf{x}_{1:n}). \quad (1)$$

The decoder (Dec) is autoregressive, which means that when decoding the current output \mathbf{y}_t , in addition to the encoder output, i.e. the hidden representations $\mathbf{h}_{1:n}$, the previously generated features $\mathbf{y}_{1:t-1}$ are also considered:

$$\mathbf{y}_t = \text{Dec}(\mathbf{h}_{1:n}, \mathbf{y}_{1:t-1}). \quad (2)$$

Some extra components and techniques are adopted in our seq2seq VC model to improve performance and stabilize training, most of which are inspired by the success of modern seq2seq TTS models [60], [61].

- A prenet containing 2 fully connected layers is added to the decoder, which serves as an information bottleneck essential for learning the autoregressive decoder.
- A linear projection layer is used to project the decoder output to have the desired dimension. To learn when to stop decoding, a separate linear projection layer is used to predict a stopping probability, which can be used with a threshold to decide when to stop decoding during inference
- To compensate for the missing context information in the autoregressive decoder, a five-layer CNN postnet is used to predict a residual that is added to the projected output.
- Introducing the reduction factor r greatly helps speed up convergence and reduce training time and memory footprint. Specifically, at each decoding step, r non-overlapping frames are predicted. Since adjacent speech frames are often correlated, this technique allows the decoder to correctly model the interaction with the hidden representation sequence.

The training objectives include an L1 and L2 loss, in combination with a weighted binary cross-entropy loss on the stop token prediction. The whole network is composed of neural networks and optimized via backpropagation.

B. RNN-Based Model

Our RNN-based seq2seq VC model is based on the Tacotron2 TTS model [61] and resembles the work in [8]. The encoder first linearly projects the input log-mel spectrogram, followed by a stack of convolutional layers, batch normalization, and ReLU activations. The output of the final convolutional layer is then passed into a bi-directional LSTM layer to generate the hidden representations.

For each decoder output step, an attention mechanism [6], [7] is used to attend to different positions of the hidden representation sequence. First, a context vector \mathbf{c}_t is calculated as a weighted sum of $\mathbf{h}_{1:n}$, where the weight is represented using an attention probability vector $\mathbf{a}_t = (a_t^{(1)}, \dots, a_t^{(n)})$. Each attention probability $a_t^{(k)}$ can be thought of as the importance of the hidden representation \mathbf{h}_k at the current time step. As in Tacotron2, we adopt the location-sensitive attention [62], which takes cumulative attention weights from previous decoder time steps as an additional feature to encourage a forward consistency to prevent repeated or missed phonemes. The context vector is then concatenated with the prenet output and passed into a stacked uni-directional LSTM network to predict the r output frames.

For seq2seq speech synthesis models, the attention alignment is usually monotonic and linear, so a guided attention loss that encourages the attention matrix to be diagonal can speed up attention learning and convergence [8], [63]. In addition, in [8], a context preservation loss was applied to maintain linguistic consistency after conversion [8].

C. Transformer-Based Model

Our Transformer-based seq2seq VC model, which we refer to as the Voice Transformer Network (VTN), is based on the Transformer architecture [59], which was originally designed for machine translation but also widely applied to other sequential modeling problems. There are several core components of the Transformer:

Multi-head attention (MHA) sublayer: The MHA layer is defined as:

$$\text{MHA}(Q, K, V) = [\text{head}_1, \dots, \text{head}_h]W^O, \quad (3)$$

$$\text{head}_i = \text{Att}(QW_i^Q, KW_i^K, VW_i^V), \quad (4)$$

where Q , K and V denote the input matrices that, following [59], are referred to as the query, key and value, respectively. MHA uses h different, learned linear projections W^Q, W^K, W^V to map the inputs to different *heads*, and then perform the Att operation in parallel. The outputs from all heads are concatenated and projected with W^O . As in [59], the Att operation is implemented scaled dot-product attention is used:

$$\text{Att}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_{\text{att}}}}\right)V, \quad (5)$$

where d_{att} is the attention dimension.

Position-wise feed-forward network (FFN) layer: The FFN layer is defined as:

$$\text{FFN}(\mathbf{x}) = \max(0, \mathbf{x}W_1 + b_1)W_2 + b_2, \quad (6)$$

which is independently applied at each time step (position) with different parameters from layer to layer.

Layer normalization and residual connection: Around either of the above-mentioned sublayers, a residual connection followed by layer normalization [64] is employed. For input X of a sublayer, the output is given as:

$$\text{LayerNorm}(X + \text{Sublayer}(X)). \quad (7)$$

Due to the residual connections, all sublayers have the same output dimension d_{model} .

Scaled positional encoding (SPE): In the original Transformer [59], since no recurrent relation is employed in the Transformer, to let the model be aware of information about the relative or absolute position of each element, the triangular (sinusoidal) positional encoding (PE) [65] is added to the inputs to the encoder and decoder. In this work, we adopt the SPE [66], which is a generalized version of the original PE that scales the encodings with a trainable weight α , so that they can adaptively fit the scales of the encoder and the decoder:

$$\text{SPE}(t) = \begin{cases} \alpha \cdot \sin\left(\frac{t}{10000 \frac{d_{\text{model}}}{2t}}\right), & \text{if } t \text{ is even,} \\ \alpha \cdot \cos\left(\frac{t}{10000 \frac{d_{\text{model}}}{2t}}\right), & \text{if } t \text{ is odd.} \end{cases} \quad (8)$$

The encoder we adopt in this work resembles the one in [67]. First, the input acoustic feature sequence is downsampled in the time and frequency axes by a fraction of 4 using two convolutional layers with stride 2×2 . While the reduction of the memory footprint is a clear benefit, a hidden representation with a low sampling rate can not only speed up attention learning

convergence due to easier attention calculation at each decoding step but also approximates phoneme-level or even character-level linguistic contents [9]. After linearly projecting to d_{model} -dimensions and adding the SPE, L identical encoder layers are stacked to form the core of the encoder. Each encoder layer consists of an MHA sublayer and an FFN sublayer, followed by a residual connection and layer normalization. The MHA layers in the encoder are *self-attention* layers since the queries, keys, and values are all from the output from the previous layer.

The decoder in this work is composed of the same number of L identical decoder layers as in the encoder. In each decoder layer, the first sublayer is the so-called *masked* self-attention MHA sublayer, where a mask is utilized such that at time step t , only vectors with time index up to and including t can be accessed. This preserves the autoregressive property of the model. Then, an MHA sublayer uses the outputs from the previous layer as queries and H as the keys and values, which ensembles the encoder–decoder attention in III-B. Finally, an FFN sublayer is used, as in the encoder. Again, all sublayers are wrapped with a residual connection and layer normalization.

In addition to the L1, L2, and weighted binary cross-entropy losses, the guided attention loss is also applied. As pointed out in [66], in Transformer-based speech synthesis, not all attention heads demonstrate diagonal alignments, so following [30], [68], the guided attention loss is applied to partial heads in partial decoder layers.

IV. PRETRAINING TECHNIQUES

In seq2seq models for speech applications, effective intermediate representations can facilitate correct attention learning that bridges the encoder and the decoder, thus crucial to success. By the definition of VC, it is natural to try to encode the linguistic contents of the source speech into the hidden representations so that they can be maintained. Thus, we conjecture that the core ability of successful seq2seq VC models is to generate and utilize high-fidelity hidden representations.

In theory, both TTS and ASR tasks aim to find a mapping between two modalities: speech and text. As speech signals contain all essential linguistic information, the hidden representation spaces induced by these two tasks should lie in the middle of the spectrum between speech and text. Thus, we conjecture that such space is desirable for seq2seq VC models, thus suitable for pretraining.

In this work, we extend the TTS-oriented pretraining technique previously proposed in [10] to both TTS and ASR, as depicted Fig. 3. Specifically, a two-stage training procedure is employed: in the first pretraining stage, a large-scale corpus is used to learn the initial seq2seq model parameters as a prior; then, in the second stage, the seq2seq VC model is initialized with the pretrained model parameters and trained with a relatively smaller VC dataset. The goal of this pretraining procedure is to provide fast, sample-efficient VC model learning, thus reducing the data size requirement and training time. In addition, this 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 pretraining corpus and the VC dataset to be parallel.

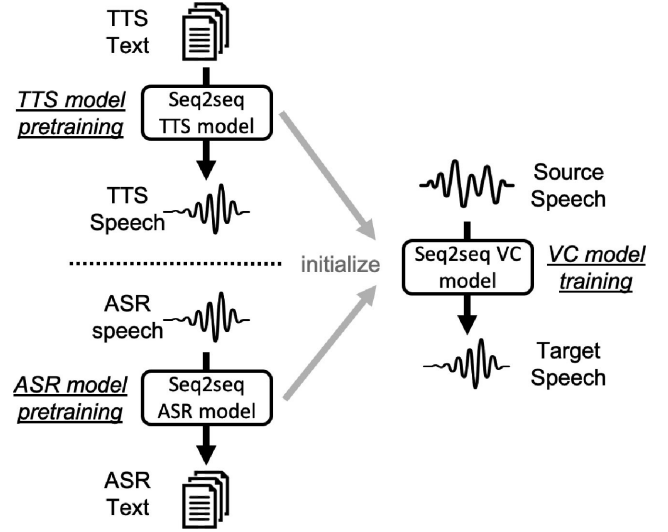


Fig. 3. Illustration of the concept of pretraining from seq2seq TTS or ASR to seq2seq VC.

Let the parallel VC dataset be $D_{\text{VC}} = \{S_{\text{src}}, S_{\text{trg}}\}$, where $S_{\text{src}}, S_{\text{trg}}$ denote the source, target speech, respectively. Our goal is to find a set of prior model parameters to train the final encoder Enc_{VC}^S and decoder Dec_{VC}^S .

A. TTS-Oriented Pretraining

In this subsection we review the TTS-oriented pretraining technique [10]. We assume that access to a large single-speaker TTS corpus $D_{\text{TTS}} = \{T_{\text{TTS}}, S_{\text{TTS}}\}$ is available, where $T_{\text{TTS}}, S_{\text{TTS}}$ denote the text and speech of the TTS speaker respectively. The pretraining can be broken down into two steps.

A.1 *Decoder pretraining*: As in A.1 in Fig. 4, the decoder is pretrained, on D_{TTS} , by training a conventional TTS model composed of a text encoder $\text{Enc}_{\text{TTS}}^T$ and a speech decoder $\text{Dec}_{\text{TTS}}^S$.

A.2 *Encoder pretraining*: Then, as in A.2 in Fig. 4, the encoder is pretrained, also on the same D_{TTS} , by training an autoencoder which takes S_{TTS} as input and output. The decoder here is the pretrained $\text{Dec}_{\text{TTS}}^S$ and we fix the parameters so that they are not updated during training. The desired pretrained encoder $\text{Enc}_{\text{TTS}}^S$ can then be obtained by minimizing the reconstruction loss.

The intuition of the encoder pretraining is to obtain an encoder capable of encoding acoustic features into hidden representations that are recognizable by the well pretrained decoder. Another interpretation is that the final pretrained encoder $\text{Enc}_{\text{TTS}}^S$ tries to mimic the text encoder $\text{Enc}_{\text{TTS}}^T$. In the first decoder pretraining step, since text itself contains pure linguistic information, the text encoder $\text{Enc}_{\text{TTS}}^T$ is ensured to learn to encode an effective hidden representation that can be consumed by the decoder $\text{Dec}_{\text{TTS}}^S$. Fixing the decoder in the encoder pretraining process, as a consequence, guarantees the encoder to behave similarly to the text encoder, which is to extract fine-grained, linguistic-information-rich representations.

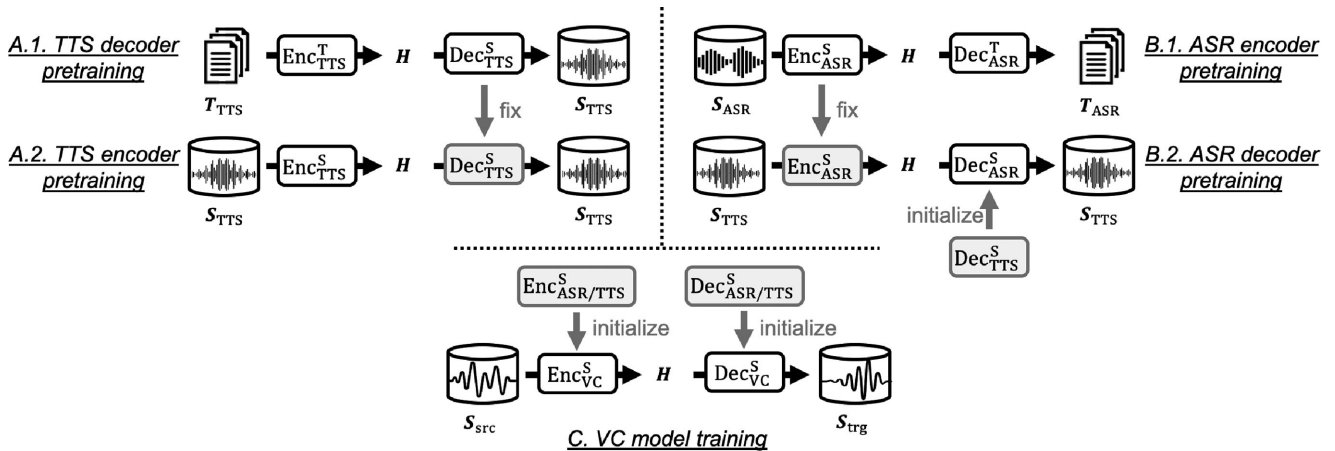


Fig. 4. Diagram of the pretraining procedures for TTS and ASR. Top left: TTS-oriented pretraining. Top right: ASR-oriented pretraining. Bottom: VC model training.

B. ASR-Oriented Pretraining

In this subsection we describe how to extend the TTS-oriented pretraining technique in IV-A to ASR. We assume that a large *multi-speaker* ASR corpus $D_{ASR} = \{S_{ASR}, T_{ASR}\}$ is available, where S_{ASR}, T_{ASR} denote the speech and text data in D_{ASR} , respectively. Similar to TTS-oriented pretraining, the ASR-oriented pretraining is again broken down into two steps.

B.1 Encoder pretraining: First, the *encoder* is pretrained, on D_{ASR} , by training a conventional ASR model consisting of a speech encoder Enc_{ASR}^S and a text decoder Dec_{ASR}^T , as in B.1 in Fig. 4.

B.2 Decoder pretraining: Differently, the decoder pretraining is performed on D_{TTS} , rather than on D_{ASR} . This is because D_{ASR} is a multi-speaker corpus, but the VC model architecture in this work focuses on one-to-one VC, i.e. modeling the conversion between one source speaker and one target speaker, thus cannot model individual speaker characteristics. Again, the decoder pretraining uses S_{TTS} as input and output, and the encoder is the pretrained Enc_{ASR}^S and kept fixed during training. To speed up convergence, we initialize the decoder with the one obtained in TTS decoder pretraining, namely Dec_{TTS}^S . The desired pretrained decoder Enc_{ASR}^S can then be obtained by minimizing the reconstruction loss. The decoder pretraining procedure is depicted in B.2 in Fig. 4.

The intuition of the ASR decoder pretraining is different from that of the TTS encoder pretraining. The ASR speech encoder Enc_{ASR}^S , trained with the ASR objective, should generate a compact hidden representation for decoding underlying linguistic contents. Such representations are believed to be easier to map to speech, thus suitable for pretraining the speech decoder Dec_{ASR}^S .

C. VC Model Training

Finally, as in [10], D_{VC} is used to train the desired VC models Enc_{VC}^S and Dec_{VC}^S , with the encoder initialized with either Enc_{TTS}^S or Enc_{ASR}^S , and the decoder with Dec_{TTS}^S or Dec_{ASR}^S , respectively. As we will show later, the pretrained model parameters serve as a very good prior to adapt to the relatively scarce VC data, achieving significantly better conversion performance.

V. EXPERIMENTAL EVALUATIONS

A. Experimental Settings

1) Data: For D_{VC} , we conducted our experiments on the CMU ARCTIC database [69], which contains parallel recordings of professional US English speakers sampled at 16 kHz. Data from four speakers were used: a male source speaker (*bdl*) and a female source speaker (*clb*), as well as a male target speaker (*rms*) and a female target speaker (*slt*). 100 utterances were selected for each validation and evaluation sets, and the remaining 932 utterances were used as training data. For D_{TTS} , we chose a US female English speaker (*judy bieber*) from the M-AILABS speech dataset [70]. With the sampling rate also at 16 kHz, the training set contained 15 200 utterances, which were roughly 32 hours long. For D_{ASR} , we used the LibriSpeech dataset [71] and pooled *train-clean-100* and *train-clean-360* together to get 460 hours of data from roughly 1170 speakers.

2) Implementation: The entire experiment was carried out on the open-source ESPnet toolkit [68], [72], including feature extraction, training and benchmarking. The official implementation has been made publicly available,¹ and since readers may access all the settings and configurations online, we omit the detailed hyperparameters here. For the VC training, 80-dimensional mel filterbanks with 1024 FFT points and a 256 point frame shift was used as the acoustic features. We used the LAMB optimizer [73] and set the learning rate to 0.001.

3) Waveform Synthesis Module: We used the Parallel WaveGAN (PWG) [74], which is a non-autoregressive variant of the WaveNet vocoder [75], [76] and enables parallel, faster than real-time waveform generation.² Since speaker-dependent neural vocoders outperform speaker-independent ones [77], we trained a speaker-dependent PWG conditioned on natural mel spectrograms, one for each target speaker. Note that we used the full training dataset, since the goal is to demonstrate the effects of various methods, so we did not train separate PWGs w.r.t. different training data sizes.

¹[Online]. Available: <https://github.com/espnet/espnet/tree/master/egs/arctic/vc1>

²We followed the open-source implementation [Online]. Available: <https://github.com/kan-bayashi/ParallelWaveGAN>

4) *Objective Evaluation Metrics*: We carried out three types of objective evaluations between the converted speech and the ground truth.

- Mel cepstrum distortion (MCD) [78]: The MCD is a commonly used measure of spectral distortion in VC, which is based on mel-cepstral coefficients (MCCs). It is defined as:

$$\text{MCD}[\text{dB}] = \frac{10}{\log 10} \sqrt{2 \sum_{d=1}^K (\text{mcc}_d^{(c)} - \text{mcc}_d^{(t)})^2}, \quad (9)$$

where K is the dimension of the MCCs and $\text{mcc}_d^{(c)}$ and $\text{mcc}_d^{(t)}$ represent the d -th dimensional coefficient of the converted MCCs and the target MCCs, respectively. In practice, MCD is calculate in a utterance-wise manner. A dynamic time warping (DTW) based alignment is performed to find the corresponding frame pairs between the non-silent converted and target MCC sequences beforehand. We used the WORLD vocoder [3] for MCC extraction and silence frame decisions, and set $K = 24$.

- F0 root mean square error (FORMSE): Since using seq2seq modeling in VC can greatly improve prosody conversion, we report the RMSE between the F0 of converted speech and that of the reference target speech. Similar to the calculation of MCD, DTW-based is performed and we take only the non-silent frames into account.
- Character/word error rate (CER/WER): The CER/WER is an underestimate of the intelligibility of the converted speech. The ASR engine is based on the Transformer architecture [67] and is trained using the LibriSpeech dataset [71]. The CER and WER for the ground-truth validation set were 0.9% and 3.8%, respectively, which could be regarded as the upper bound.

Note that to avoid overfitting, we used the validation set MCD as the criterion for model selection, and the best performing models were proceeded to generate the samples for the subjective test.

5) *Subjective Evaluation Methods*: The following subjective evaluations were performed using the open-source toolkit [79] which implements the ITU-T Recommendation P.808 [80] for subjective speech quality assessment in the crowd using the Amazon Mechanical Turk (Mturk), and screens the obtained data for unreliable ratings. We recruited more than fifty listeners.³

- The mean opinion score (MOS) test on naturalness: Subjects were asked to evaluate the naturalness of the converted and natural speech samples on a scale from 1 (completely unnatural) to 5 (completely natural).
- The VCC [32] style test on similarity: This paradigm was adopted by the VCC organizing committee. Listeners were given a pair of speech utterances consisting of a natural speech sample from a target speaker and a converted speech sample. Then, they were asked to determine whether the pair of utterances can be produced by the same speaker,

with 4-level confidence of their decision, i.e., sure or not sure.

B. Effectiveness of TTS-Oriented Pretraining on RNN and Transformer Based Models

First, we show that TTS-oriented pretraining is a technique effective on not only VTN but also RNN-based seq2seq VC models. The objective results are in Table I. First, without pretraining, both VTN and RNN could not stay robust against the reduction of training data. The performance dropped dramatically with the reduction of training data, where a similar trend was also reported in [37]. This identifies the data efficiency problem of seq2seq VC. By incorporating TTS-oriented pretraining, both VTN and RNN exhibited a significant improvement in all objective measures, where the effectiveness was robust against the reduction in the size of training data. With only 80 utterances, both models can achieve comparable performance to that of using the full training dataset except the FORMSE, wherein the case of VTN, the intelligibility is even better.

The subjective results are in Table III. Without pretraining, the VTN and RNN suffered from about 1.2 and 0.8 MOS points drop when the training data reduces from 932 to 80 utterances. On the other hand, with TTS-oriented pretraining applied, the naturalness of VTN and RNN improved by more than 1 point with 932 utterances and more than 2 points with 80 utterances. Moreover, when the training data reduces, there was only a very limited performance drop. These results demonstrate the effectiveness of the TTS-oriented pretraining technique.

C. Comparison of TTS-Oriented and ASR-Oriented Pretraining

Next, we compare the effectiveness of TTS-oriented and ASR-oriented pretraining. From Tables I and III, with the full training set, ASR-oriented pretraining could bring almost the same amount of improvement compared to TTS-oriented pretraining. However, as the size of training data reduces, the performance of the ASR-oriented pretrained model dropped significantly, except FORMSE. This shows that ASR-oriented pretraining lacks the robustness essential for practical VC.

To investigate the failure of ASR-oriented pretrained models against limited training data, we chose one sentence from the evaluation set and show the ASR results of the converted samples using TTS-oriented and ASR-oriented pretrained VTNs in Table II. Although TTS-oriented pretraining could not ensure complete linguistic consistency, the errors were minor and possibly due to the imperfect ASR engine used for evaluation, thus the result seems reasonable. On the other hand, the recognition result of the ASR-oriented pretrained model with 80 utterances that had no connection to the source sentence. We conclude that linguistic consistency is poorly maintained under the limited data scenario using ASR-oriented pretraining.

D. Comparison of RNN and Transformer Based Models

In [10], the VTN was shown to outperform the RNN-based model [8], while it was not clear whether the improvement came from different model architectures or the pretraining technique.

³A demo web page with samples used for subjective evaluation [Online]. Available: <https://unilight.github.io/Publication-Demos/publications/vtn-taslp/index.html>

TABLE I

VALIDATION-SET OBJECTIVE EVALUATION RESULTS OF VTNS WITH NO PRETRAINING, TTS-ORIENTED PRETRAINING, ASR-ORIENTED PRETRAINING, AND RNN-BASED MODELS WITH NO PRETRAINING AND TTS-ORIENTED PRETRAINING, WHICH ARE TRAINED ON DIFFERENT CONVERSION PAIRS AND DIFFERENT SIZES OF DATA. BOLD FONT INDICATES THE BEST PERFORMANCE ACROSS THE AVERAGE SCORES

Model	Pretrain	Conversion pair		932 training utterances				250 training utterances				80 training utterances				
		Source	Target	MCD	FORMSE	CER	WER	MCD	FORMSE	CER	WER	MCD	FORMSE	CER	WER	
VTN	None	clb(F)	slt(F)	6.60	24.54	12.4	20.2	7.43	25.37	29.2	42.3	8.23	26.65	65.3	87.6	
			rms(M)	6.83	24.02	21.4	33.0	7.83	24.86	53.2	73.3	8.68	27.20	71.8	94.8	
		bdl(M)	slt(F)	7.33	23.36	23.1	33.7	8.31	23.40	52.9	75.9	8.74	24.81	73.9	95.7	
			rms(M)	7.37	22.68	28.4	43.3	8.30	23.78	56.2	78.4	9.14	24.38	79.0	102.8	
	Average				7.03	23.65	21.3	32.6	7.97	24.35	47.9	67.5	8.70	25.76	72.5	95.2
	TTS	clb(F)	slt(F)	6.02	23.94	5.5	9.1	6.41	24.86	5.2	9.7	6.66	27.24	10.4	14.7	
			rms(M)	6.22	24.29	6.8	11.9	6.75	24.91	12.8	21.3	6.94	27.86	12.5	22.0	
		bdl(M)	slt(F)	6.33	22.44	5.0	7.6	6.71	23.24	4.8	8.1	7.07	23.75	9.7	13.6	
			rms(M)	6.69	22.63	7.3	12.7	7.13	23.32	11.3	18.0	7.39	24.13	17.2	26.2	
	Average				6.32	23.33	6.2	10.3	6.75	24.08	8.5	14.3	7.02	25.75	12.5	19.1
	ASR	clb(F)	slt(F)	6.11	24.03	4.8	10.9	6.84	24.78	15.9	26.0	8.28	27.13	72.1	97.6	
			rms(M)	6.22	24.15	8.1	16.0	7.08	24.89	27.2	43.2	7.93	26.57	60.2	86.2	
bdl(M)		slt(F)	6.50	22.35	5.7	11.1	7.33	23.65	26.1	39.8	8.18	24.23	58.2	80.7		
		rms(M)	6.68	22.46	9.1	15.6	7.58	23.24	32.9	51.6	8.22	24.25	59.7	82.9		
Average				6.38	23.25	6.9	13.4	7.21	24.14	25.5	40.2	8.15	25.55	62.6	86.9	
RNN	None	clb(F)	slt(F)	6.77	24.81	7.1	12.1	7.29	25.02	15.4	24.0	7.76	25.04	38.6	56.8	
			rms(M)	6.80	23.54	11.6	19.7	7.49	24.84	24.7	38.0	7.98	27.67	48.9	68.7	
		bdl(M)	slt(F)	7.45	23.37	23.4	32.6	8.06	24.53	37.1	54.4	8.44	24.40	65.6	93.8	
			rms(M)	7.62	23.96	20.0	32.4	8.25	24.32	47.2	90.2	8.52	25.13	59.7	81.5	
	Average				7.16	23.92	15.5	24.2	7.77	24.68	31.1	51.7	8.18	25.56	53.2	75.2
	TTS	clb(F)	slt(F)	6.29	24.62	5.6	10.1	6.63	23.99	7.4	12.7	6.92	26.40	14.0	22.2	
			rms(M)	6.35	23.58	8.3	16.1	6.88	24.30	17.0	27.7	7.08	26.54	29.0	44.0	
		bdl(M)	slt(F)	6.74	22.89	8.2	13.9	7.08	23.11	11.3	19.8	7.46	23.60	16.3	23.8	
			rms(M)	6.97	22.36	15.1	26.3	7.39	23.34	21.1	32.4	7.57	23.30	25.4	39.6	
	Average				6.59	23.36	9.3	16.6	7.00	23.69	14.2	23.2	7.26	24.96	21.2	32.4

TABLE II

ASR-BASED RECOGNITION RESULTS OF VTN CONVERTED SAMPLES FROM THE EVALUATION SET OF THE CLB-SLT CONVERSION PAIR. THE ERRORS ARE IN UPPERCASE

Description	Training data size	Recognition result
Ground truth	-	the history of the eighteenth century is written ernest prompted
TTS-oriented pretraining	932	the history of the eighteenth century is written IN IS prompted TO
	250	the history of the eighteenth century is written IN HIS PROMPTER
	80	the history of the eighteenth century is written ON HIS PROMPT
ASR-oriented pretraining	932	the history of the eighteenth century is written EARNEST prompted
	250	the history of the eighteenth CENTURY'S RADIANCE prompted
	80	IT DISTURBED the DAY TO HIMSELF TO REJOIN HIM IN NORTH'S LIBRARY

TABLE III

EVALUATION-SET NATURALNESS AND SIMILARITY SUBJECTIVE EVALUATION RESULTS OF VTNS WITH NO PRETRAINING, TTS PRETRAINING, ASR PRETRAINING, AND RNN-BASED MODELS WITH NO PRETRAINING AND TTS PRETRAINING, WHICH ARE AVERAGED OVER ALL CONVERSION PAIRS AND DIFFERENT SIZES OF DATA

Model	Pretraining	932 training utterances		80 training utterances	
		Naturalness	Similarity	Naturalness	Similarity
Analysis-synthesis		4.45 ± 0.14	-	-	-
VTN	None	3.19 ± 0.23	61% ± 14%	1.96 ± 0.16	44% ± 13%
	TTS	4.34 ± 0.15	80% ± 11%	4.11 ± 0.09	68% ± 8%
	ASR	4.25 ± 0.16	77% ± 10%	3.38 ± 0.20	53% ± 12%
RNN	None	2.33 ± 0.20	40% ± 12%	1.57 ± 0.14	33% ± 15%
	TTS	3.91 ± 0.19	68% ± 13%	3.71 ± 0.09	58% ± 10%

To make a clearer comparison, we applied TTS-oriented pretraining to both VTNs and RNNs. From Tables I and III, it was shown that without TTS-oriented pretraining, VTNs were less robust to training data reduction than RNNs in terms of objective measures but better in terms of subjective measures.

This is possibly because that a more complex model like VTN is capable of generating better-sounding voices while being more prone to overfitting since it lacks attention regularizations such as the location-sensitive location, as suggested in [68]. As we applied TTS-oriented pretraining to both VTN and RNN, it could

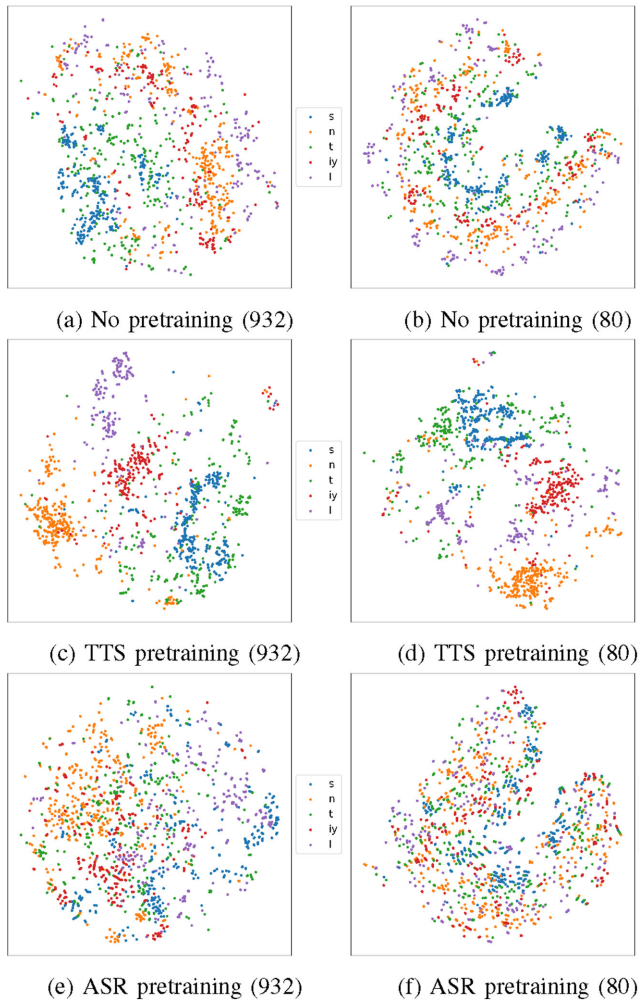


Fig. 5. Visualizations of hidden representations extracted from VTNs with no pretraining, TTS-oriented pretraining, and ASR-oriented pretraining. The validation set from *clb* was used. The numbers in the parenthesis indicate the number of training utterances.

be clearly observed that VTNs outperformed RNNs in terms of all objective measures except FORMSE and subjective scores. This is possibly due to the use of MCD as the model selection criterion.

E. Visualizing the Hidden Representation Space

In Section IV-A, we suspected that applying the TTS-oriented pretraining technique results in an encoder that can extract linguistic-information-rich representation. To demonstrate this tendency, we extracted the hidden representations with the trained encoders using the validation set from the *clb* speaker as input, and visualized them using the t-SNE method [81]. We used the phoneme labels that come with the CMU ARCTIC dataset as ground truth and colored the 5 most common phonemes and their corresponding hidden representations to simplify the plots. Note that for encoders with a reduction factor greater than 1, the corresponding label was decided with majority voting. For example, if the encoder reduction factor is 4, and the labels of the four frames that corresponds to a hidden representation are

“s,” “s,” “s,” “a,” then the label of that hidden representation will be set to “s”.

The resulting plots are shown in Fig. 5. It could be clearly observed that compared to no pretraining, the hidden representation spaces learned from TTS-oriented pretraining demonstrated a strong degree of clustering effect, where points correspond to the same phoneme were close to each other. This tendency was consistent in the cases of both 932 and 80 training utterances. On the other hand, ASR-oriented pretraining yielded a much scatter hidden representation space even with 932 training utterances.

This analysis suggests that the TTS-oriented pretraining technique can result in a more discretized representation space, which matches our initial assumption. We may further conclude that, by looking together with the objective and subjective results in Tables I and III, the degree of clustering effect somehow reflect the goodness of the hidden representations for seq2seq VC.

VI. CONCLUSION

In this work, we evaluated the pretraining techniques for addressing the problem of data efficiency in seq2seq VC. Specifically, a unified, two-stage training strategy that first pretrains the decoder and the encoder subsequently followed by initializing the VC model with the pretrained model parameters was proposed. ASR and TTS were chosen as source tasks to transfer knowledge from, and the RNN and VTN architectures were implemented. Through objective and subjective evaluations, it was shown that the TTS-oriented pretraining strategy can greatly improve the performance in terms of speech intelligibility and quality when applied to both RNNs and VTNs, and the performance could stay without significant degradation even with limited training data. As for ASR-oriented pretraining, the robustness was not so good with the reduction of training data size. Also, VTNs performed inferior to RNNs without pretraining but superior with TTS-oriented pretraining. The visualization experiment suggested that the TTS-oriented pretraining could learn a linguistic-information-rich hidden representation space while the ASR-oriented pretraining lacks such ability, which lets us imagine what an ideal hidden representation space would be like.

In the future, we plan to extend our pretraining technique to more flexible training conditions, such as many-to-many [82] or nonparallel training [37].

REFERENCES

- [1] Y. Stylianou, O. Cappe, and E. Moulines, “Continuous probabilistic transform for voice conversion,” *IEEE Speech Audio Process.*, vol. 6, no. 2, pp. 131–142, Mar. 1998.
- [2] T. Toda, A. W. Black, and K. Tokuda, “Voice conversion based on maximum-likelihood estimation of spectral parameter trajectory,” *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 15, no. 8, pp. 2222–2235, Nov. 2007.
- [3] M. Morise, F. Yokomori, and K. Ozawa, “WORLD: A vocoder-based high-quality speech synthesis system for real-time applications,” *Inst. Electron., Inf. Commun. Engineers Trans. Inf. Syst.*, 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 Commun.*, 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 *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 3104–3112.
- [6] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *Proc. Int. Conf. Learn. Representations*, 2014.
- [7] T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," in *Proc. Conf. Empirical Methods Nat. Lang. Process.*, Sep. 2015, pp. 1412–1421.
- [8] K. Tanaka, H. Kameoka, T. Kaneko, and N. Hojo, "ATTS2S-VC: Sequence-to-sequence voice conversion with attention and context preservation mechanisms," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, May 2019, pp. 6805–6809.
- [9] J. Zhang, Z. Ling, L. Liu, Y. Jiang, and L. Dai, "Sequence-to-sequence acoustic modeling for voice conversion," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 3, pp. 631–644, Mar. 2019.
- [10] W.-C. Huang, T. Hayashi, Y.-C. Wu, H. Kameoka, and T. Toda, "Voice transformer network: Sequence-to-sequence voice conversion using transformer with text-to-speech pretraining," in *Proc. Interspeech*, 2020, pp. 4676–4680.
- [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," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2019, pp. 6785–6789.
- [12] O. Russakovsky *et al.*, "ImageNet large scale visual recognition challenge," *Int. J. Comput. Vis.*, vol. 115, no. 3, pp. 211–252, 2015.
- [13] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 3320–3328.
- [14] J. Donahue *et al.*, "DeCAF: A deep convolutional activation feature for generic visual recognition," in *Proc. 31st Int. Conf. Mach. Learn.*, 2014, pp. 647–655.
- [15] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 580–587.
- [16] R. Girshick, "Fast R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 1440–1448.
- [17] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 91–99.
- [18] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 3431–3440.
- [19] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2961–2969.
- [20] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 2414–2423.
- [21] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 694–711.
- [22] M. Peters *et al.*, "Deep contextualized word representations," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics: Hum. Lang. Technol.*, Jun. 2018, pp. 2227–2237.
- [23] J. Howard and S. Ruder, "Universal language model fine-tuning for text classification," in *Proc. 56th Annu. Meeting Assoc. Comput. Linguistics*, Jul. 2018, pp. 328–339.
- [24] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, "Improving language understanding by generative pre-training," Tech. Rep., OpenAI, 2018.
- [25] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics: Hum. Lang. Technol.*, Jun. 2019, pp. 4171–4186.
- [26] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," Tech. Rep., OpenAI, 2019.
- [27] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," in *Proc. Interspeech*, 2019, pp. 3465–3469.
- [28] M. Ravanelli and Y. Bengio, "Learning speaker representations with mutual information," in *Proc. Interspeech*, 2019, pp. 1153–1157.
- [29] H.-T. Luong and J. Yamagishi, "Nautilus: A versatile voice cloning system," 2020, *arXiv:2005.11004*.
- [30] S. Karita *et al.*, "A comparative study on transformer vs RNN in speech applications," *IEEE Automat. Speech Recognit. Understanding Workshop (ASRU)*, 2019, pp. 449–456.
- [31] T. Toda *et al.*, "The voice conversion challenge 2016," in *Proc. Interspeech*, 2016, pp. 1632–1636.
- [32] J. Lorenzo-Trueba *et al.*, "The voice conversion challenge 2018: Promoting development of parallel and nonparallel methods," in *Proc. Odyssey Speaker Lang. Recognit. Workshop*, 2018, pp. 195–202.
- [33] L. Sun, K. Li, H. Wang, S. Kang, and H. Meng, "Phonetic posteriorgrams for many-to-one voice conversion without parallel data training," in *Proc. IEEE Int. Conf. Multimedia Expo.*, 2016, pp. 1–6.
- [34] G. Zhao, S. Ding, and R. Gutierrez-Osuna, "Foreign accent conversion by synthesizing speech from phonetic posteriorgrams," in *Proc. Interspeech*, 2019, pp. 2843–2847.
- [35] R. L. Leshem and R. Giryas, "Taco-vc: A single speaker tacotron based voice conversion with limited data," 2019, *arXiv:1904.03522*.
- [36] 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.
- [37] J.-X. Zhang, Z.-H. Ling, and L.-R. Dai, "Non-parallel sequence-to-sequence voice conversion with disentangled linguistic and speaker representations," *IEEE/ACM Transact. Audio, Speech, Lang. Process.*, vol. 28, pp. 540–552, 2020.
- [38] H. Kameoka, K. Tanaka, T. Kaneko, and N. Hojo, "ConvS2S-VC: Fully Convolutional Sequence-to-Sequence Voice Conversion," *IEEE/ACM Transact. Audio, Speech, Lang. Process.*, vol. 28, pp. 1849–1863, 2020.
- [39] G. E. Dahl, D. Yu, L. Deng, and A. Acero, "Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 20, no. 1, pp. 30–42, Jan. 2012.
- [40] G. Hinton *et al.*, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal Process. Mag.*, vol. 29, no. 6, pp. 82–97, Nov. 2012.
- [41] J. Chorowski, R. J. Weiss, S. Bengio, and A. van den Oord, "Unsupervised speech representation learning using wavenet autoencoders," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 12, pp. 2041–2053, Dec. 2019.
- [42] S. Pascual, M. Ravanelli, J. Serrà, A. Bonafonte, and Y. Bengio, "Learning problem-agnostic speech representations from multiple self-supervised tasks," in *Proc. Interspeech*, 2019, pp. 161–165.
- [43] A. Tjandra, B. Sisman, M. Zhang, S. Sakti, H. Li, and S. Nakamura, "VQ-VAE unsupervised unit discovery and multi-scale Code2Spec inverter for zerospeech challenge 2019," in *Proc. Interspeech*, 2019, pp. 1118–1122.
- [44] Y.-A. Chung and J. Glass, "Generative pre-training for speech with autoregressive predictive coding," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2020, pp. 3497–3501.
- [45] Y.-A. Chung and J. Glass, "Improved speech representations with multi-target autoregressive predictive coding," in *Proc. 58th Ann. Meet. Assoc. Comput. Linguist.*, 2020, pp. 2353–2358.
- [46] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," 2018, *arXiv:1807.03748*.
- [47] A. Baevski, S. Schneider, and M. Auli, "vq-wav2vec: Self-supervised learning of discrete speech representations," in *Proc. Int. Conf. Learn. Representation*, 2020.
- [48] B. Sisman, J. Yamagishi, S. King, and H. Li, "An overview of voice conversion and its challenges: From statistical modeling to deep learning," *IEEE/ACM Transact. Audio, Speech, Lang. Process.*, vol. 29, pp. 132–157, 2021.
- [49] J. Wu, Z. Wu, and L. Xie, "On the use of i-vectors and average voice model for voice conversion without parallel data," in *Proc. Asia-Pacific Signal Inf. Process. Assoc.*, 2016, pp. 1–6.
- [50] X. Tian, E. S. Chng, and H. Li, "A speaker-dependent WaveNet for voice conversion with non-parallel data," in *Proc. Interspeech*, 2019, pp. 201–205.
- [51] Y. Zhou, X. Tian, H. Xu, R. K. Das, and H. Li, "Cross-lingual voice conversion with bilingual phonetic posteriorgram and average modeling," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2019, pp. 6790–6794.
- [52] S. Liu, Y. Cao, and H. Meng, "Multi-target emotional voice conversion with neural vocoders," 2020, *arXiv:2004.03782*.
- [53] M. Zhang, B. Sisman, S. S. Rallabandi, H. Li, and L. Zhao, "Error reduction network for DBLSTM-based voice conversion," in *Proc. Asia-Pacific Signal Inf. Process. Assoc.*, 2018, pp. 823–828.
- [54] M. Zhang, X. Wang, F. Fang, H. Li, and J. Yamagishi, "Joint training framework for text-to-speech and voice conversion using multi-source tacotron and WaveNet," in *Proc. Interspeech*, 2019, pp. 1298–1302.

- [55] H. Luong and J. Yamagishi, "Bootstrapping non-parallel voice conversion from speaker-adaptive text-to-speech," in *Proc. Autom. Speech Recognit. Understanding Workshop*, 2019, pp. 200–207.
- [56] S.-W. Park, D.-Y. Kim, and M.-C. Joe, "Cotatron: Transcription-guided speech encoder for any-to-many voice conversion without parallel data," in *Proc. Interspeech*, 2020, pp. 4696–4700.
- [57] 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.
- [58] P. Narayanan, P. Chakravarty, F. Charette, and G. Puskorius, "Hierarchical sequence to sequence voice conversion with limited data," 2019, *arXiv:1907.07769*.
- [59] A. Vaswani *et al.*, "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5998–6008.
- [60] Y. Wang *et al.*, "Tacotron: Towards end-to-end speech synthesis," in *Proc. Interspeech*, 2017, pp. 4006–4010.
- [61] J. Shen *et al.*, "Natural TTS synthesis by conditioning WaveNet on MEL spectrogram predictions," in *Proc. IEEE Inter. Conf. Acoust., Speech Signal Process.*, 2018, pp. 4779–4783.
- [62] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," in *Proc. Conf. Neural Inf. Process. Syst.*, 2015, pp. 577–585.
- [63] H. Tachibana, K. Uenoyma, and S. Aihara, "Efficiently trainable text-to-speech system based on deep convolutional networks with guided attention," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2018, pp. 4784–4788.
- [64] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," 2016, *arXiv:1607.06450*.
- [65] J. Gehring, M. Auli, D. Grangier, D. Yarats, and Y. N. Dauphin, "Convolutional sequence to sequence learning," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 1243–1252.
- [66] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, "Neural speech synthesis with transformer network," in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 6706–6713.
- [67] L. Dong, S. Xu, and B. Xu, "Speech-transformer: A. no-recurrence sequence-to-sequence model for speech recognition," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2018, pp. 5884–5888.
- [68] T. Hayashi *et al.*, "ESPnet-TTS: Unified, reproducible, and integratable open source end-to-end text-to-speech toolkit," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, 2020, pp. 7654–7658.
- [69] J. Kominek and A. W. Black, "The CMU ARCTIC speech databases," in *Proc. 5th Ind. Sci. Cong. Assoc. Workshop Speech Synth.*, 2004, pp. 223–224.
- [70] Munich Artificial Intelligence Laboratories GmbH, the M-AILABS Speech Dataset." 2019, Accessed: Nov. 30, 2019. [Online]. Available: <https://www.caito.de/2019/01/the-m-ailabs-speech-dataset/>
- [71] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "LibriSpeech: An ASR corpus based on public domain audio books," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2015, pp. 5206–5210.
- [72] S. Watanabe *et al.*, "ESPnet: End-to-end speech processing toolkit," in *Proc. Interspeech*, 2018, pp. 2207–2211.
- [73] Y. You *et al.*, "Large batch optimization for deep learning: Training bert in 76 minutes," in *Proc. Int. Conf. Learn. Representations*, 2020.
- [74] R. Yamamoto, E. Song, and J. Kim, "Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, 2020, pp. 6199–6203.
- [75] A. v. d. Oord *et al.*, "Wavenet: A generative model for raw audio," 2016, *arXiv:1609.03499*.
- [76] A. Tamamori, T. Hayashi, K. Kobayashi, K. Takeda, and T. Toda, "Speaker-dependent wavenet vocoder," in *Proc. Interspeech*, 2017, pp. 1118–1122.
- [77] T. Hayashi, A. Tamamori, K. Kobayashi, K. Takeda, and T. Toda, "An investigation of multi-speaker training for WaveNet vocoder," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2017, pp. 712–718.
- [78] R. Kubichek, "Mel-cepstral distance measure for objective speech quality assessment," in *Proc. IEEE Pacific Rim Conf. Commun. Comput. Signal Process.*, 1993, pp. 125–128.
- [79] B. Naderi and R. Cutler, "An open source implementation of itu-t recommendation P. 808 with validation," 2020, *arXiv:2005.08138*.
- [80] ITU-T Recommendation P. 808, *Subjective evaluation of speech quality with a crowdsourcing approach*, Standard, 2018.
- [81] L. v. d. Maaten and G. Hinton, "Visualizing data using t-sne," *J. Mach. Learn. Res.*, vol. 9, pp. 2579–2605, Nov. 2008.

- [82] H. Kameoka, W.-C. Huang, K. Tanaka, T. Kaneko, N. Hojo, and T. Toda, "Many-to-many voice transformer network," *IEEE/ACM Transact. Audio, Speech, Lang. Process.*, pp. 1, 2020.



was the recipient of the Best Student Paper Award in ISCSLP2018. He is a Co-Organizer of the Voice Conversion Challenge 2020.




Tomoki Hayashi received the B.E. degree in engineering, the M.E. degree in information science, and the Ph.D. degree in information science from Nagoya University, Nagoya, Japan, in 2014, 2016, and 2019, respectively. His research interests include statistical speech and audio signal processing. He is currently a Postdoctoral Researcher with Nagoya University and the Chief Operating Officer of Human Dataware Laboratory Company, Ltd. He was the recipient of the IEEE SPS Japan 2020 Young Author Best Paper Award.



Yi-Chiao Wu received the B.S. and M.S. degrees in engineering from the School of Communication Engineering, National Chiao Tung University, Hsinchu, Taiwan, in 2009 and 2011, respectively. He is currently working toward the Ph.D. degree with the Graduate School of Informatics, Nagoya University, Nagoya, Japan. He was with Realtek, ASUS, and Academia Sinica for five years. His research interests include speech generation applications based on machine learning methods, such as voice conversion and speech enhancement.



Hirokazu Kameoka (Senior Member, IEEE) received the B.E., M.S., and Ph.D. degrees, from the University of Tokyo, Tokyo, Japan, in 2002, 2004, and 2007, respectively. He is currently a Senior Distinguished Researcher with NTT Communication Science Laboratories, Nippon Telegraph and Telephone Corporation, and an Adjunct Associate Professor with the National Institute of Informatics, Tokyo, Japan. From 2011 to 2016, he was an Adjunct Associate Professor with the University of Tokyo. He is the author or co-author of about 150 journal articles and peer-reviewed conference proceedings. His research interests include audio, speech, and music signal processing, and machine learning. He has been an Associate Editor for the IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH, AND LANGUAGE PROCESSING since 2015, a Member of the IEEE Audio and Acoustic Signal Processing Technical Committee since 2017, and a Member of the IEEE Machine Learning for Signal Processing Technical Committee since 2019. He was the recipient of 17 awards, including the IEEE Signal Processing Society 2008 SPS Young Author Best Paper Award.



Tomoki Toda (Senior Member, IEEE) received the B.E. degree from Nagoya University, Nagoya, Japan, in 1999 and the D.E. degree from the Nara Institute of Science and Technology (NAIST), Ikoma, Japan, in 2003. He is currently a Professor with Information Technology Center, Nagoya University. He was a Research Fellow with the Japan Society for the Promotion of Science from 2003 to 2005. From 2005 to 2011, he was an Assistant Professor, and from 2011 to 2015, was an Associate Professor with NAIST. His research interests include statistical approaches to speech, music, and environmental sound processing. He was the recipient of the IEEE SPS 2009 Young Author Best Paper Award and the 2013 EURASIP-ISCA Best Paper Award, *Speech Communication* journal.