

NAUTILUS: A Versatile Voice Cloning System

Hieu-Thi Luong , *Member, IEEE*, and Junichi Yamagishi , *Senior Member, IEEE*

Abstract—We introduce a novel speech synthesis system, called NAUTILUS, that can generate speech with a target voice either from a text input or a reference utterance of an arbitrary source speaker. By using a multi-speaker speech corpus to train all requisite encoders and decoders in the initial training stage, our system can clone unseen voices using untranscribed speech of target speakers on the basis of the backpropagation algorithm. Moreover, depending on the data circumstance of the target speaker, the cloning strategy can be adjusted to take advantage of additional data and modify the behaviors of text-to-speech (TTS) and/or voice conversion (VC) systems to accommodate the situation. We test the performance of the proposed framework by using deep convolution layers to model the encoders, decoders and WaveNet vocoder. Evaluations show that it achieves comparable quality with state-of-the-art TTS and VC systems when cloning with just five minutes of untranscribed speech. Moreover, it is demonstrated that the proposed framework has the ability to switch between TTS and VC with high speaker consistency, which will be useful for many applications.

Index Terms—Voice cloning, text-to-speech, voice conversion, speaker adaptation, neural network.

I. INTRODUCTION

SPEECH synthesis is the technology of generating speech from an input interface. In its narrow sense, speech synthesis is used to refer to text-to-speech (TTS) systems [1], which play an essential role in a spoken dialog system as a way for machine-human communication. In its broader definition, speech synthesis can refer to all kinds of speech generation interfaces like voice conversion (VC) [2], video-to-speech [3], [4], and others [5]–[7]. Recent state-of-the-art speech synthesis systems can generate speech with natural sounding quality, some of which are indistinguishable from recorded speech [8], [9]. Deep neural networks are used in various components of these speech synthesis systems. Many use sequence-to-sequence (seq2seq) models to unfold a compact phoneme sequence into acoustic features in the case of TTS [9], [10] or to handle the misalignment of acoustic sequences in the case of VC [11]–[13]. A neural

vocoder, which generates waveforms sample-by-sample [14]–[16], is also a staple of many high-quality speech-generation recipes [9], [17]. Generally speaking, the performance of deep learning approaches are high when training on a large amount of data. For speech generation models, this means that we need many hours of speech from a target speaker to train a model. This limits the ability to scale the technology to many different voices.

Besides improving the naturalness, cloning new voices with a small amount of data is also an active research topic. While there are many different approaches proposed to tackle this problem, they all share the same fundamental principle which is using an abundant corpus to compensate for the lack of data of a target speaker [18]. For neural TTS, we can fine-tune all or part of a well-trained acoustic model using transcribed speech from a target speaker [19]. For neural VC, we can pool the speech data of multiple source and target speakers and share knowledge learned from each [20]. In most of these cases, the data used for training or adaptation is either paired or labeled. However, as all acoustic characteristics of a speaker are fully contained within speech signals, we should hypothetically be able to clone voices by using untranscribed speech only, and this would greatly reduce the cost of building speech generation systems. Disentangling speaker characteristics from linguistic information and representing it as a speaker vector is hence a popular way for cloning voices [21]. Another approach is to use labels auto-generated by speaker-independent automatic speech recognition (ASR) trained on large-scale multi-speaker corpora [19]. Either way, the cloning method is usually formulated for a specific data scenario of a specific speech generation system (either TTS or VC), while a true data-efficient method should work on extremely limited data and also abundant data with or without labels.

From the perspective of voice cloning, TTS and VC can be regarded as similar systems that use different inputs for generating speech with a target voice. They share almost the same objective as well as many functional components, but they are normally treated as different systems and are modeled using vastly different frameworks. In this work, we present a novel speech generation system, called NAUTILUS, which can clone voices with a small amount of untranscribed speech and be used for both TTS and VC. It is expected to have state-of-the-art (SOTA) quality and highly consistent speaker similarity when switching between modes.¹ More importantly, this combination has the ability to clone unseen voices with a versatile strategy that could be adjusted to accommodate the data situation of the

Manuscript received May 20, 2020; revised September 3, 2020 and October 12, 2020; accepted October 15, 2020. Date of publication October 30, 2020; date of current version November 18, 2020. This work was supported in part by JST CREST VoicePersonae project under Grant JPMJCR18A6, Japan, in part by MEXT KAKENHI Grants 16H06302, 17H04687, 18H04120, 18H04112, and 18KT0051, Japan, and in part by SOKENDAI (The Graduate University for Advanced Studies), Japan. (*Corresponding author: Hieu-Thi Luong.*)

Hieu-Thi Luong is with the National Institute of Informatics, and with the Department of Informatics, SOKENDAI (The Graduate University for Advanced Studies), Tokyo 101-8340, Japan (e-mail: luonghieuthi@nii.ac.jp).

Junichi Yamagishi is with the National Institute of Informatics and with the Department of Informatics, SOKENDAI (The Graduate University for Advanced Studies), Tokyo 101-8340, Japan (e-mail: jyamagis@nii.ac.jp).

Digital Object Identifier 10.1109/TASLP.2020.3034994

¹The basis of the voice cloning method for TTS was proposed in [22] and as a proof-of-concept it was also shown that the same principle is applicable to VC in [23]. This new work builds upon the methodology and presents a SOTA unified voice cloning system for TTS and VC.

target speakers. Our experiments show that the proposed system is able to capture unique and subtle speaker characteristics such as L2 accents.

This paper is structured as follows: Section II reviews works on TTS and VC in the context of cloning voices. Section III explains the principles of our framework. Section IV gives details on the NAUTILUS system used in this paper. Section V presents experiment scenarios and their evaluations. Section VI provides further analysis and discussion. We conclude our findings in Section VII.

II. RELATED WORK ON VOICE CLONING

A. Definition of Voice Cloning

The term *voice cloning* is used to refer to a specific speaker adaptation scenario for TTS with untranscribed speech in several works [21], [24]. However in pop culture, it is loosely used to describe technology that resembles VC. In this paper, we use *voice cloning* as an umbrella term that indicates any type of system that generates speech imitating the voice of a particular speaker. The main difference between *voice cloning* and *speech synthesis* is that the former puts an emphasis on the identity of the target speaker [25], while the latter sometimes disregards this aspect for naturalness [26]. Given this definition, a voice cloning can be a TTS, a VC, or any type of speech synthesis system [4], [7]. The NAUTILUS system is designed to be expandable to other input interfaces. However, we focus on TTS and VC, which are two common speech synthesis tasks, in this work as they play an irreplaceable role in our voice cloning framework.

The performance of a voice cloning system is judged on many aspects. As a speech generation system, the naturalness and similarity to target speakers are important [9]. As a computer system, a small memory footprint [18] and fast computing time [21], [27] are desirable for practical reasons. However, the defining property of voice cloning compared with generic speech synthesis is its data efficiency as this determines its scalability [28]. While data efficiency can be interpreted as using as little data as possible [18], a better voice cloning system should not only work in a situation with extremely limited amount of data but also be able to take advantage of abundant speech data [28] when they become available regardless of the availability of transcriptions [22].

B. Training Voice Conversion System for a Target Speaker

The conventional VC approach is text-dependent, i.e., it expects training data to be parallel utterances of source and target speakers [29], [30]. As obtaining these utterances is expensive and labor-intensive work, a parallel VC system has to commonly be built with as little as five minutes of data from a speaker [31]. This is inconvenient and it limits the quality of VC systems in general. Many have worked on methodologies for building VC systems with non-parallel utterances [32]. With HMM models, we can formulate a transformation function to adapt pretrained models using non-parallel speech [33], [34]. With recent deep representation learning approaches, the popular method for non-parallel VC is training a speaker-disentangled linguistic representation either implicitly or explicitly. For implicit cases, Hsu

et al. [35] used variational auto-encoder (VAE), while Kameoka *et al.* [32] used generative adversarial network (GAN) to train a *many-to-many* non-parallel VC system. These methods use multi-speaker data, conditional labels, and various regularizations to encourage a model to disentangle linguistic content from speaker characteristics via a self-supervised training process. For explicit cases, Sun *et al.* [36] used phonetic posteriorgrams (PPG) obtained from an ASR model to train an *any-to-one* non-parallel VC system. As the ASR model is speaker-independent, a PPG-based VC system can theoretically convert the speech of arbitrary source speakers into a target speaker. As PPG is a stand-in for text, the adaptation techniques used for TTS such as adapting an average [37] or a multi-speaker [38] acoustic model can also be applied for PPG-based VC systems [20].

Even though a typical VC system is only trained on speech data, recent studies have suggested that using transcriptions of training data [13], [39] or jointly training TTS along with VC [40] can further improve naturalness of the generated speech. In our previous work [23], we established a methodology to bootstrap VC from TTS by utilizing the pretrained linguistic latent space. This paper builds upon this method by introducing an auxiliary phoneme recognition module and many new techniques to improve overall performance.

C. Adapting Text-to-Speech System to an Unseen Target

A TTS system is typically trained on dozens of hours of transcribed speech [9], [41]. Due to the high requirement for quantity and quality, a professional voice actor is commonly commissioned to record such data in a controlled environment. This makes the conventional approach ill-fitted for the voice cloning task in which we do not have controls over target speaker, recording environment, or the amount of data. To build a TTS system for speakers with a limited amount of labeled data, we can adapt a pretrained model. The initial model can be trained on the data of a single speaker [42] or data pooled from multiple speakers [37], [43]. This simple fine-tuning produces a high-quality model when the data of target speakers is sufficient (e.g., one hour) [28]. When the data is extremely limited (e.g., one minute), we can restrict the tuning to certain components instead of the entire network to prevent overfitting [28], [43], [44]. In summary, speaker adaptation transfers knowledge learned from abundant data of one or multiple speakers to reduce demand on a target.

The costly part of the voice cloning system is the data collecting process, especially the transcription of speech. Theoretically speaking, as speaker characteristics are self-contained within an utterance we should be able to clone voices without using text. One practical approach is obtaining automatically annotated transcriptions using a SOTA ASR system [19]. However ASR-predicted transcriptions contain incorrect annotations, which affects the performance of the adaptation. Moreover, this approach assumes that a well-trained ASR is obtainable for the target language, which makes it impractical for low-resource languages [26] or performing cross-language adaptation [24], [45]. Given the disentanglement ability of deep learning models, another approach is to train a speaker-adaptive model

conditioned on a speaker representation extracted from speech [21], [46]. The speaker representation can be an *i-vector* [47], *d-vector* [18], [48], or *x-vector* [49], which are all byproducts of speaker recognition systems. This approach has a computational advantage in that it does not involve an optimization loop [21]. However, the drawback is its limited scalability; in other words the speaker similarity seems to stop improving when more than a few seconds of speech is used [28]. The basis of our backpropagation-based unsupervised adaptation method, with high scalability, was proposed in previous publications [22], [50]. This paper tests the same method on a more elaborate and integrated speech generation system to refine the quality and speaker similarity of the target speakers.

D. TTS as Speech-Chain Component

Even though TTS and ASR, two essential modules of spoken dialog systems, are placed at the two ends of the human-machine communication interface and compliment each other, historically, they are built independently under different frameworks [1], [51]. Recent end-to-end speech models have reduced the technical difference between TTS and ASR systems and opened up the possibility of integrating them into a single ecosystem. Tjandra *et al.* [52] developed the Speech Chain model which consists of a TTS and ASR that consume each other’s output as their own inputs. Karita *et al.* [53] factorized TTS and ASR into encoders and decoders and then jointly trained them all together by putting a constraint on the common latent space. The purpose of these unified systems is combining resources and enabling semi-supervised training.

Similar to the situation with ASR, several studies have tried to combine VC with TTS [13], [40] or bootstrapping VC from TTS [23], [39], [54]. However, their focus was on leveraging a TTS-like system for VC [13], [54] or vice versa [40], [55] in a data-abundant scenario (target speakers with a reasonable amount of transcribed speech), and they disregard the data efficiency aspect as well as the application synergy between the two systems. Hypothetically speaking, given a perfect ASR system, there is no difference between TTS and VC systems. Specifically, the PPG-based VC system [36] is essentially a TTS model stacked on top of an ASR model. Polyak *et al.* [55] trained a TTS with the target voice by combining any-to-one VC and robot-voice TTS systems [55]. In this paper, we focus not only on improving the performance of TTS and VC individually but also on developing a unified system which can perform both tasks with high consistency. Such systems would be useful for many practical application scenarios.

III. VERSATILE VOICE CLONING FRAMEWORK

Our proposed system is a multimodal neural network [5], [6], that can be used as TTS [22] or VC [23]. It is not just a combination of conventional TTS and VC systems [40] but a carefully designed system that has the ability to clone unseen voices using either transcribed or untranscribed speech [22]. The core concept is to train a latent linguistic embedding (LLE) to use as a stand-in for text when transcription is difficult to obtain. The architecture of our multimodal system resembles

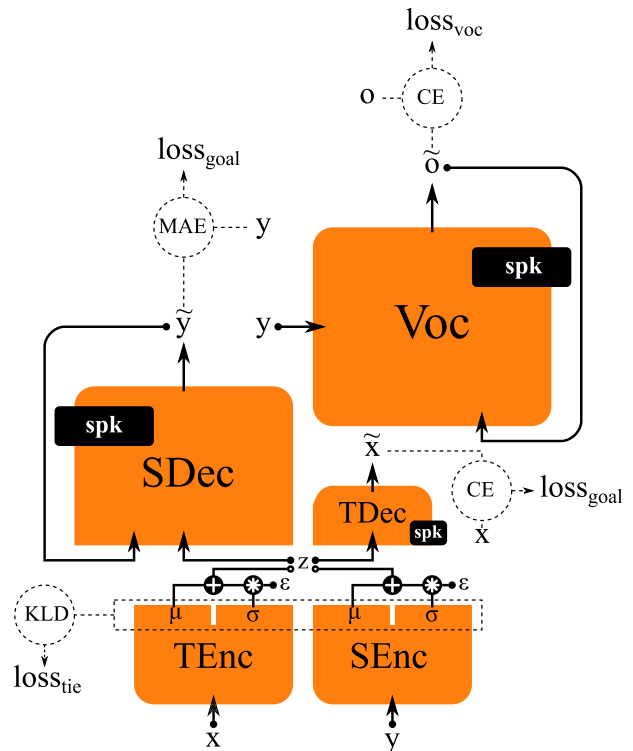


Fig. 1. The proposed system comprises of a text encoder (*TEnc*), a speech encoder (*SEnc*), a text decoder (*TDec*), a speech decoder (*SDec*), and a neural vocoder (*Voc*). Where x is a text (phoneme) representation, y is a speech (acoustic) representation, o is a waveform representation, while z is a latent linguistic embedding. \tilde{x} , \tilde{y} , and \tilde{o} are approximations of the respective representations produced by the neural networks. $loss_{goal}$ is a placeholder for either $loss_{tts}$, $loss_{sts}$, $loss_{stt}$, or $loss_{ttt}$ depending on the encoder/decoder combination. Specific to the experiments in this paper, the speech generation tasks use the mean absolute error (MAE) while the speech recognition tasks use cross entropy (CE) as a cost function. CE is also used to train the neural vocoder, while the KL divergence (KLD) is used as the latent tying loss $loss_{tie}$. The black box with the word “spk” indicates the module contains speaker-dependent components. The encoders output the mean (μ) and standard deviation (σ) of the latent features and then generate the features by using a random value (ϵ) drawn from a standard normal distribution.

the model proposed by Karita *et al.* [53]; however, they focus on the performance of ASR system instead of speaker adaptation. While the emphasis on linguistic latent features is similar to the PPG-based VC system proposed by Sun *et al.* [36], their phonetic representation extractor is trained independently with the VC model while our linguistic latent features are jointly trained with the speech generation model. Given the similarity in techniques, we will compare our system with the PPG-based VC system in the experiments.

A. Training the Text-Speech Multimodal Neural Network

The main components of the framework are presented in Fig. 1. The multimodal neural network is essential for our voice cloning methodology. While the neural vocoder is optional, we included it since it is necessary for generating high-quality speech in most recent setups [9], [17]. The proposed system contains four modules, which are encoders and decoders of either text, x , or speech, y . In combinations of encoders and

decoders, the modules can perform four transformations: text-to-speech (TTS), speech-to-speech (STS), speech-to-text (STT), and text-to-text (TTT). These modules have synergies when trained together. The speech encoder helps the TTS system adapt with untranscribed speech [22], while the text encoder helps the VC system disentangle speaker from the linguistic [23]. The text decoder is the new addition in this paper. While Karita *et al.* [53] use a similar setup for speech recognition, we focus on speech generation and the text decoder is only used as a regularizer.

Our methodology is designed around the training of a speaker-disentangled LLE, z . The LLE in our setup plays the same role as the PPG proposed for VC [35]. However, the LLE is jointly trained with the speech generation modules and contains linguistic information as a whole (instead of phoneme). There are several ways to train the multimodal neural network. It can be trained stochastically [56], step-by-step [54], or jointly [50], [53]. We proposed two methods for the joint training in our previous work [50]: 1) *joint-goal* where several losses calculated between an output inferred by each decoder and its ground-truth are combined, and 2) *tied-layer*, where the distance or distortion between two latent spaces obtained from encoders are constrained to be identical. Using one or the other is enough [22], [50], but as they are complementary, we could use them together:

$$\text{loss}_{train} = \text{loss}_{goals} + \beta \text{loss}_{tie} \quad (1)$$

Here, loss_{goals} is a weighted combination of several types of loss_{goal} , which is a placeholder for the training losses created by combining different encoders and decoders of the multimodal networks. Specifically, given the text-speech multimodal system illustrated in Fig. 1, we used the following equation as the joint-goal loss to train the initial model:

$$\text{loss}_{goals} = \text{loss}_{tts} + \alpha_{sts} \text{loss}_{sts} + \alpha_{stt} \text{loss}_{stt}, \quad (2)$$

where loss_{tts} is the TTS loss defined by the text encoder and speech decoder, it is also used as the anchor to adjust other hyperparameters; loss_{sts} is the STS loss defined by the speech encoder and speech decoder, it is de-emphasized by the weighting parameter α_{sts} ; and loss_{stt} is the STT loss defined by the speech encoder and text decoder, even though speech-to-text is not a target task, loss_{stt} is included to encourage the latent space to focus more on phonemes (but not entirely). Some other works have shown that an auxiliary phoneme classifier helps in boosting the quality of speech generation systems in general [13]. The TTT loss defined by the text encoder and text decoder, loss_{ttt} , is not included as we do not think that it helps.

In each training step, we calculate each term of the loss_{train} using a transcribed speech sample and then optimize all parameters in a supervised manner. Karita *et al.* [53] used a similar loss to jointly train their system but with one important difference, two separated speech samples, one with its transcription and another without, are used to calculate a single training loss. Specifically, loss_{tts} , loss_{stt} , and loss_{tie} are calculated using the transcribed sample, while loss_{sts} and loss_{ttt} are calculated on the untranscribed sample. This semi-supervised training strategy was proposed to take advantage of an abundant unlabeled corpus

[53]. Our system can also benefit from this semi-supervised strategy, but we only focus on supervised training in this work.

For the tied-layer loss, we calculated the symmetrized Kullback-Leibler divergence between the outputs of the text and speech encoders instead of the asymmetric one [22]:

$$\begin{aligned} \text{loss}_{tie} = & \frac{1}{2} L_{KLD}(TEnc(\mathbf{x}), SEnc(\mathbf{y})) \\ & + \frac{1}{2} L_{KLD}(SEnc(\mathbf{y}), TEnc(\mathbf{x})) \quad (3) \end{aligned}$$

The constraints help obtaining a consistent latent space between the text and speech encoders. Through experiments we found that KL divergence is an effective tied-layer loss [22].²

As the text and speech encoders output the mean (μ) and standard deviation (σ) of the LLE similarly to the VAE network, we need to apply the reparameterization trick so that the network can be trained with the backpropagation algorithm:

$$z = \mu + \sigma \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, 1). \quad (4)$$

The same process is used in the inference step to generate an LLE sequence. As ϵ is drawn from a normal distribution, this trick can also be interpreted as a noise augmenting process, which means the text and speech decoders are trained in a denoising fashion. This, in turn, makes them robust to unseen samples, which is helpful for speaker adaptation. To push the speech generation system toward an E2E setup, we include a neural vocoder to generate a waveform from the acoustic representation instead of using a conventional vocoder. In the training stage, the neural vocoder is trained separately from the rest of the system on natural speech samples:

$$\text{loss}'_{train} = \text{loss}_{voc} \quad (5)$$

We used an autoregressive WaveNet [14] conditioned on mel-spectrogram and trained on multi-speaker corpus as the neural vocoder in this paper. However, our voice cloning procedure is applicable to any type of neural vocoder.

B. Speaker Adaptation Framework

The multimodal network trained in the previous stage is essentially a multi-speaker TTS/VC system; however our goal is to perform voice cloning for unseen speakers. The following subsections describe the cloning protocols for the unsupervised scenario, which uses untranscribed speech, and the supervised scenario, which uses transcribed speech.

1) *Cloning Voice Using Untranscribed Speech*: The core mechanism for unsupervised speaker adaptation is the same as our prior work [22], [23]; however, the detail of the executions have been updated. The voice cloning stage now contains three steps, which takes the neural vocoder into account.

Step 1 - Adaptation: This is essentially our legacy unsupervised adaptation stage [22] in which the speech decoder and neural vocoder are adapted separately. We first remove all speaker

²Karita *et al.* [53] reported that KL divergence is unstable for training. The reason for this contrast is that in their work the autoencoder-based latent space is assumed to be Gaussian distribution while in our case it is forced to be an isotropic Gaussian distribution through VAE-like structure [57].

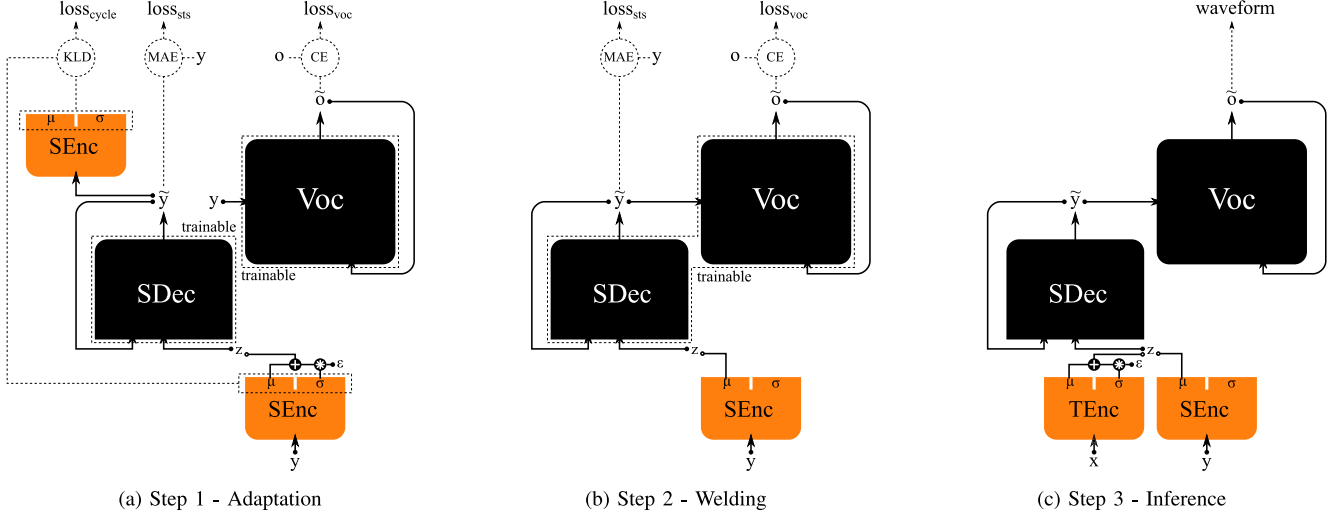


Fig. 2. Cloning procedure with untranscribed speech of the target speaker. The black background indicates modules that were or will be adapted to target speaker’s data, while the orange background indicates modules that were trained on a multi-speaker corpus in the training stage, and are, supposedly, universal. The trainable modules in each step are indicated by the dashed border with the word “trainable” near it. In the *welding* and *inference* step, the mean-value LLE tactic is applied to the speech encoder by assigning zero to the ϵ instead of sampling from a normal distribution.

components and then fine-tune the remaining parameters of the speech decoder using the following loss:

$$\text{loss}_{adapt} = \text{loss}_{sts} + \beta \text{loss}_{cycle} \quad (6)$$

The speech distortion loss loss_{sts} by itself is enough for the adaptation [22], but we further add a linguistic cycle consistent term loss_{cycle} to try to improve the performance. loss_{cycle} is the KL divergence between LLE distributions of natural speech and reconstructed speech as follows:

$$\begin{aligned} \text{loss}_{cycle} = & \frac{1}{2} L_{KLD}(SEnc(\mathbf{y}), SEnc(\tilde{\mathbf{y}})) \\ & + \frac{1}{2} L_{KLD}(SEnc(\tilde{\mathbf{y}}), SEnc(\mathbf{y})) \end{aligned} \quad (7)$$

Even though both loss_{sts} and loss_{cycle} try to force the reconstructed features to be close to natural speech, they focus on different aspects; loss_{sts} is either l_1 or l_2 frame-based hard distortion of the acoustic features, while loss_{cycle} focuses on linguistic content with soft divergence. We adapt the neural vocoder in a similar manner using its goal loss:

$$\text{loss}'_{adapt} = \text{loss}_{voc} \quad (8)$$

As a neural vocoder depends on speech only, it can be used in an unsupervised adaptation strategy. This is a simple yet effective approach [17].

Step 2 - Welding: Even though fine-tuning the acoustic model and the neural vocoder separately can produce sufficient quality [17], there are still mismatches between the generated features and the natural features used to train the vocoder. For text-to-speech systems, Zhao *et al.* [58] fine-tuned an acoustic model with the losses propagating from a neural vocoder, while Ping *et al.* [41] jointly trained them together. For voice conversion, due to the duration mismatch between source and target

utterances, Huang *et al.* proposed that the WaveNet vocoder be fine-tuned by using reconstructed acoustic features of a target speaker [59]. Motivated by them, we deploy a “welding” strategy, illustrated in Fig. 2b, that conducts fine-tuning by using the reconstructed features of the target speaker in a similar way to Huang’s approach [59], but, for both the speech decoder and neural vocoder like Ping’s method [41] based on the loss function below:

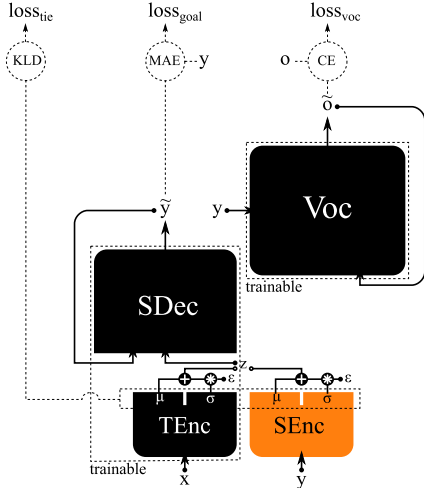
$$\text{loss}_{weld} = \text{loss}_{sts} + \gamma \text{loss}_{voc}, \quad (9)$$

where loss_{sts} is included to preserve the acoustic space even after the welding process as the speech decoder is assumed to be autoregressive in the domain.

Two practical tactics are further introduced for this step. 1) *mean-value LLE*: to let the acoustic model learn fine-grained details, we remove the sampling process from the speech encoder and use the mean value instead. 2) *mix-in*: as losses propagating from the neural vocoder can overpower the speech decoder [58], we propose a mix-in tactic, inspired by drop-out, to ease this problem. Specifically the output of the speech decoder is randomly mixed with natural frames by a percentage to reduce the amount of losses propagated back and act as a form a regularization to prevent overfitting to the generated frames.

Step 3 - Inference: Even though we use the speech encoder to tune the speech decoder and neural vocoder in the adaptation and welding steps, the text encoder can utilize these tuned modules without any further adjustment in inference (See Fig. 2c) thanks to the consistency between the latent spaces of the text and speech encoders. As our cloning method tunes entire modules, the more data available, the better the performance.

2) *Alternative Strategy to Cloning voices With Transcribed Speech:* The strategy for supervised speaker adaptation using transcribed speech was also refined compared with our previous work [22]. Instead of using exactly the same strategy as those



(a) Step 1 - Adaptation (supervised alternative)

Fig. 3. Cloning procedure with transcribed speech of a target speaker. The welding and inference steps are identical to the procedure with untranscribed speech. In this figure, $loss_{goal}$ is a placeholder for either $loss_{tts}$ or $loss_{sts}$ depending on the encoder/decoder combination.

for the above unsupervised strategy, we first tune the speech decoder and text encoder together using the transcribed speech since transcriptions could benefit the TTS system.

Step 1 - Adaptation (supervised alternative): The supervised strategy for the adaptation step is illustrated in Fig. 3a. We adapt both the speech decoder and text encoder using the following function.

$$loss''_{adapt} = loss_{tts} + \alpha loss_{sts} + \beta loss_{tie} \quad (10)$$

The optimizing loss is similar to that used in the training stage (Equation 1). We use $loss_{sts}$ and $loss_{tie}$ to maintain the linguistic latent space for VC. The *welding* and *inference* steps are the same as the unsupervised strategy.

IV. DETAILS OF NAUTILUS SYSTEM

The methodology explained in Section III can be applied to any neural architecture from the conventional acoustic model [22] to end-to-end (E2E) model [9]. Next we give the details on our system used in the experiments. It is not a fully E2E system but inspired by the E2E model in various ways.

A. Text-Speech Multimodal System

Our system is shown in Fig. 4. The text representation x is a phoneme sequence and the speech representation y is mel-spectrogram.

1) *Text Encoder*: the text encoder transforms a compact phoneme sequence x into the LLE sequence z , which has the same length as the acoustic sequence. Our specifications for the text encoder are illustrated in Fig. 4a. The input phoneme sequence is represented as one-hot vectors. As engineered linguistic features are no longer provided, *tenc-linguistic-context* is used to learn the linguistic context. This is a direct imitation of Tacotron 2 [9] but with quasi-RNN [60] used in place of the

standard RNN to speed up the training. The attention mechanism is essential in a E2E setup to unroll the phoneme sequence; our setup, however, uses an explicit duration/alignment module called “*tenc-alignment*” in training and inference to have direct control over the generated sample prosody.³ The coarse linguistic features, then, go through several dilated convolution layers called “*tenc-latent-context*” to capture the local context and smooth out the coarseness. *tenc-latent-context* has essentially the same design as the acoustic models used in our prior work [22], which used residual, skip connection and filter-gate function (Fig. 4a in [22]) to help the gradient flow:

$$h_l = \tanh(W_l^f h_{l-1} + c_l^f) \odot \sigma(W_l^g h_{l-1} + c_l^g), \quad (11)$$

where h_l is the output of the l -th layer, and W_l^f , W_l^g , c_l^f , and c_l^g are the weights and biases for filters and gates. The output of the text encoder consists of the mean and standard deviation of a text-encoded LLE sequence.

2) *Speech Decoder*: the speech decoder takes in an LLE sequence z to generate a respective acoustic sequence \tilde{y} with a particular voice. It is essentially a multi-speaker speech synthesis model and there are three components that significantly affect the performance: temporal context capturing [62], autoregressive mechanism [61], [63], and speaker modeling [44]. *sdec-context-blk* captures LLE temporal context by using time-domain convolution layers, which also contain speaker biases in their filters and gates (Fig. 4b in [22]):

$$h_l = \tanh(W_l^f h_{l-1} + c_l^f + b_l^{f,(k)}) \odot \sigma(W_l^g h_{l-1} + c_l^g + b_l^{g,(k)}), \quad (12)$$

where $b_l^{f,(k)}$ and $b_l^{g,(k)}$ are the speaker biases of k -th speaker in the training speaker pool. The effective type of speaker component depends on the network structure as well as the acoustic features [44]. We previously found that speaker biases work the best for our setup [22].

An autoregressive mechanism is introduced to improve the overall naturalness. *sdec-prenet* is responsible for the autoregressive dependency that captures the past outputs using causal layers. This is a direct imitation of the *AudioEnc* proposed by Tachinaba *et al.* [27]. The layers in *sdec-prenet* use the highway function in the same way as [27] as follows:

$$h_l^f = W_l^f h_{l-1}, \quad (13)$$

$$h_l^g = \sigma(W_l^g h_{l-1}), \quad (14)$$

$$h_l = h_l^f \odot h_l^g + h_{l-1} \odot (1 - h_l^g) \quad (15)$$

The linguistic context and the past-state token are fed into more causal layers before being transformed into the acoustic features. The architecture of the speech decoder is shown in Fig. 4b. We use the mean absolute error (MAE) as the loss function for the speech generation goals. In the adaptation stage, speaker biases are removed from the speech decoder.

³The *tenc-alignment* could be replaced with attention mechanism for convenience, and this could also potentially improve the quality further [61].

from *tenc-linguistic-context* and *sdec-prenet*, for which the rate was set at 0.5. The training was stopped when loss on validation stopped improving for ten consecutive epochs.

One hundred speakers of the VCTK corpus [66] were used to train the multi-speaker text-speech system and the WaveNet vocoder. The sampling rate was converted to the target scenarios. Among the remaining speakers, one male and one female with an American accent were used as targets for an experiment described in Section V-B. All common sentences were removed from the training so they could be used for evaluation. As VCTK lacks diversity in linguistic content, we first used 24-kHz LibriTTS corpus [67] to warm-up the text-speech network. Only *train-clean-100* and *train-clean-360* sets, which are 245 hours in total, were used to reduce the warming time. The phoneme alignments of each corpus were extracted using an ASR model trained on the same corpus using the KALDI toolkit [68]. For evaluated utterances, the model trained on the LibriTTS corpus is used to extract their phoneme alignments.

There were two voice cloning experiments, scenarios A and B. For the voice cloning stage, the number of epochs was fixed to create a uniform process. Specifically, for scenario A described in Section V-A, we first adapted the text-speech model for 256 epochs, the vocoder for 128 epochs, and then welded them together for 64 more. For scenario B described in Section V-B, the number of epochs was 256, 64, and 32, respectively. The *mix-in* rate in the welding step was set at 0.9.

For the inference stage, the speech encoder used its mean output for VC while text encoder sampled a LLE sequence from Gaussian distributions for TTS as shown in Fig. 2c. To maintain stochasticity but reduce the chance of sampling undesirable outliers, we multiplied the standard deviation output of the text encoder by 0.1 before random sampling.

D. Evaluation Measurements

We treated our system as a whole, instead of focusing on individual techniques, and we compared it with other third-party systems. For objective evaluations, we used an ASR model⁴ to calculate the word error rate (WER) of the generated speech. Note that WER should be treated as a secondary reference since it is highly sensitive to the training data of the ASR model. As a large-scale English corpus of native speakers was used to train the speech recognition model, we can interpret lower WER as indicating better pronunciation and/or greater similarity to the voices of speakers in the training set. For subjective evaluations, we used MOS on a 5-point scale for quality and DMOS on a 4-point scale for speaker similarity [31]. In most of the questions on speaker similarity, participants were asked to compare the speaker similarity of a generated utterance with a natural utterance. However, scenario A included additional questions for comparing speaker similarity between generated utterances. In scenario B, the participants were also asked to do several AB tests on quality and speaker similarity. In the AB test, two speech samples were shown at each test page and

TABLE I
TARGET SPEAKERS OF SCENARIO A

Speaker	Database	Gender	Accent	Quantity	Duration
VCC2TF1	VCC2018	female	American	81 utt.	5.2 min
VCC2TF2	VCC2018	female	American	81 utt.	5.0 min
VCC2TM1	VCC2018	male	American	81 utt.	5.2 min
VCC2TM2	VCC2018	male	American	81 utt.	5.3 min

participants were asked to choose the better of the two. These questions were used to highlight the fine-grained differences between generation systems. Each participant in our subjective listening tests was asked to do ten sessions.

V. EXPERIMENT SCENARIOS AND EVALUATIONS

As our system can clone voices by using either transcribed or untranscribed speech and can be used as TTS or VC systems, it would be difficult to evaluate all of these tasks in a single experiment. Therefore, we tested its performance and versatility under two separate scenarios. The first scenario focuses more on VC and cloning voices with untranscribed speech, while the second scenario focuses more on TTS and performance of the supervised and unsupervised speaker adaptation strategies.⁵

A. Cloning Voices Using Untranscribed Speech

In the first scenario, scenario A, we tested the ability to clone voices by using a small amount of untranscribed speech (about five minutes). A system showing good performance under this scenario is expected to have the capability to clone thousands of voices efficiently and cheaply.

1) *Experiment Setups*: we re-enacted the SPOKE task of Voice Conversion Challenge 2018 (VCC2018) [31] for this scenario. The original goal of the task was to build VC systems for 4 target English speakers (2 males and 2 females) using 81 utterances (Table I). These systems were used to convert the speech of 4 source speakers (2 males and 2 females) into each of the target voices. We followed the VCC2018 guideline [31] faithfully with one extension – we evaluated TTS systems as well as VC systems at the same time. These TTS systems were required to train on the untranscribed speech of the target speakers. In the inference stage, transcriptions of source utterances were used to generate speech with TTS systems. As there were only 35 unique sentences, we generated each sentence twice. In summary, each TTS system produced 70 utterances for each target speaker while each VC system produced 140 utterance. We split each VC system into two entities, one for same-gender conversion denoted by the superscript “=” and the other for cross-gender denoted by “×”.

2) *Evaluated Systems*: We evaluated the following TTS and VC systems in scenario A:

- **XV**: a speaker-adaptive E2E TTS system using the x-vector [18], [21], [49]. XV was used as a third-party unsupervised TTS baseline. We used the *libritts.tacotron2.v1* model and the speaker-independent WaveNet vocoder *libritts.wavenet.mol.v1* which were trained on the LibriTTS

⁴A chain system based on TDNN-F pretrained on the Librispeech corpus [69] was used for the calculation (<http://kaldi-asr.org/models/m13>).

⁵The generated speech samples of both experiment scenarios are available at <https://nii-yamagishilab.github.io/sample-versatile-voice-cloning/>

corpus to realize this approach. Both are available at the ESPnet [70] repository.⁶ As the x-vector is utterance-based, we randomly picked five utterances (about ten seconds) from the training pool of target speakers to extract the x-vector each time we generate an utterance.

- **N10**: the winner of the VCC2018 SPOKE task. N10 contains a PPG-based acoustic model [36] and a fine-tuned WaveNet vocoder [17]. It uses a speaker-independent ASR model trained on hundreds of hours of labeled data to extract PPG from speech. N10 clones voice without using the speech data of source speakers.
- **N13N17 (NR)**: the runners-up of the VCC2018 SPOKE task in terms of quality and similarity, respectively. To reduce the amount of systems, we treat them as one (denotes as NR) and use N13 in the quality evaluation while using N17 [71] in the similarity evaluation.
- **VCA_u**: VC mode of the NAUTILUS system which was adapted to target speaker by using the unsupervised strategy described in Section III-B1. The letter “A” as in “*any-to-one*” indicates that the model is not trained on source speakers. The word *unsupervised* means that the cloning is performed with untranscribed speech in the context of our current work. It is operated at 22.05 kHz to be compatible with the target speakers.
- **TTS_u**: TTS mode of the NAUTILUS system which was adapted by using the unsupervised strategy. As we did not train an automatic duration model, we used the duration extracted from the same-gender source speakers to generate speech from text. This means that TTS_u shares the same duration model as VCA_u[−] (and other same-gender VC systems). This reduces the difference in experimental conditions between them and allows us to make more insightful observations.
- **T00** and **S00**: natural utterances of the target and source speakers used as references, respectively.

3) *Evaluation*: Twenty-eight native English speakers participated in the subjective test for scenario A. They were asked to answer 18 quality and 22 similarity questions in each session. In summary, each system was judged 560 times for each measurement, while natural speech systems (T00 and S00) were judged 280 times. The objective and subjective evaluation results are shown in Table II and Fig. 5 with many interesting observations. a) XV had better quality but worse similarity than the runners-up of VCC2018, while it received the lowest WER for certain speakers. One possible explanation is that the utterances generated by XV had the characteristics of the speakers in LibriTTS corpus instead of those of the target speakers, which makes its utterances more compatible with ASR model trained on LibriSpeech. The subjective evaluation of the XV speech samples supports this speculation. b) Our systems had high scores in both subjective measurements. Interestingly our TTS system had a lower WER than our VC systems. c) Even though we had a lower score for quality than did N10, the similarity seemed to be higher. d) Our TTS and VC systems had highly consistent results, while there was a gap between the same-gender and cross-gender

TABLE II
WORD ERROR RATE FOR OBJECTIVE EVALUATION OF SCENARIO A

System	Target speakers (%WER)			
	VCC2TF1	VCC2TF2	VCC2TM1	VCC2TM2
XV	3.25	2.98	3.66	10.57
N10 ⁼	9.21	7.99	11.79	9.89
N10 [×]	9.62	11.52	8.67	9.21
N13 ⁼	23.31	21.68	31.57	27.37
N13 [×]	32.25	24.80	21.41	26.96
N17 ⁼	25.47	24.39	33.47	23.71
N17 [×]	38.08	31.44	35.23	25.88
VCA _u [−]	25.34	26.02	27.37	25.75
VCA _u [×]	30.62	27.51	23.71	22.63
TTS _u	7.72	8.40	6.23	7.18
	Source speakers (%WER)			
	VCC2SF3	VCC2SF4	VCC2SM3	VCC2SM4
S00	5.69	4.88	5.69	7.32

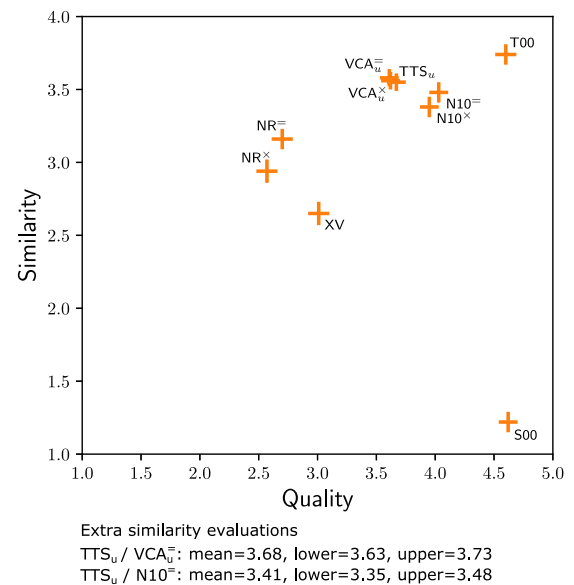


Fig. 5. Subjective results of scenario A. Lines indicate 95% confidence interval. Cross-gender and Same-gender conversion of VC systems were treated as separate entities.

subsystems of N10. The extra similarity evaluations, between the generated systems, presented in Fig. 5, shows similar results. The similarity between our TTS_u and VCA_u[−] systems was higher than the similarity between TTS_u and N10[−].

4) *Scenario Conclusion*: Even though the naturalness of our voice cloning system was slightly worse than N10 (again the best system at VCC2018), generally speaking it has achieved performance that is comparable to SOTA systems considering the difference in experimental conditions (e.g., the amount of data used in the training stage). More importantly, our system can seamlessly switch between TTS and VC modes with high consistency in terms of speaker characteristics. This is a desirable trait that would be useful for many applications.

B. Capturing Unique Speaker Characteristics

As mentioned earlier, the way voice cloning is differentiated from speech synthesis is that it should prioritize capturing the unique characteristics of target speakers. While it is easy for

⁶[Online]. Available: <https://github.com/espnet/espnet>

TABLE III
TARGET SPEAKERS OF SCENARIO B

Speaker	Database	Gender	Accent/L1	Quantity	Duration
p294	VCTK	female	American	325 utt.	11.2 min
p345	VCTK	male	American	325 utt.	11.0 min
MF6	EMIME	female	Mandarin	145 utt.	10.2 min
MM6	EMIME	male	Mandarin	145 utt.	11.3 min

TABLE IV
WORD ERROR RATE FOR OBJECTIVE EVALUATION OF SCENARIO B

System	Target speakers (%WER)			
	VCTK-p294	VCTK-p345	EMIME-MF6	EMIME-MM6
NAT*	6.09	8.69	56.24	43.39
XV	3.50	24.05	5.33	3.81
FT	13.39	20.09	57.53	42.01
VCM _u	22.22	24.05	27.70	27.09
VCM _s	23.29	24.81	29.07	29.53
TTS _u	8.37	9.74	13.39	14.92
TTS _s	9.28	10.05	36.38	38.20
Source speakers (%WER)				
	VCTK-p299	VCTK-p311	-	-
SRC**	5.64	6.51	-	-

*calculated on all training utterances of target speakers.

**calculated on natural utterances of source speakers.

listeners to grasp general global characteristics (e.g., average pitch), it is more difficult to notice local subtle traits (e.g., pronunciation of particular words) with just a single reference utterance. We could use famous individuals as targets [25], but this assumes that listeners would be familiar with them. In scenario B, we therefore used non-native speakers as targets to highlight their unique characteristics. This is convenient for subjective evaluation as native speakers can generally spot their distinctiveness without any explanation about the linguistic aspect of it [72]. In simple words, the goal of scenario B was to reproduce the accent of non-native speakers. This scenario is closely related to reducing accents [73], [74] or controlling accents [24] tasks.

1) *Experiment Setups*: the target speakers for this scenario included two American and two non-native English speakers who use Mandarin as their native language. Each speaker had about 10 minutes of speech as listed in Table III. As the base model was trained with native speakers of English, the speakers from the VCTK corpus represented the standard easy task while the speakers from the EMIME corpus [75] represented difficult and unique target speakers. The evaluated systems were required to be built with either the transcribed or untranscribed speech of the targets. Twenty common sentences from the VCTK corpus were used for the evaluations. Each sentence was generated twice by each TTS system, which totaled 40 utterances. In the case of VC, one female (p299) and one male (p311) with a general American accent included in the training pool are used as source speakers.

2) *Evaluated Systems*: The following TTS and VC systems were used for the evaluation in scenario B:

- **XV**: the same x-vector system in scenario A is reused as the unsupervised baseline of TTS.
- **FT**: a fine-tuned E2E TTS system was used as the supervised baseline. We used *ljspeech.tacotron2.v3*, implemented with ESPnet [76], as the initial model. It was

trained with 24 hours of the transcribed speech of a female speaker from the LJSpeech corpus [77]. An initial WaveNet vocoder was also trained with the same corpus. When cloning voices, we fine-tuned both acoustic and vocoder models with the transcribed speech of the targets. This system represented a simple supervised approach by fine-tuning a well-trained single speaker model [19].

- **VCM_u**: VC mode of the NAUTILUS system which was adapted to target speaker by using the unsupervised strategy described in Section III-B1. The letter “M” as in “many-to-one” indicates that the source speakers were included in the training pool of the base model. The system was operated in 24kHz.
- **VCM_s**: VC mode of the NAUTILUS system which was adapted to target speaker by using the alternative supervised strategy described in Section III-B2. The supervised strategy is expected to be more relevant to TTS, but we included its VC counterpart as an anchor for comparison.
- **TTS_u**: TTS mode of the NAUTILUS system which was adapted by using the unsupervised strategy. The duration was extracted from the source speakers of VC. This means our TTS and VC systems share the same duration model.
- **TTS_s**: TTS mode of the NAUTILUS system which was adapted by using the alternative supervised strategy.
- **NAT**: the natural utterances of the target speakers.

3) *Evaluation*: Thirty-two native speakers took part in our subjective evaluation for scenario B. As the participants were native English speakers living in Japan and many work as English teachers, we expected that they could quickly pick up on the non-native accents. Each session had 18 quality and 18 similarity questions that contain utterances of both native and non-native speakers. Besides the standard MOS tests, we also included several AB tests in this scenario. In summary, each system was evaluated 640 times for each assessment. The objective evaluation result are listed in Table III, and the subjective evaluation results are shown in Fig. 6. Here the results of native and non-native speakers are separately shown.

For the standard case with native target speakers, the subjective results show high MOS scores for most systems as shown in Fig. 6a. The new results here are comparisons between supervised and unsupervised approaches. Comparing the XV and FT systems, which represent unsupervised and supervised TTS baselines, we see that the fine-tuned one was significantly better than the speaker embedding one as it benefited from all ten minutes of data. Similar to scenario A, XV system has better WER than FT for many targets. Among our systems, the difference between the supervised and unsupervised strategies was marginal, but they were all better than the supervised baseline FT. One hypothesis is that our approaches are less sensitive to overfitting thanks to the multi-speaker corpus, speaker factorization and denoising training while FT has a higher possibility of overfitting when using ten minutes of speech [19], [54]. These observations are also supported by AB-preference tests (See the bottom part of Fig. 6a).

For the challenging case with non-native target speakers, the subjective results revealed more interesting tendencies (Fig. 6b). This scenario not only showed the robustness of the voice




Fig. 6. Subjective evaluations of scenario B. The lines used to form the crosses in figure (a) and (b) indicate 95% confidence interval.

cloning methods but also the listeners’ behaviors. First, we can see that our systems had higher similarity scores than the TTS baselines, FT and XV. The differences between our supervised and unsupervised strategies was more profound in the non-native cases. TTS_s seemed to have higher similarity than TTS_u . Next, we see that the natural speech of the non-native speakers (NAT) had lower quality scores than their native counterpart. This would be because our native listeners perceived the “quality” of speech with strong non-native accents as low, which made the quality and similarity results in this case no longer a positive correlation. The average per-listener results for non-native NAT are plotted in Fig. 6c. A negative correlation was found even for the subjective results of the TTS baselines, FT and XV, indicating that higher-quality speech corresponded to less accented speech and hence lower speaker similarity to non-native target speakers. This highlights the pros and cons of these adaptation methods. Interestingly, WER of TTS_s was worse than that of TTS_u while the natural speech (NAT) had even worse score in the non-native case. This can be interpreted as that TTS_s produces pronunciation which is more similar to the natural speech than TTS_u , which means TTS_s is better at capturing non-standard speaker characteristics.

In summary, the proposed system had higher speaker similarity than the baseline systems. Our TTS system, in particular, benefited from the supervised strategy although the improvement was relatively small. Regarding the TTS_u and the other two VC systems that had slightly better quality than the natural speech, we suspect that this is due to the reduced/lack of accents of their generated speech. This hints at potential uses for other accent-related tasks [73].

4) *Scenario Conclusion*: The subjective results have shown that the fine-tuning approach is better at capturing unique speaker characteristics than the speaker embedding approach when data are sufficient. Our systems, in particular, achieved high performance for native speakers as well as non-native speakers. Moreover our cloning strategy can be adjusted to take advantage of the transcriptions if they are available. In the meantime, the experiment also points out the limitations of the

subjective evaluation. While the current quality and similarity questions work well for native speakers, listeners’ judgements were biased when they needed to evaluate the voices of non-native speakers.

VI. ANALYSIS AND DISCUSSION

A. Training Robust and Consistent Linguistic Latent Spaces

The linguistic latent spaces obtained in the initial training stage have critical effect on the performance of the proposed system, as the rest of the voice cloning procedure functions on the assumption that LLE is a speaker-disentangled linguistic feature. Therefore, the training of the text-speech multimodal system must be carefully designed to guarantee that objective. If we only consider the text encoder and the speech decoder, then the proposed system is just a multi-speaker TTS model which lacks the ability to adapt with untranscribed speech. By adding a speaker-independent speech encoder, we provide a backdoor for unsupervised speaker adaptation, which is the topic explored in our previous publications [22], [50]. If we only consider the speech encoder and speech decoder, then it is not much different from a VAE-based multi-speaker non-parallel VC system [35]. However, to avoid the weakness of self-supervised models, which is the dependence on regularization to shape the latent space indirectly, we jointly trained the STS stack with the text encoder and transcribed speech in a supervised fashion. This ensures that the latent spaces will contain linguistic information which in turn guarantees a high performance for VC [23].

By jointly training the text and speech encoder, we help the speech encoder to learn a speaker-disentangled representation, as it is forced to approximate the text encoder, which is speaker-independent by nature. Fig. 7a shows the training curves of the TTS and STS goals, both of which descend over time and gradually converge to each other. In practice, we have to de-emphasize $loss_{sts}$ with a weighting parameter and observe the progress of the training curves, as there is a risk that the training will focus on optimizing $loss_{sts}$ and abandon $loss_{tts}$ completely, as there is always an easy and uninteresting solution to the autoencoder task.

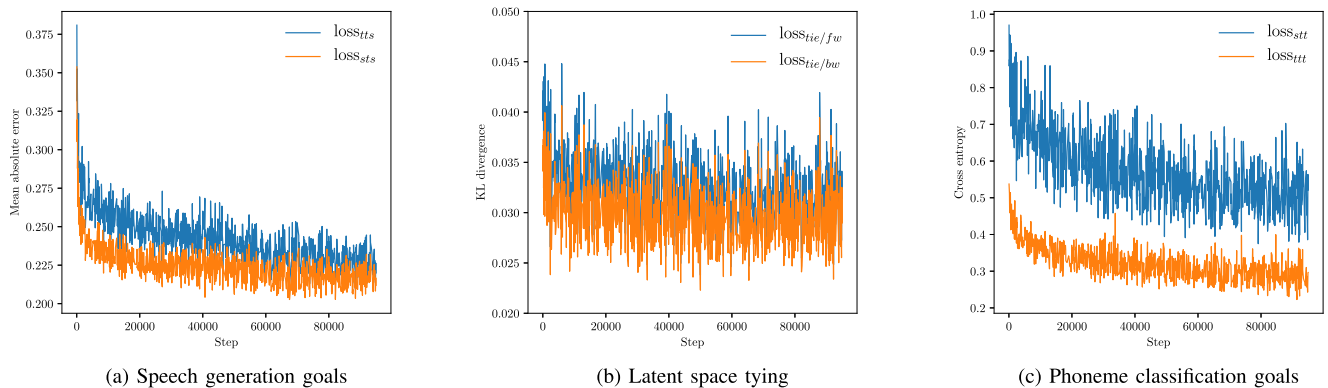


Fig. 7. Training curves of different losses available in the training stage of the text-speech multimodal system. The data points are from the training of the model used in scenario A.

TABLE V
DESCRIPTION AND WORD ERROR RATES OF EVALUATED SETUPS

Setup	Denote	TTS _u		VCM _u	
		p294	p345	p294	p345
NAUTILUS	N	8.37	9.74	22.22	24.05
N - welding	A	8.98	9.44	23.90	24.35
N - loss _{cycle} in adaptation	B	8.98	11.11	22.68	22.98
N - text decoder in training	C	12.33	12.02	25.57	25.88
N - all above (A, B & C)	D	10.35	13.85	27.55	29.38

In summary, a robust and speaker-disentangled latent linguistic representation is guaranteed by strategic placement of speaker components, joint training of the TTS and STS stacks, and use of a large-scale transcribed multi-speaker corpus. Furthermore, the tied-layer loss is used in conjunction with the join-goal losses to encourage a consistent representation between text-encoded and speech-encoded latent spaces. Fig. 7b shows the forward and backward KL divergence learning curves, which reveals that a small gap still exists between the two. Finally, the text decoder was used to force the LLE to focus more on phonetic information by adding loss_{stt} to the optimizing loss. Interestingly, even though loss_{ttt} is not optimized, it is still better than loss_{stt}, as can be seen in Fig. 7c.

B. Effects of Auxiliary Techniques on Word Error Rates

Beside the new architecture (Section IV) and the large-scale training corpus, which are the main contributors to the improved performance compared with our previous work [22], [23], the NAUTILUS system also incorporates many new auxiliary/optionally techniques. In this section, we investigate their effect on the word error rate of generated speech samples. Specifically, several slightly different setups of the proposed system were evaluated in the unsupervised speaker adaptation scenario. The experimental environment of scenario B was reused, but we only evaluated the two native English speakers, as their results are easier to interpret. Table V lists the WER of the generated utterances produced by different setups. Setup N is the unsupervised voice cloning process described in Section III-B1, A does not include the welding step, B removes loss_{cycle} from the adaptation loss, C is not trained with the text decoder, while

D removes all three elements from the procedure. In the other words, D is the most similar to the setup used in our previous publication [22], [23]. Interestingly, setup A and B have WER not much different from N, while setup C and D are significantly worse than the others even though they were all trained and adapted on the same data. These results suggest that the text decoder plays a significant role in improving the pronunciations of the generated utterances. By comparing setup C and D, we can see that the welding step and the linguistic cycle consistent also have positive impact on WER but their effects are smaller and more situational.

Having these complementary techniques at disposal can be useful for squeezing out the last bit of performance in production. If we have reliable automatic metrics, the cloning strategy can be personalized to accommodate a specific target speaker or a specific application scenario by searching for the optimal setup and hyperparameters for the particular situation, which is a topic that we will explore more in our future works. The speech samples of these setup can also be found in the accommodated web page.

C. LLE of the Supervised and Unsupervised Models

As mentioned in earlier sections, the architecture of the NAUTILUS system used in this paper is not an E2E system, which is inconvenient for practical applications but it allows us to have more control over the duration of the generated utterances. In this section, we look into the linguistic latent spaces of the adapted models to understand the behaviors of the supervised and unsupervised cloning strategies. Fig. 8 shows selected dimensions of the 64-dimensional LLE sequences generated by either the text or speech encoder of models adapted to either p294 or MF6. For each target speaker, we used speaker-independent speech encoder and an utterance not included in their adaptation data to generate a speech-encoded LLE sequence and then used either the supervised or unsupervised text encoder and the phoneme (text) and duration information of the same utterance to generate text-encoded LLE sequences. This arrangement guarantees the LLE sequences generated from the encoders are aligned,

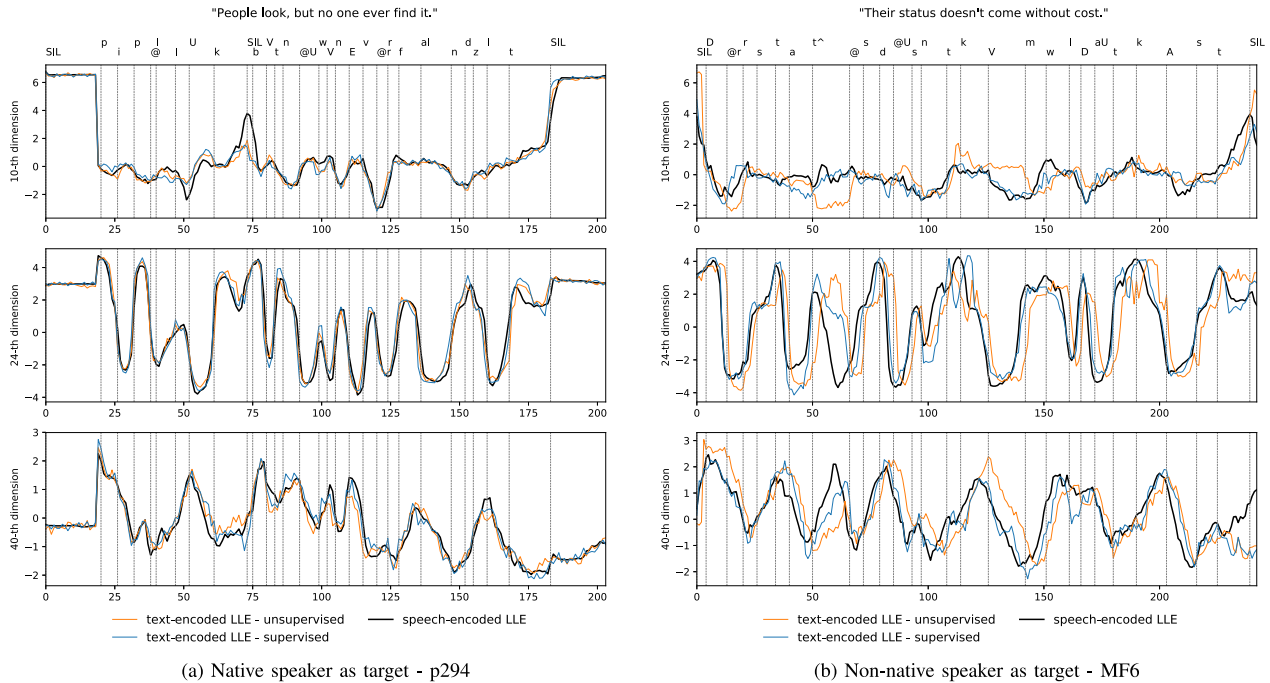


Fig. 8. Examples of 64-dimensional LLE sequences generated by the text and speech encoders of models adapted using either the supervised or unsupervised cloning strategy. An utterance of the target speaker was used to generate the speech-encoded LLE with the speech encoder, while text (phoneme) and alignment information extracted from the same utterance were used to generate the text-encoded LLE with either the supervised or unsupervised text encoder.

which helps to highlight differences between the supervised and unsupervised text encoders.

Even though we referred to the outputs of both the text and speech decoder as LLE, they actually represent slightly different concepts. The speech-encoded LLE represents the sound spoken in an utterance input, while the text-encoded LLE represents the sound that we want to generate from a symbolic phoneme input. Fig. 8a shows all three LLE sequences are well-aligned to each other in the case of p294. This suggests that the unsupervised speaker-independent text encoder was able to correctly map the symbolic phoneme to the actual spoken sound when the target is a native English speaker, which left little room for the supervised strategy to improve upon. It is expected as the text and speech encoders were initially trained on transcribed speech of a large-scale native speaker corpus. In contrast, Fig. 8b shows clear misalignments between the LLE sequences; the text-encoded LLE sequence of the supervised model seems to align to the speech-encoded LLE sequence better than its unsupervised counterpart. From this figure, we can see that the supervised strategy adjusted the text encoder to map the symbolic phoneme to the actual (wrong) sound spoken by MF6, which helps to improve the speaker similarity but degrades the quality (or pronunciation) of generated utterances. The latent spaces of the models adapted to p345 and MM6, while not presented in this paper, also show similar patterns.

VII. CONCLUSION

In this paper, we showed that our voice cloning system, “NAUTILUS”, can achieve state-of-the-art performance. More

importantly, it can act as a text-to-speech or voice conversion system with high consistency in terms of speaker characteristics when switching between the two. With the versatile cloning strategy, which can be adjusted to specific data situation of a target speaker, it is potentially useful for many other interesting tasks like accent reduction [73] or cross-lingual voice cloning [78], [79]. For future work, we will focus on evaluating our systems by using different architectures for text-speech systems [10], [54] or neural vocoders [16], [80] to solve specific voice cloning scenarios [23], [24]. Finally given the multimodal structure, extending our system to other speech generation tasks (e.g., video-to-speech [3]) would be a natural direction toward a unified voice cloning framework.

REFERENCES

- [1] K. Tokuda, T. Yoshimura, T. Masuko, T. Kobayashi, and T. Kitamura, “Speech parameter generation algorithms for HMM-based speech synthesis,” in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2000, pp. 1315–1318.
- [2] A. Kain and M. W. Macon, “Spectral voice conversion for text-to-speech synthesis,” in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 1998, pp. 285–288.
- [3] T. L. Cornu and B. Milner, “Reconstructing intelligible audio speech from visual speech features,” in *Proc. INTERSPEECH*, 2015, pp. 3355–3359.
- [4] D. Michelsanti, O. Slizovskaia, G. Haro, E. Gómez, Z.-H. Tan, and J. Jensen, “Vocoder-based speech synthesis from silent videos,” in *Proc. INTERSPEECH*, 2020, pp. 3530–3534.
- [5] K. Kinoshita, M. Delcroix, A. Ogawa, and T. Nakatani, “Text-informed speech enhancement with deep neural networks,” in *Proc. INTERSPEECH*, 2015, pp. 1760–1764.
- [6] J.-C. Hou, S.-S. Wang, Y.-H. Lai, Y. Tsao, H.-W. Chang, and H.-M. Wang, “Audio-visual speech enhancement using multimodal deep convolutional neural networks,” *IEEE Trans. Emerg. Topics Comput. Intell.*, vol. 2, no. 2, pp. 117–128, Apr. 2018.

- [7] G. Krishna, C. Tran, Y. Han, and M. Carnahan, "Speech synthesis using EEG," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 1235–1238.
- [8] Y. Wang *et al.*, "Tacotron: Towards end-to-end speech synthesis," in *Proc. INTERSPEECH*, 2017, pp. 4006–4010.
- [9] J. Shen *et al.*, "Natural TTS synthesis by conditioning WaveNet on Mel spectrogram predictions," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 4779–4783.
- [10] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, "Neural speech synthesis with transformer network," in *Proc. AAAI Conf. AI*, 2019, pp. 6706–6713.
- [11] 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.
- [12] K. Tanaka, H. Kameoka, T. Kaneko, and N. Hojo, "ATTS2S-VC: Sequence-to-sequence voice conversion with attention and context preservation mechanisms," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2019, pp. 6805–6809.
- [13] J. Zhang, Z. Ling, and L.-R. Dai, "Non-parallel sequence-to-sequence voice conversion with disentangled linguistic and speaker representations," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 28, pp. 540–552, 2019.
- [14] A. van den Oord *et al.*, "WaveNet: A generative model for raw audio," 2016, *arXiv:1609.03499*.
- [15] R. Prenger, R. Valle, and B. Catanzaro, "Waveglow: A flow-based generative network for speech synthesis," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Brighton, U.K., 2019, pp. 3617–3621.
- [16] X. Wang, S. Takaki, and J. Yamagishi, "Neural source-filter waveform models for statistical parametric speech synthesis," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 28, no. 1, pp. 402–415, Nov. 2019.
- [17] L.-J. Liu, Z.-H. Ling, Y. Jiang, M. Zhou, and L.-R. Dai, "Wavenet vocoder with limited training data for voice conversion," in *Proc. INTERSPEECH*, 2018, pp. 1983–1987.
- [18] Y. Jia *et al.*, "Transfer learning from speaker verification to multispeaker text-to-speech synthesis," in *Proc. NeurIPS*, 2018, pp. 4480–4490.
- [19] K. Inoue, S. Hara, M. Abe, T. Hayashi, R. Yamamoto, and S. Watanabe, "Semi-supervised speaker adaptation for end-to-end speech synthesis with pretrained models," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 7634–7638.
- [20] X. Tian, J. Wang, H. Xu, E. S. Chng, and H. Li, "Average modeling approach to voice conversion with non-parallel data," in *Proc. Odyssey*, 2018, pp. 227–232.
- [21] S. Arik, J. Chen, K. Peng, W. Ping, and Y. Zhou, "Neural voice cloning with a few samples," in *Proc. Neural Inf. Process. Syst.*, 2018, pp. 10 040–10 050.
- [22] H.-T. Luong and J. Yamagishi, "A unified speaker adaptation method for speech synthesis using transcribed and untranscribed speech with backpropagation," 2019, *arXiv:1906.07414*.
- [23] H.-T. Luong and J. Yamagishi, "Bootstrapping non-parallel voice conversion from speaker-adaptive text-to-speech," in *Proc. Autom. Speech Recognit. Understanding*, 2019, pp. 200–207.
- [24] Y. Zhang *et al.*, "Learning to speak fluently in a foreign language: Multilingual speech synthesis and cross-language voice cloning," *Proc. INTERSPEECH*, pp. 2080–2084, 2019.
- [25] J. Lorenzo-Trueba, F. Fang, X. Wang, I. Echizen, J. Yamagishi, and T. Kinnunen, "Can we steal your vocal identity from the internet?: Initial investigation of cloning obamas voice using gan, wavenet and low-quality found data," in *Proc. Odyssey*, 2018, pp. 240–247.
- [26] A. Gutkin, L. Ha, M. Jansche, K. Pipatsrisawat, and R. Sproat, "TTS for low resource languages: A Bangla synthesizer," in *Proc. LREC*, 2016, pp. 2005–2010.
- [27] H. Tachibana, K. Uenoyama, and S. Aihara, "Efficiently trainable text-to-speech system based on deep convolutional networks with guided attention," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 4784–4788.
- [28] Y. Chen *et al.*, "Sample efficient adaptive text-to-speech," in *Proc. ICLR*, 2019.
- [29] Y. Stylianou, O. Cappe, and E. Moulines, "Statistical methods for voice quality transformation," in *Proc. EUROSPEECH*, 1995, pp. 447–450.
- [30] 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, 2007.
- [31] J. Lorenzo-Trueba *et al.*, "The voice conversion challenge 2018: Promoting development of parallel and nonparallel methods," in *Proc. Odyssey*, 2018, pp. 195–202.
- [32] H. Kameoka, T. Kaneko, K. Tanaka, and N. Hojo, "StarGAN-VC: Non-parallel many-to-many voice conversion using star generative adversarial networks," in *IEEE Spoken Lang. Technol. Workshop*, 2018, pp. 266–273.
- [33] Y. Chen, M. Chu, E. Chang, J. Liu, and R. Liu, "Voice conversion with smoothed GMM and MAP adaptation," in *Proc. EUROSPEECH*, 2003, pp. 2413–2416.
- [34] T. Toda, Y. Ohtani, and K. Shikano, "Eigenvoice conversion based on Gaussian mixture model," in *Proc. INTERSPEECH*, 2006, pp. 2446–2449.
- [35] C.-C. Hsu, H.-T. Hwang, Y.-C. Wu, Y. Tsao, and H.-M. Wang, "Voice conversion from non-parallel corpora using variational auto-encoder," in *Proc. Asia-Pacific Signal Inf. Process. Assoc.*, 2016, pp. 1–6.
- [36] 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. Integr. Comput. Mater. Eng.*, 2016, pp. 1–6.
- [37] J. Yamagishi and T. Kobayashi, "Average-voice-based speech synthesis using HSM-based speaker adaptation and adaptive training," *IEICE Trans. Inf. Syst.*, vol. 90, no. 2, pp. 533–543, 2007.
- [38] H.-T. Luong, S. Takaki, G. E. Henter, and J. Yamagishi, "Adapting and controlling DNN-based speech synthesis using input codes," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2017, pp. 4905–4909.
- [39] D. Wang *et al.*, "End-to-end voice conversion via cross-modal knowledge distillation for dysarthric speech reconstruction," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 7744–7748.
- [40] 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.
- [41] W. Ping, K. Peng, and J. Chen, "ClariNet: Parallel wave generation in end-to-end text-to-speech," in *Proc. Int. Conf. Learn. Representations*, 2019, pp. 1–15.
- [42] Z. Huang, H. Lu, M. Lei, and Z. Yan, "Linear networks based speaker adaptation for speech synthesis," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 5319–5323.
- [43] Y. Fan, Y. Qian, F. K. Soong, and L. He, "Multi-speaker modeling and speaker adaptation for DNN-based TTS synthesis," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2015, pp. 4475–4479.
- [44] H.-T. Luong and J. Yamagishi, "Scaling and bias codes for modeling speaker-adaptive DNN-based speech synthesis systems," in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2018, pp. 610–617.
- [45] Y.-N. Chen, Y. Jiao, Y. Qian, and F. K. Soong, "State mapping for cross-language speaker adaptation in TTS," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2009, pp. 4273–4276.
- [46] S. Takaki, Y. Nishimura, and J. Yamagishi, "Unsupervised speaker adaptation for DNN-based speech synthesis using input codes," in *Proc. Asia-Pacific Signal Inf. Process. Assoc.*, 2018, pp. 649–658.
- [47] Z. Wu, P. Swietojanski, C. Veaux, S. Renals, and S. King, "A study of speaker adaptation for DNN-based speech synthesis," in *Proc. INTERSPEECH*, 2015, pp. 879–883.
- [48] R. Doddipatla, N. Braunschweiler, and R. Maia, "Speaker adaptation in DNN-based speech synthesis using d-vectors," in *Proc. INTERSPEECH*, 2017, pp. 3404–3408.
- [49] E. Cooper *et al.*, "Zero-shot multi-speaker text-to-speech with state-of-the-art neural speaker embeddings," in *Proc. ICASSP*, 2020, pp. 6184–6188.
- [50] H.-T. Luong and J. Yamagishi, "Multimodal speech synthesis architecture for unsupervised speaker adaptation," in *Proc. INTERSPEECH*, 2018, pp. 2494–2498.
- [51] M. Gales and Steve Young, *The Application of Hidden Markov Models in Speech Recognition. Foundations Trends Signal Process.*, vol. 1, no. 3, pp. 195–304, 2008.
- [52] A. Tjandra, S. Sakti, and S. Nakamura, "Listening while speaking: Speech chain by deep learning," in *Proc. Autom. Speech Recognit. Understanding*, 2017, pp. 301–308.
- [53] S. Karita, S. Watanabe, T. Iwata, M. Delcroix, A. Ogawa, and T. Nakatani, "Semi-supervised end-to-end speech recognition using text-to-speech and autoencoders," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2019, pp. 6166–6170.
- [54] 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," 2019, *arXiv:1912.06813*.
- [55] A. Polyak and L. Wolf, "Attention-based wavenet autoencoder for universal voice conversion," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2019, pp. 6800–6804.
- [56] B. Li and H. Zen, "Multi-language multi-speaker acoustic modeling for LSTM-RNN based statistical parametric speech synthesis," in *Proc. INTERSPEECH*, 2016, pp. 2468–2472.

[57] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *Proc. ICLR*, 2014, pp. 1–14.

[58] Y. Zhao, S. Takaki, H.-T. Luong, J. Yamagishi, D. Saito, and N. Minematsu, "Wasserstein GAN and waveform loss-based acoustic model training for multi-speaker text-to-speech synthesis systems using a wavenet vocoder," *IEEE Access*, vol. 6, pp. 60 478–60 488, 2018.

[59] W.-C. Huang *et al.*, "Refined wavenet vocoder for variational autoencoder based voice conversion," in *Proc. EUSIPCO*, 2019, pp. 1–5.

[60] J. Bradbury, S. Merity, C. Xiong, and R. Socher, "Quasi-recurrent neural networks," in *Proc. ICLR*, 2017, pp. 1–12.

[61] O. Watts, G. E. Henter, J. Fong, and C. Valentini-Botinhao, "Where do the improvements come from in sequence-to-sequence neural TTS?" in *Proc. SSW*, 2019, pp. 217–222.

[62] H. Zen and H. Sak, "Unidirectional long short-term memory recurrent neural network with recurrent output layer for low-latency speech synthesis," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2015, pp. 4470–4474.

[63] X. Wang, S. Takaki, and J. Yamagishi, "An autoregressive recurrent mixture density network for parametric speech synthesis," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2017, pp. 4895–4899.

[64] T. Hayashi, A. Tamamori, K. Kobayashi, K. Takeda, and T. Toda, "An investigation of multi-speaker training for wavenet vocoder," in *Proc. Autom. Speech Recognition Understanding*, 2017, pp. 712–718.

[65] K. Richmond, R. A. Clark, and S. Fitt, "Robust LTS rules with the combilex speech technology lexicon," in *Proc. INTERSPEECH*, 2009, pp. 1295–1298.

[66] C. Veaux, J. Yamagishi, and K. MacDonald, "CSTR VCTK corpus: English multi-speaker corpus for CSTR voice cloning toolkit," 2017, <http://dx.doi.org/10.7488/ds/1994>.

[67] H. Zen *et al.*, "LibriTTS: A corpus derived from librispeech for text-to-speech," in *Proc. INTERSPEECH*, 2019, pp. 1526–1530.

[68] D. Povey *et al.*, "The Kaldi speech recognition toolkit," in *Proc. Autom. Speech Recognition Understanding*, 2011.

[69] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2015, pp. 5206–5210.

[70] S. Watanabe *et al.*, "ESPnet: End-to-end speech processing toolkit," *Proc. INTERSPEECH*, pp. 2207–2211, 2018.

[71] Y.-C. Wu, P. L. Tobing, T. Hayashi, K. Kobayashi, and T. Toda, "The NU non-parallel voice conversion system for the voice conversion challenge 2018," in *Proc. Odyssey*, 2018, pp. 211–218.

[72] A. C. Janska and R. A. Clark, "Native and non-native speaker judgements on the quality of synthesized speech," in *Proc. INTERSPEECH*, 2010, pp. 1121–1124.

[73] S. Aryal and R. Gutierrez-Osuna, "Can voice conversion be used to reduce non-native accents?" in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2014, pp. 7879–7883.

[74] Y. Oshima, S. Takamichi, T. Toda, G. Neubig, S. Sakti, and S. Nakamura, "Non-native text-to-speech preserving speaker individuality based on partial correction of prosodic and phonetic characteristics," *IEICE Trans. Inf. Syst.*, vol. 99, no. 12, pp. 3132–3139, 2016.

[75] M. Wester and H. Liang, "The EMIME mandarin bilingual database," 2011. [Online]. Available: <http://hdl.handle.net/1842/4862>

[76] T. Hayashi *et al.*, "ESPnet-TTS: Unified, reproducible, and integratable open source end-to-end text-to-speech toolkit," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 7654–7658.

[77] K. Ito, "The LJ speech dataset," 2017. [Online]. Available: <https://keithito.com/LJ-Speech-Dataset/>

[78] M. Abe, K. Shikano, and H. Kuwabara, "Statistical analysis of bilingual speakers speech for cross-language voice conversion," *J. Acoust. Soc. Am.*, vol. 90, no. 1, pp. 76–82, 1991.

[79] 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.

[80] S. Mehri *et al.*, "SampleRNN: An unconditional end-to-end neural audio generation model," in *Proc. Int. Conf. Learn. Representations*, 2017, pp. 1–11.



Hieu-Thi Luong (Member, IEEE) received a Ph.D. degree from the Graduate University for Advanced Studies (SOKENDAI), Japan, in 2020 for a thesis focuses on unifying the methodology of cloning voices using text-to-speech and voice conversion systems. He received B.E and M.E degrees in computer science from Vietnam National University, Ho Chi Minh City, University of Science, Vietnam in 2014 and 2016 respectively while working on speech technology systems for Vietnamese language.

In 2017, he was awarded a Japanese Government (Monbukagakusho: MEXT) Scholarship to pursue a PhD degree in statistical speech synthesis and machine learning in Tokyo, Japan.



Junichi Yamagishi (Senior Member, IEEE) received the Ph.D. degree from the Tokyo Institute of Technology (Tokyo Tech), Meguro, Tokyo, Japan, in 2006. His Ph.D. dissertation pioneered speaker-adaptive speech synthesis. He is currently a Professor with the National Institute of Informatics, Chiyoda, Tokyo, Japan. Since 2006, he has authored or co-authored over 250 refereed papers in international journals and conferences.

Prof. Yamagishi was a recipient of the Tejima Prize as the best Ph.D. thesis of Tokyo Tech in 2007. He received the Itakura Prize from the Acoustic Society of Japan in 2010, the Kiyasu Special Industrial Achievement Award from the Information Processing Society of Japan in 2013, the Young Scientists' Prize from the Minister of Education, Science and Technology in 2014, the JSPS Prize from the Japan Society for the Promotion of Science in 2016, and the 17th DOCOMO Mobile Science Award from the Mobile Communication Fund, Japan in 2018.

He was one of the organizers for special sessions on Spoofing and Countermeasures for the Automatic Speaker Verification at INTERSPEECH 2013, the 1st/2nd/3rd ASVspoof Evaluation at INTERSPEECH 2015/2017/2019, the Voice Conversion Challenge 2016/2018 at INTERSPEECH 2016 and the ISCA Speaker Odyssey 2018, the VoicePrivacy Challenge 2020 at INTERSPEECH 2020, Deep Learning for Speech Synthesis at the IEEE Spoken Language Technology (SLT) Workshop 2018. He also served as one of technical program chair/area chair for Interspeech 2012, IEEE ASRU 2019, and IEEE SLT Workshop 2021. He also served as one of organizing committees for the ISCA Speech Synthesis Workshop 2019/2021 and the ISCA Speaker Odyssey 2020. He was an Associate Editor of the IEEE/ACM Transactions on Audio, Speech, and Language Processing, a Lead Guest Editor of the IEEE Journal of Selected Topics in Signal Processing Special Issue on Spoofing and Countermeasures for Automatic Speaker Verification, and a member of the Speech and Language Processing Technical Committee of the IEEE Signal Processing Society. He is now the Chairperson of ISCA Special Interest Group: Speech Synthesis (SynSig), a member of the Technical Committee for the Asia-Pacific Signal and Information Processing Association Multimedia Security and Forensics, a Senior Area Editor of the IEEE/ACM Transaction on Audio, Speech, and Language Processing, a member of the Advisory Committee for the Computer Speech and Language (CSL) Special Issue on Advances in Automatic Speaker Verification Anti-Spoofing, and a Guest Editor of CSL Special Issue on Voice Privacy.