
Meta-StyleSpeech : Multi-Speaker Adaptive Text-to-Speech Generation

Dongchan Min¹ Dong Bok Lee¹ Eunho Yang^{1,2} Sung Ju Hwang^{1,2}

Abstract

With rapid progress in neural text-to-speech (TTS) models, personalized speech generation is now in high demand for many applications. For practical applicability, a TTS model should generate high-quality speech with only a few audio samples from the given speaker, that are also short in length. However, existing methods either require to fine-tune the model or achieve low adaptation quality without fine-tuning. In this work, we propose StyleSpeech, a new TTS model which not only synthesizes high-quality speech but also effectively adapts to new speakers. Specifically, we propose Style-Adaptive Layer Normalization (SALN) which aligns gain and bias of the text input according to the style extracted from a reference speech audio. With SALN, our model effectively synthesizes speech in the style of the target speaker even from a single speech audio. Furthermore, to enhance StyleSpeech’s adaptation to speech from new speakers, we extend it to Meta-StyleSpeech by introducing two discriminators trained with style prototypes, and performing episodic training. The experimental results show that our models generate high-quality speech which accurately follows the speaker’s voice with single short-duration (1-3 sec) speech audio, significantly outperforming baselines.

1. Introduction

In the past few years, the fidelity and intelligibility of speech produced by neural text-to-speech (TTS) synthesis models have shown dramatic improvements. Furthermore, a number of applications such as AI voice assistant services and audio navigation systems have been actively developed and deployed to real-world, attracting increasing demand of TTS

synthesis. The majority of the TTS models aim to synthesize high quality speech of a single speaker from the given text (Oord et al., 2016; Wang et al., 2017; Shen et al., 2018; Ren et al., 2019; 2020) and have been extended to support multi speakers (Gibiensky et al., 2017; Ping et al., 2017; Chen et al., 2020).

Meanwhile, there is an increasing demand for *personalized speech generation*, which requires TTS models to generate high-quality speech that well captures the voice of the given speaker with only a few samples of the speech data. However, natural human speech is highly expressive and contains rich information, including various factors such as the speaker identity and prosody. Thus, generating personalized speech from a few speech audios, potentially even from a single audio, is an extremely challenging task. To achieve this goal, the TTS model should be able to generate speech of multiple speakers and adapt well to an unseen speaker’s voice.

A popular approach to handle this challenge is to pre-train the model on a large dataset consisting of the speech from many speakers and fine-tune the model with a few audio samples of a target speaker (Chen et al., 2019; Arik et al., 2018; Chen et al., 2021). However, this approach requires audio samples and corresponding transcripts of the target speaker to fine-tune, as well as hundreds of fine-tuning steps, which limits its applicability to real-world scenarios. Another approach is to use a piece of reference speech audio to extract a latent vector that captures the voice of the speaker such as speaker identity, prosody and speaking style (Skerry-Ryan et al., 2018; Wang et al., 2018; Jia et al., 2018). Specifically, these models are trained to synthesize speech conditioned on a *latent style vector* extracted from the speech of the given speaker, in addition to text input. This style-based approach has shown convincing results in expressive speech generation, and is able to adapt to new speakers without fine-tuning. However, they heavily rely on the diversity of the speakers in the source dataset, and thus often show low adaptation performance on new speakers.

Meta-learning (Thrun & Pratt, 1998), or learning to learn, has attracted attention of many researchers recently, as it allows the trained model to rapidly adapt to new tasks only with a few examples. In this regard, some of the few-shot generative models have utilized meta-learning for improved

¹Graduate School of AI, Korea Advanced Institute of Science and Technology (KAIST), Seoul, South Korea ²AITRICS, Seoul, South Korea. Correspondence to: Dongchan Min <alsehdc95@kaist.ac.kr>, Sung Ju Hwang <sjhwang82@kaist.ac.kr>.

generalization (Rezende et al., 2016; Bartunov & Vetrov, 2018; Clouâtre & Demers, 2019). Closely related to few-shot generation, few-shot classification also is the most extensively studied problem for meta-learning. In particular, metric-based methods (Snell et al., 2017; Misra, 2020), which meta-learn a space where the instances from the same class are embedded closer while instances that belong to different classes are embedded farther apart, have shown to achieve high performance. However, the existing methods on few-shot generation and classification mostly targets image domains, and are not straightforwardly applicable to TTS.

To overcome these difficulties, we propose *StyleSpeech*, a high-quality and expressive multi-speaker adaptive text-to-speech generation model. Our model is inspired by Karras et al. (2019) proposed for image generation, that are shown to generate surprisingly realistic photos of human faces. Specifically, we propose *Style-Adaptive Layer Normalization* (SALN) which aligns gain and bias of the text input according to a style vector extracted from a reference speech audio. Additionally, we further propose *Meta-StyleSpeech*, which is *meta-learned* with discriminators to further improve the model’s ability to adapt to new speakers that have not been seen during training. In particular, we perform episodic training by simulating the one-shot adaptation case in each episode, while additionally training two discriminators, a style discriminator and a phoneme discriminator with adversarial loss. Furthermore, the style discriminator learns a set of style prototypes enforcing the generator to generate speech from each speaker to be embedded closer to its correct style prototype which can be a voice identity of the speaker.

Our main contributions are as follows:

- We propose *StyleSpeech*, a high-quality and expressive multi-speaker adaptive TTS model, which can flexibly synthesize speech with the style extracted from a single short-length reference speech audio.
- We extend *StyleSpeech* to *Meta-StyleSpeech* which adapts well to speech from unseen speakers, by introducing the phoneme and style discriminators with style prototypes and an episodic meta-learning algorithm
- Our proposed models achieve start-of-the-art TTS performance across multiple tasks, including multi-speaker speech generation and one-shot short-length speaker adaptation.

2. Related Work

Text-to-Speech Neural TTS models have shown a rapid progress, including WaveNet (Oord et al., 2016), DeepVoice1, 2, 3 (Arik et al., 2017; Gibiansky et al., 2017;

Ping et al., 2017), Char2Wav (Sotelo et al., 2017) and Tacotron1, 2 (Wang et al., 2017; Shen et al., 2018). These models mostly resort to autoregressive generation of mel-spectrogram, which suffers from slow inference speed and a lack of robustness (word missing and skipping). Recently, several works such as Paraneet (Peng et al., 2020) and FastSpeech1 (Ren et al., 2019) have proposed non-autoregressive TTS models to handle such issues and achieve fast inference speed and improved robustness over autoregressive models. Besides, FastSpeech2 (Ren et al., 2020) extend FastSpeech1 by using additional pre-obtained acoustic features such as pitch and energy so that they show more expressive speech generation. However, these models only support a single speaker system. In this work, we base our model on FastSpeech2 and we propose an additional component to generate various voice of multi speakers.

Speaker Adaptation As the demand for personalized speech generation have increased, adaptation of the TTS models to new speakers has been extensively studied. A popular approach is to train the model on a large multi-speakers dataset and then fine-tune the whole model (Chen et al., 2019; Arik et al., 2018) or only parts of the model (Moss et al., 2020; Zhang et al., 2020; Chen et al., 2021). As an alternative approach, some of recent works attempted to model the style directly from a speech audio sample. For example, Nachmani et al. (2018) extend VoiceLoop (Taigman et al., 2018) with an additional speaker encoder. Tacotron-based approaches (Skerry-Ryan et al., 2018; Wang et al., 2018) use a piece of reference speech audio to extract the style vector and synthesize speech with the style vector, in addition to text input. Moreover, GMVAE-Tacotron (Hsu et al., 2019) present a variational approach with Gaussian Mixture prior in style modeling. To certain extent, these methods appear to be able to generate speech, that are not limited to trained speakers. However, they often achieve low adaptation performance on unseen speakers, especially when the reference speech is short in length. To handle this issue, we introduce a novel method to flexibly synthesize speech with the style vector and propose meta-learning to improve adaptation performance of our model on unseen speakers.

Meta Learning In recent years, diverse meta-learning algorithms have been proposed, mostly focusing on the few-shot classification problem. Among them, metric-based meta-learning (Vinyals et al., 2016; Snell et al., 2017; Ore-shkin et al., 2018) aims to learn the embedding space where the instances of the same class are embedded closer, while the instances belonging to different classes are embedded farther apart. On the other hand, some of existing works have studied the problem of few-shot generation. For one-shot image generalization, Rezende et al. (2016) propose sequential generative models that are built on the principles of

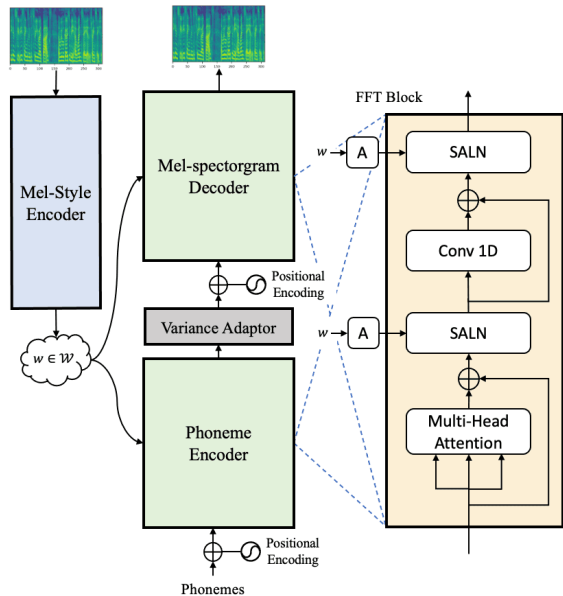


Figure 1. **The architecture of StyleSpeech.** The mel-style encoder extracts the style vector from a reference speech sample, and the generator converts the phoneme sequence into speech of various voices through the Style-Adaptive Layer Normalization (SALN).

feedback and the attention mechanism. Bartunov & Vetrov (2018) developed a hierarchical Variational Autoencoder for few-shot image generation, while Reed et al. (2018) propose the extension of Pixel-CNN with neural attention for few-shot auto-regressive density modeling. However, all these methods focus on the image domain. In contrast, our work focuses on the few-shot, short-length adaptation of TTS models, which has been relatively overlooked.

3. StyleSpeech

In this section, we first describe the architecture of StyleSpeech for multi-speaker speech generation. StyleSpeech is comprised of a mel-style encoder and a generator. The overall architecture of StyleSpeech is shown in Figure 1.

3.1. Mel Style Encoder

The *mel-style encoder*, Enc_s , takes a reference speech X as input. The goal of the mel-style encoder is to extract a vector $w \in \mathbb{R}^N$ which contains the style such as speaker identity and prosody of given speech X . Similar to Arik et al. (2018), we design the mel-style encoder to comprise of the following three parts:

- 1) Spectral processing: We first input the mel-spectrogram into fully-connected layers to transform each frames of mel-spectrogram into hidden sequences.
- 2) Temporal processing: We then use gated CNNs (Dauphin

et al., 2017) with residual connection to capture the sequential information from the given speech.

- 3) Multi-head self-attention: Then we apply a multi-head self-attention with residual connection to encode the global information. In contrast to Arik et al. (2018) where the multi-head self attention is applied across audio samples, we apply it at the frame level so that the mel-style encoder can better extract style information even from a short speech sample. Then, we temporally average the output of the self-attention to get a one-dimensional style vector w .

3.2. Generator

The *generator*, G , aims to generate a speech \tilde{X} given a phoneme (or text) sequence t and a style vector w . We build the base generator architecture upon FastSpeech2 (Ren et al., 2020), which is one of the most popular single-speaker models in non-autoregressive TTS. The model consists of three parts; a phoneme encoder, a mel-spectrogram decoder and a variance adaptor. The phoneme encoder converts a sequence of phoneme embedding into a hidden phoneme sequence. Then, the variance adaptor predicts different variances in the speech such as pitch and energy in phoneme-level¹. Furthermore, the variance adaptor predict a duration of each phonemes to regulate the length of the hidden phoneme sequence into the length of speech frames. Finally, the mel-spectrogram decoder converts the length-regulated phoneme hidden sequence into mel-spectrogram sequence. Both the phoneme encoder and mel-spectrogram decoder are composed of Feed-Forward Transformer blocks (FFT blocks) based on the Transformer (Vaswani et al., 2017) architecture. However, this model does not generate speech with diverse speakers, and thus we propose a novel component to support multi-speaker speech generation, in the following paragraph.

Style-Adaptive Layer Norm Conventionally, the style vector is provided to the generator simply through either the concatenation or the summation with the encoder output or the decoder input. In contrast, we apply an alternative approach by proposing the *Style-Adaptive Layer Normalization* (SALN). SALN receives the style vector w and predicts the gain and bias of the input feature vector. More precisely, given feature vector $\mathbf{h} = (h_1, h_2, \dots, h_H)$ where H is the dimensionality of the vector, we derive the normalized vector $\mathbf{y} = (y_1, y_2, \dots, y_H)$ as follows:

$$\mathbf{y} = \frac{\mathbf{h} - \mu}{\sigma} \quad (1)$$

$$\mu = \frac{1}{H} \sum_{i=1}^H h_i, \quad \sigma = \sqrt{\frac{1}{H} \sum_{i=1}^H (h_i - \mu)^2}$$

¹In multi-style generation, we find that it is straightforward to predict the phoneme-level information than speech frame-level information as in FastSpeech2.

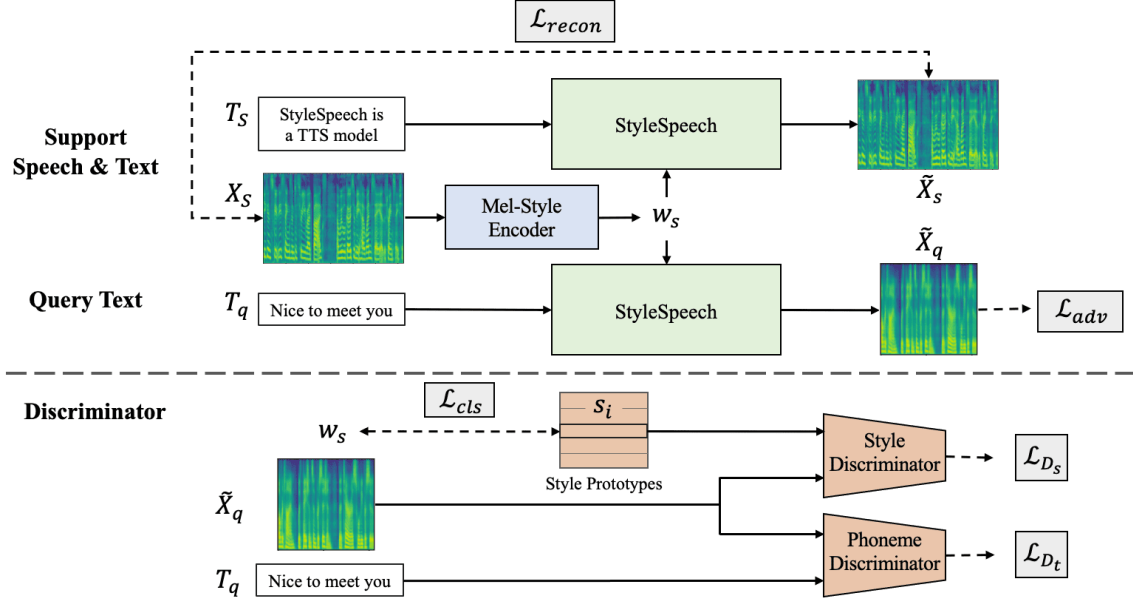


Figure 2. **Overview of Meta-StyleSpeech.** The generator use the style vector extracted from support speech and the query text to synthesize the query speech. The style discriminator learns a set of style prototypes for each speaker and enforces the generated speech to be gathered around the style prototype of the target speaker. The phoneme discriminator distinguish the real speech and generated speech condition on the input text.

Then, we compute the gain and bias with respect to the style vector w .

$$SALN(\mathbf{h}, w) = g(w) \cdot \mathbf{y} + b(w) \quad (2)$$

Unlike the fixed gain and bias as in LayerNorm (Ba et al., 2016), $g(w)$ and $b(w)$ can adaptively perform scaling and shifting of the normalized input features based on the style vector. We substitute SALN for layer normalizations in FFT blocks in the phoneme encoder and the mel-spectrogram decoder. The affine layer which convert the style vector into bias and gain is a single fully connected layer. By utilizing SALN, the generator can synthesize various styles of speech of multiple speakers given the reference audio sample in addition to the phoneme input.

3.3. Training

In the training process, both the generator and the mel-style encoder are optimized by minimizing a reconstruction loss between a mel-spectrogram synthesized by the generator and a ground truth mel-spectrogram². We use the L_1 distance as a loss function, as follows:

$$\tilde{X} = G(t, w) \quad w = Enc_s(X) \quad (3)$$

$$\mathcal{L}_{recon} = \mathbb{E} \left[\left\| \tilde{X} - X \right\|_1 \right] \quad (4)$$

²The reconstruction loss includes pitch, energy, and duration losses as in FastSpeech2, however, we do not present them for brevity.

where \tilde{X} is a generated mel-spectrogram given the phoneme input, t , and the style vector, w , which extracted from a ground truth mel-spectrogram, X .

4. Meta-StyleSpeech

Although StyleSpeech can adapt to the speech from a new speaker by utilizing SALN, it may not generalize well to the speech from an unseen speaker with a shifted distribution. Furthermore, it is difficult to generate the speech to follow the voice of the unseen speaker, especially with few speech audio samples that are also short in length. Thus, we further propose *Meta-StyleSpeech*, which is meta-learned to further improve the model’s ability to adapt to unseen speakers. In particular, we assume that only a single speech audio sample of the target speaker is available. Thus, we simulate one-shot learning for new speakers via episodic training. In each episode, we randomly sample one support (speech, text) sample, (X_s, t_s) , and one query text, t_q , from the target speaker i . Our goal is then to generate the query speech \tilde{X}_q from the query text t_q and the style vector w_s which is extracted from the support speech X_s . However, a challenge here is that we can not apply reconstruction loss on \tilde{X}_q , since no ground-truth mel-spectrogram is available. To handle this issue, we introduce an additional adversarial network with two discriminators; a style discriminator and a phoneme discriminator.

4.1. Discriminators

The *style discriminator*, D_s , predicts whether the speech follows the voice of the target speaker. The discriminator has similar architecture with mel-style encoder except it contains a set of *style prototypes* $S = \{s_i\}_{i=1}^K$, where $s_i \in \mathbb{R}^N$ denotes the style prototype for the i th speaker and K is the number of speakers in the training set. Given the style vector, $w_s \in \mathbb{R}^N$, as input, the style prototype s_i is learned with following classification loss.

$$\mathcal{L}_{cls} = -\log \frac{\exp(w_s^T s_i)}{\sum_{i'} \exp(w_s^T s_{i'})} \quad (5)$$

In detail, the dot product between the style vector and all style prototypes is computed to produce style logits, followed by cross entropy loss that encourages the style prototype to represent the target speaker’s common style such as speaker identity.

The style discriminator then maps the generated speech \tilde{X}_q to a M -dimensional vector $h(\tilde{X}_q) \in \mathbb{R}^M$ and compute a single scalar with the style prototype. The key idea here is to enforce the generated speech to be gathered around the style prototype for each speaker. In other words, the generator learns how to synthesize speech that follows the common style of the target speaker from a single short reference speech sample. Similar to the idea of Miyato & Koyama (2018), the output of the style discriminator is then computed as:

$$D_s(\tilde{X}_q, s_i) = w_0 s_i^T V h(\tilde{X}_q) + b_0 \quad (6)$$

where $V \in \mathbb{R}^{N \times M}$ is a linear layer and w_0 and b_0 are learnable parameters. The discriminator loss function of D_s then becomes

$$\mathcal{L}_{D_s} = \mathbb{E}_{t,w,s_i \sim S} [(D_s(X_s, s_i) - 1)^2 + D_s(\tilde{X}_q, s_i)^2] \quad (7)$$

The discriminator loss follows LS-GAN (Mao et al., 2017), which replace the binary cross-entropy terms of the original GAN (Goodfellow et al., 2014) objective with least squares loss functions.

The *phoneme discriminator*, D_t , takes \tilde{X}_q and t_q as inputs, and distinguishes the real speech from the generated speech given the phoneme sequence t_q as the condition. In particular, the phoneme discriminator consists of fully-connected layers and is applied at frame level. Since we know the duration of each phoneme, we can concatenate each frame in mel-spectrogram with the corresponding phoneme. The discriminator then computes scalars for each frame and averages them to get a single scalar. The final discriminator loss function for D_t then is given as:

$$\mathcal{L}_{D_t} = \mathbb{E}_{t,w} [(D_t(X_s, t_s) - 1)^2 + D_t(\tilde{X}_q, t_q)^2] \quad (8)$$

Finally, the generator loss for query speech can be defined as the sum of the adversarial loss for each discriminator, as follows:

$$\mathcal{L}_{adv} = \mathbb{E}_{t,w,s_i \sim S} [(D_s(G(t_q, w_s), s_i) - 1)^2] + \mathbb{E}_{t,w} [(D_t(G(t_q, w_s), t_q) - 1)^2]. \quad (9)$$

Furthermore, we additionally apply a reconstruction loss for the support speech, as we empirically find that it improves the quality of the generated mel-spectrograms.

$$\mathcal{L}_{recon} = \mathbb{E} [\|G(t_s, w_s) - X_s\|_1] \quad (10)$$

4.2. Episodic meta-learning

Overall, we conduct the meta-learning of Meta-StyleSpeech by alternating between the updates of the generator and mel-style encoder that minimizes \mathcal{L}_{recon} and \mathcal{L}_{adv} losses, and the updates of the discriminators that minimizes \mathcal{L}_{D_s} , \mathcal{L}_{D_t} and \mathcal{L}_{cls} losses. Thus, the final meta-training loss to minimize is defined as :

$$\mathcal{L}_G = \alpha \mathcal{L}_{recon} + \mathcal{L}_{adv} \quad (11)$$

$$\mathcal{L}_D = \mathcal{L}_{D_s} + \mathcal{L}_{D_t} + \mathcal{L}_{cls} \quad (12)$$

We set $\alpha = 10$ in our experiments.

5. Experiment

In this section, we evaluate the effectiveness of Meta-StyleSpeech and StyleSpeech on few-shot text-to-speech synthesis tasks. The audio samples are available at <https://stylespeech.github.io/>.

5.1. Experimental Setup

Datasets We train StyleSpeech and Meta-StyleSpeech on **LibriTTS** dataset (Zen et al., 2019), which is a multi-speaker English corpus derived from LibriSpeech (Panayotov et al., 2015). LibriTTS contains 110 hours audios of 1141 speakers and their corresponding text transcripts. We split the dataset into a training and a validation (test) set, and use the validation set for the evaluation on the trained speakers. For evaluation of the models’ performance on unseen speaker adaptation tasks, we use the **VCTK** (Yamagishi et al., 2019) dataset which contains audios of 108 speakers.

Preprocessing We convert the text sequences into the phoneme sequences with an open-source grapheme-to-phoneme tool³ and take the phoneme sequences as input. We downsampled an audio to 16kHz and trimmed leading and trailing silence using Librosa (McFee et al., 2015). We

³<https://github.com/Kyubyong/g2p>

Meta-StyleSpeech : Multi-Speaker Adaptive Text-to-Speech Generation

Model	MOS (\uparrow)	MCD (\downarrow)	WER (\downarrow)
GT	4.37 \pm 0.15	-	-
GT <i>mel</i> + Vocoder	4.02 \pm 0.15	-	-
DeepVoice3	2.23 \pm 0.15	4.92 \pm 0.22	36.56
GMVAE	3.28 \pm 0.20	4.81 \pm 0.24	24.49
Multi-speaker FS2(<i>vanilla</i>)	3.53 \pm 0.13	4.50 \pm 0.21	16.49
Multi-speaker FS2+ <i>d</i> -vector	3.46 \pm 0.13	4.53 \pm 0.21	16.51
StyleSpeech	3.84 \pm 0.14	4.49 \pm 0.22	16.79
Meta-StyleSpeech	3.89\pm0.12	4.29\pm 0.21	15.68

Table 1. MOS, MCD and WER for seen speakers.

Model	SMOS (\uparrow)	Sim (\uparrow)
GT	4.75 \pm 0.14	-
GT <i>mel</i> + Vocoder	4.57 \pm 0.14	-
DeepVoice3	-	0.65
GMVAE	3.17 \pm 0.06	0.76
Multi-speaker FS2(<i>vanilla</i>)	3.41 \pm 0.16	0.80
Multi-speaker FS2+ <i>d</i> -vector	2.42 \pm 0.08	0.66
StyleSpeech	3.83 \pm 0.16	0.80
Meta-StyleSpeech	4.08\pm0.15	0.81

Table 2. SMOS and Sim for seen speakers.

extract a spectrogram with a FFT size of 1024, hop size of 256, and window size of 1024 samples. Then, we convert it to a mel-spectrogram with 80 frequency bins. In addition, we average ground-truth pitch and energy by duration to get phoneme-level pitch and energy.

Implementation Details The generator in StyleSpeech uses 4 FFT blocks on both phoneme encoder and mel-spectrogram decoder following FastSpeech2. In addition, we add pre-nets to the phoneme encoder and the mel-spectrogram decoder. In detail, the encoder pre-net consists of two convolution layers and a linear layer with residual connection and the decoder pre-net consists of two linear layers. The architecture of pitch, energy and duration predictor in the variance adaptor are the same as those of FastSpeech2. However, instead of using 256 bins for pitch and energy, we use a 1D convolution layer to add real or predicted pitch and energy directly to the phoneme encoder output (La’ncucki, 2020). For the mel-style encoder, the dimensionality of all latent hidden vectors are set to 128, including the size of the style vector. Furthermore, we use Mish (Misra, 2020) activation for both the generator and the mel-style encoder.

The style discriminator has the similar architecture as the mel-style encoder, but we use 1D convolutions instead of gated CNNs. The phoneme discriminator consist of fully-connected layers. Specifically, the mel-spectrogram passes through two fully connected layers with 256 hidden dimensions before concatenating with phoneme embeddings, and then through three fully connected layers with 512 hidden dimensions and one final projection layer. We use Leaky-ReLU activation for both discriminators. In addition, We apply spectral normalization (Miyato et al., 2018) in all the

Model	MOS (\uparrow)	MCD (\downarrow)	WER (\downarrow)
GT	4.40 \pm 0.13	-	-
GT <i>mel</i> + Vocoder	4.03 \pm 0.12	-	-
GMVAE	3.15 \pm 0.17	5.54 \pm 0.26	23.86
Multi-speaker FS2(<i>vanilla</i>)	3.69 \pm 0.15	4.97 \pm 0.23	17.35
Multi-speaker FS2+ <i>d</i> -vector	3.74 \pm 0.14	5.03 \pm 0.22	17.55
StyleSpeech	3.77 \pm 0.15	5.01 \pm 0.23	17.51
Meta-StyleSpeech	3.82\pm0.15	4.95\pm0.24	16.79

Table 3. MOS, MCD and WER for unseen speakers.

layers for both discriminators except style prototypes. The more details about the architectures are in the supplementary material.

In the training process, we train StyleSpeech for 100k steps. For Meta-StyleSpeech, we start from pretrained StyleSpeech that is trained for 60k steps, and then meta-train the model for additional 40k steps via episodic training. We find that meta training from pretrained StyleSpeech helps obtain better stability in training. Furthermore, we use the predicted pitch and energy from the variance adaptor to generate both support and query speech in Meta-StyleSpeech training. In addition, we train our models with a minibatch size of 48 for StyleSpeech and 20 for Meta-StyleSpeech using the Adam optimizer. The parameters we use for the Adam optimizer are $\beta_1 = 0.9, \beta_2 = 0.98, \epsilon = 10^{-9}$. The learning rate of generator and mel-style encoder follows Vaswani et al. (2017), while the learning rate of discriminator is fixed as 0.0002. We use MelGAN (Kumar et al., 2019) as the vocoder to convert the generated mel-spectrograms into audio waveforms.

Baselines We compare our model with several baselines. **1) GT (*oracle*):** This is the Ground-Truth speech. **2) GT *mel* (*oracle*):** This is the speech synthesized by MelGAN vocoder using Ground-Truth mel-spectrogram. **3) Deep-voice3 (Ping et al., 2017):** This is a multi speaker TTS model which learns a look-up table to map embeddings for different speaker identity. Since DeepVoice3 is able to generate only seen speakers, we only compare against DeepVoice3 on seen speaker evaluation task. **4) GMVAE (Hsu et al., 2019):** This is a multi-speaker TTS model based on the Tacotron with a variational approach using Gaussian Mixture prior. **5) Multi-speaker FS2 (*vanilla*):** This is a multi-speaker FastSpeech2, which adds the style vector to the encoder output and the decoder input. The style vector is extracted by the mel-style encoder. **6) Multi-speaker FS2 + *d*-vector:** This is same as **5) Multi-speaker FS2 (*vanilla*)** except that the style vector is extracted from a pre-trained speaker verification model as suggested in Jia et al. (2018). **6) StyleSpeech:** This is our proposed model which generate multi-speaker speech from a single speech audio with the style-adaptive layer normalization and the mel-style encoder. **7) Meta-StyleSpeech:** This is our proposed model which extends StyleSpeech to perform meta training, with




Figure 3. TSNE visualization of the style vectors for unseen speakers (VCTK) with (a) and without (b) meta learning, and (c) GMVAE.

Metric	SMOS (\uparrow)				Sim (\uparrow)				Accuracy (\uparrow)				
	Length	<1 sec	1~3 sec	1 sen.	2 sen.	<1 sec	1~3 sec	1 sen.	2 sen.	<1 sec	1~3 sec	1 sen.	2 sen.
GMVAE		2.85 \pm 0.12	3.01 \pm 0.12	2.91 \pm 0.16	3.11 \pm 0.10	0.629	0.695	0.748	0.765	20.75%	30.49%	28.33%	46.15%
Multi-speaker FS2(<i>vanilla</i>)		3.14 \pm 0.17	3.63 \pm 0.16	3.31 \pm 0.14	3.36 \pm 0.12	0.713	0.735	0.775	0.773	64.80%	73.80%	72.60%	81.40%
Multi-speaker FS2+d-vector		1.85 \pm 0.12	2.08 \pm 0.16	2.11 \pm 0.16	2.12 \pm 0.14	0.601	0.603	0.619	0.616	2.40%	3.80%	5.60%	5.60%
StyleSpeech		3.32 \pm 0.16	4.13 \pm 0.16	3.50\pm0.10	3.46 \pm 0.12	0.725	0.756	0.791	0.795	77.60%	85.00%	83.46%	85.19%
Meta-StyleSpeech		3.66\pm0.13	4.19\pm0.14	3.43 \pm 0.14	3.81\pm0.12	0.738	0.779	0.813	0.815	82.60%	90.20%	88.66%	91.20%

Table 4. Adaptation performance on speech from unseen speakers with varying length of reference audios.

Metric		MCD (\downarrow)	Sim (\uparrow)	Accuracy (\uparrow)
Gender	Male	4.69 \pm 0.23	0.76	91%
	Female	4.87 \pm 0.23	0.75	89%
Accent	American	4.80 \pm 0.23	0.77	91%
	British	4.83 \pm 0.21	0.75	91%
	Indian	5.28 \pm 0.31	0.74	86%
	African	5.06 \pm 0.24	0.76	93%
	Australian	5.01 \pm 0.30	0.75	95%

Table 5. Adaptation performance depends on gender and accents.

two additional discriminators to guide the generator.

Evaluation Metric The evaluation of the TTS models is very challenging, due to its subjective nature in the evaluation of the perceptual quality of generated speech. However, we use a sufficient number of metrics to evaluate the performance of our model. Specifically, for subjective evaluation, we conduct human evaluations with **MOS** (mean opinion score) for naturalness and **SMOS** (similarity MOS) for similarity. Both metrics are rated in 1-to-5 scale and reported with the 95% confidence intervals (CI). 50 judges were participated in each experiment, where they were allowed to evaluate each audio sample once, that are presented in random orders.

In addition to subjective evaluation, we also conduct objective evaluation with quantitative measures. **MCD** evaluates the compatibility between the spectra of two audio sequences. Since the sequences are not aligned, we perform


Dynamic Time Warping to align the sequences prior to comparison. In addition, **WER** validates an intelligibility of generated speech. We use a pre-trained ASR model DeepSpeech2 (Amodei et al., 2016) to compute Word Error Rate (WER). Note that both MCD and WER are not absolute metrics for evaluating speech quality, so we only use them for relative comparisons.

Beyond the quality evaluation, we also evaluate how similar the generated speech is, to the voice of the target speaker. In particular, we use the speaker verification model based on Wan et al. (2018) to extract x-vectors from generated speech as well as the actual speech of the target speaker. We then compute cosine similarity between them and denote it as **Sim**. The Sim scores are between -1 and 1, with higher scores indicating higher similarity. Since the verification model can be used for any speakers, we use it to evaluate both the trained speakers and new speakers.


Furthermore, we use a speaker classifier that identifies which speaker an audio sample belongs to. We train the classifier on VCTK dataset and the classifier achieve 99% accuracy for validation set. We then calculate its **accuracy** to conduct the evaluation for adaptability of our model to new speakers. Better adaptation would result in higher classification accuracy.

5.2. Evaluation on Trained Speakers

Before investigating the ability of our model on unseen speaker adaptation, we first evaluate the quality of speech



(a) **Left:** A reference speech sample, **Right:** Generated mel-spectrogram



(b) Ground truth mel-spectrogram

Figure 4. Mel-spectrogram generation with a reference speech shorter than 1 second.

synthesized for seen (trained) speakers. To this end, we randomly draw one audio sample as reference speech for each 100 seen speakers in validation set of LibriTTS dataset. Then, we synthesize speech using the reference speech audio and the given text. Table 1 shows the results of MOS, MCD, and WER evaluation on different models for the LibriTTS dataset. Our StyleSpeech and Meta-StyleSpeech outperform the baseline method in all three metrics, indicating that the proposed models synthesize higher quality speech.

We also evaluate the similarity between synthesized speech for seen speakers and the reference speech. Table 2 shows the results of SMOS and Sim score on different models for the LibriTTS dataset. Our StyleSpeech variants also achieve higher similarity scores to the reference speech than the other text-to-speech baseline methods. We thus conclude that our models are more effective in style transfer from reference speech samples. Furthermore, on both experiments, Meta-StyleSpeech achieves the best performance, which shows the effectiveness of the meta-learning.

5.3. Unseen Speaker Adaptation

We now validate the adaptation performance of our model on unseen speakers. To this end, we first evaluate the quality of generated speech of unseen speakers. In this experiment, we also randomly draw one audio sample as reference for each 108 unseen speaker in the VCTK datasets and then generate speech using the reference speech and the given text. Table 3 shows the results of MOS, MCD and WER evaluation

	MCD (\downarrow) (seen)	MCD (\downarrow) (unseen)	Accuracy (\uparrow) (<1 sec)
Meta-StyleSpeech	4.29\pm0.21	4.95 \pm 0.24	82.60%
StyleSpeech	4.49 \pm 0.22	5.01 \pm 0.23	77.60%
w/o D_t	4.85 \pm 0.21	5.53 \pm .27	78.00%
w/o D_s	4.51 \pm 0.20	5.17 \pm 0.24	80.40%
w/o \mathcal{L}_{cls}	4.32 \pm 0.20	4.85\pm0.24	80.60%

Table 6. Ablation study for verifying the effectiveness of the phoneme discriminator, the style discriminator and the style prototypes.

on speech from unseen speakers. Meta-StyleSpeech also achieves the best generation quality in all three metrics, largely outperforming the baselines.

As done in the seen-speaker experiments, we also evaluate the similarity to reference speech for unseen speakers. Following Nachmani et al. (2018), we expect that the ability to adapt to new speakers may depend on the length of the reference speech audio. Thus, we perform the experiments while varying the length of the reference speech audio from unseen speakers. We use four different lengths: <1 sec, $1\sim 3$ sec, 1 sentence and 2 sentences. We set 1 sentence as a speech audio sample which is longer than 3 seconds in length and, for 2 sentences, we simply concatenate two speech audio samples.

The results of SMOS, Sim, and Accuracy are presented in Table 4. As shown in the result, our model, Meta-StyleSpeech, significantly outperforms the baseline as well as StyleSpeech on reference speech of any lengths. Specifically, it achieves high adaptation performance even with the reference speech that is shorter than 1 second. Figure 4 shows the example of generated speech from the reference speech shorter than 1 second, using Meta-StyleSpeech. We observe that our model generates high quality speech with sharp harmonics and well resolved formants which is comparable to ground-truth mel-spectrogram.

In addition, we also conduct adaptation evaluation depends on various styles such as gender and accent. The result is shown in Table 5. We can see that Meta-StyleSpeech shows balanced results for gender. Moreover, the model also shows high adaptation performance for all accents. However, we can see that slightly lower performance on Indian accent. This could be because the Indian accent have more dramatic variation in speaking style than other accents.

Visualization of style vector To better understand effectiveness of meta-learning, we visualize the style vectors. In figure 3, we demonstrate the t-SNE projection (Maaten & Hinton, 2008) of style vectors from unseen speakers from the VCTK datasets. In particular, we select 10 female speakers who have similar accents and voices which are difficult to distinguish by the model. We can see that while

StyleSpeech clearly better separates the style vectors when compared with GMVAE, Meta-StyleSpeech trained with meta-learning achieves even better clustered style vectors.

5.4. Ablation Study

We further conduct an ablation study to verify the effectiveness of each components in our model, including the text discriminator, the style discriminator and the style prototypes. In the ablation study, we use two metrics; MCD and Accuracy. The results are shown in Table 6. Both the quality of generated speech and adaptation ability of the model are significantly dropped when removing the text discriminator. This indicate that the role of the text discriminator is very important in meta-training. We also find that removing the style discriminator and the style prototypes result in the performance drop on unseen speaker adaptation. From this result, we find that the style discriminator is important in helping the model adapt to speech from unseen speakers and the style prototypes further enhance the adaptation performance of the model.

6. Conclusion

We have proposed StyleSpeech, a multi-speaker adaptive TTS model which can generate high quality and expressive speech from a single short-duration audio sample of the target speaker. In particular, we propose a style-adaptive layer normalization (SALN) to generate various styles of speech of multiple speakers. Furthermore, we extend StyleSpeech to Meta-StyleSpeech with additional discriminators and meta-learning, to improve its adaptation performance on unseen speakers. Specifically, we simulate one-shot episodic training and the style discriminator utilizes a set of style prototypes to enforce the generator to generate speech that follows the common style of each speaker. The experimental results demonstrate that StyleSpeech and Meta-StyleSpeech can synthesize high-quality speech given the reference audios from both seen and unseen speakers. Moreover, Meta-StyleSpeech achieves significantly improved adaptation performance on the speech from unseen speakers, even with a single reference speech with the length of less than one second. For future work, we plan to further improve Meta-StyleSpeech to perform controllable speech generation by disentangling its latent space, to enhance its practicality in diverse real-world applications.

Acknowledgements This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2019-0-00075, Artificial Intelligence Graduate School Program (KAIST)), and the Engineering Research Center Program through the National Research Foundation of Korea (NRF) funded by the Korean Government

(MSIT (NRF-2018R1A5A1059921)). We sincerely thank the anonymous reviewers for their constructive comments which helped us significantly improve our paper during the rebuttal period.

References

- Amodei, D., Ananthanarayanan, S., Anubhai, R., Bai, J., Battenberg, E., Case, C., Casper, J., Catanzaro, B., Chen, J., Chrzanowski, M., Coates, A., Diamos, G., Elsen, E., Engel, J. H., Fan, L., Fougner, C., Hannun, A. Y., Jun, B., Han, T., LeGresley, P., Li, X., Lin, L., Narang, S., Ng, A. Y., Ozair, S., Prenger, R., Qian, S., Raiman, J., Satheesh, S., Seetapun, D., Sengupta, S., Wang, C., Wang, Y., Wang, Z., Xiao, B., Xie, Y., Yogatama, D., Zhan, J., and Zhu, Z. Deep speech 2 : End-to-end speech recognition in english and mandarin. In Balcan, M. and Weinberger, K. Q. (eds.), *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pp. 173–182. JMLR.org, 2016. URL <http://proceedings.mlr.press/v48/amodei16.html>.
- Arik, S. Ö., Chrzanowski, M., Coates, A., Diamos, G. F., Gibiansky, A., Kang, Y., Li, X., Miller, J., Ng, A. Y., Raiman, J., Sengupta, S., and Shoeybi, M. Deep voice: Real-time neural text-to-speech. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 195–204. PMLR, 2017. URL <http://proceedings.mlr.press/v70/arik17a.html>.
- Arik, S. Ö., Chen, J., Peng, K., Ping, W., and Zhou, Y. Neural voice cloning with a few samples. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 10040–10050, 2018.
- Ba, J., Kiros, J., and Hinton, G. E. Layer normalization. *ArXiv*, abs/1607.06450, 2016.
- Bartunov, S. and Vetrov, D. P. Few-shot generative modelling with generative matching networks. In Storkey, A. J. and Pérez-Cruz, F. (eds.), *International Conference on Artificial Intelligence and Statistics, AISTATS 2018, 9-11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain*, volume 84 of *Proceedings of Machine Learning Research*, pp. 670–678. PMLR, 2018. URL <http://proceedings.mlr.press/v84/bartunov18a.html>.

- Chen, M., Tan, X., Ren, Y., Xu, J., Sun, H., Zhao, S., and Qin, T. Multispeech: Multi-speaker text to speech with transformer. In *INTERSPEECH*, 2020.
- Chen, M., Tan, X., Li, B., Liu, Y., Qin, T., sheng zhao, and Liu, T.-Y. Adaspeech: Adaptive text to speech for custom voice. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=Drynvt7gg4L>.
- Chen, Y., Assael, Y. M., Shillingford, B., Budden, D., Reed, S. E., Zen, H., Wang, Q., Cobo, L. C., Trask, A., Laurie, B., Gülçehre, Ç., van den Oord, A., Vinyals, O., and de Freitas, N. Sample efficient adaptive text-to-speech. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=rkzjUoAcFX>.
- Clouâtre, L. and Demers, M. Figr: Few-shot image generation with reptile. *ArXiv*, abs/1901.02199, 2019.
- Dauphin, Y. N., Fan, A., Auli, M., and Grangier, D. Language modeling with gated convolutional networks. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 933–941. PMLR, 2017. URL <http://proceedings.mlr.press/v70/dauphin17a.html>.
- Gibiansky, A., Arik, S. Ö., Diamos, G. F., Miller, J., Peng, K., Ping, W., Raiman, J., and Zhou, Y. Deep voice 2: Multi-speaker neural text-to-speech. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 2962–2970, 2017.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. C., and Bengio, Y. Generative adversarial nets. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pp. 2672–2680, 2014.
- Hsu, W., Zhang, Y., Weiss, R. J., Zen, H., Wu, Y., Wang, Y., Cao, Y., Jia, Y., Chen, Z., Shen, J., Nguyen, P., and Pang, R. Hierarchical generative modeling for controllable speech synthesis. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=rygkk305YQ>.
- Jia, Y., Zhang, Y., Weiss, R. J., Wang, Q., Shen, J., Ren, F., Chen, Z., Nguyen, P., Pang, R., Lopez-Moreno, I., and Wu, Y. Transfer learning from speaker verification to multispeaker text-to-speech synthesis. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 4485–4495, 2018.
- Karras, T., Laine, S., and Aila, T. A style-based generator architecture for generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pp. 4401–4410. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.00453.
- Kumar, K., Kumar, R., de Boissiere, T., Gestin, L., Teoh, W. Z., Sotelo, J., de Brébisson, A., Bengio, Y., and Courville, A. C. Melgan: Generative adversarial networks for conditional waveform synthesis. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E. B., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 14881–14892, 2019.
- La’ncucki, A. Fastpitch: Parallel text-to-speech with pitch prediction. *ArXiv*, abs/2006.06873, 2020.
- Maaten, L. V. D. and Hinton, G. E. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, 2008.
- Mao, X., Li, Q., Xie, H., Lau, R. Y. K., Wang, Z., and Smolley, S. P. Least squares generative adversarial networks. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pp. 2813–2821. IEEE Computer Society, 2017. doi: 10.1109/ICCV.2017.304. URL <https://doi.org/10.1109/ICCV.2017.304>.
- McFee, B., Raffel, C., Liang, D., Ellis, D., McVicar, M., Battenberg, E., and Nieto, O. librosa: Audio and music signal analysis in python. 2015.
- Misra, D. Mish: A self regularized non-monotonic activation function. In *31st British Machine Vision Conference 2020, BMVC 2020, Virtual Event, UK, September 7-10, 2020*. BMVA Press, 2020. URL <https://www.bmvc2020-conference.com/assets/papers/0928.pdf>.

- Miyato, T. and Koyama, M. Gans with projection discriminator. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=ByS1VpgrZ>.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=BlQRgziT->.
- Moss, H. B., Aggarwal, V., Prateek, N., González, J., and Barra-Chicote, R. BOFFIN TTS: few-shot speaker adaptation by bayesian optimization. In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, pp. 7639–7643. IEEE, 2020. doi: 10.1109/ICASSP40776.2020.9054301. URL <https://doi.org/10.1109/ICASSP40776.2020.9054301>.
- Nachmani, E., Polyak, A., Taigman, Y., and Wolf, L. Fitting new speakers based on a short untranscribed sample. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3680–3688. PMLR, 2018. URL <http://proceedings.mlr.press/v80/nachmani18a.html>.
- Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., and Kavukcuoglu, K. Wavenet: A generative model for raw audio. *ArXiv*, abs/1609.03499, 2016.
- Oreshkin, B. N., López, P. R., and Lacoste, A. TADAM: task dependent adaptive metric for improved few-shot learning. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 719–729, 2018.
- Panayotov, V., Chen, G., Povey, D., and Khudanpur, S. Librispeech: An ASR corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2015, South Brisbane, Queensland, Australia, April 19-24, 2015*, pp. 5206–5210. IEEE, 2015. doi: 10.1109/ICASSP.2015.7178964.
- Peng, K., Ping, W., Song, Z., and Zhao, K. Non-autoregressive neural text-to-speech. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7586–7598. PMLR, 2020.
- Ping, W., Peng, K., Gibiansky, A., Arik, S. Ö., Kannan, A., Narang, S., Raiman, J., and Miller, J. Deep voice 3: 2000-speaker neural text-to-speech. *ArXiv*, abs/1710.07654, 2017.
- Reed, S. E., Chen, Y., Paine, T., van den Oord, A., Eslami, S. M. A., Rezende, D. J., Vinyals, O., and de Freitas, N. Few-shot autoregressive density estimation: Towards learning to learn distributions. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- Ren, Y., Ruan, Y., Tan, X., Qin, T., Zhao, S., Zhao, Z., and Liu, T. Fastspeech: Fast, robust and controllable text to speech. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E. B., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 3165–3174, 2019.
- Ren, Y., Hu, C., Tan, X., Qin, T., Zhao, S., Zhao, Z., and Liu, T. Fastspeech 2: Fast and high-quality end-to-end text to speech. *ArXiv*, abs/2006.04558, 2020.
- Rezende, D. J., Mohamed, S., Danihelka, I., Gregor, K., and Wierstra, D. One-shot generalization in deep generative models. In Balcan, M. and Weinberger, K. Q. (eds.), *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pp. 1521–1529. JMLR.org, 2016.
- Shen, J., Pang, R., Weiss, R. J., Schuster, M., Jaitly, N., Yang, Z., Chen, Z., Zhang, Y., Wang, Y., Ryan, R., Saurous, R. A., Agiomyrgiannakis, Y., and Wu, Y. Natural TTS synthesis by conditioning wavenet on MEL spectrogram predictions. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2018, Calgary, AB, Canada, April 15-20, 2018*, pp. 4779–4783. IEEE, 2018. doi: 10.1109/ICASSP.2018.8461368.
- Skerry-Ryan, R. J., Battenberg, E., Xiao, Y., Wang, Y., Stanton, D., Shor, J., Weiss, R. J., Clark, R., and Saurous, R. A. Towards end-to-end prosody transfer for expressive speech synthesis with tacotron. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference*

- on *Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4700–4709. PMLR, 2018.
- Snell, J., Swersky, K., and Zemel, R. S. Prototypical networks for few-shot learning. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 4077–4087, 2017.
- Sotelo, J., Mehri, S., Kumar, K., Santos, J. F., Kastner, K., Courville, A. C., and Bengio, Y. Char2wav: End-to-end speech synthesis. In *ICLR*, 2017.
- Taigman, Y., Wolf, L., Polyak, A., and Nachmani, E. Voiceloop: Voice fitting and synthesis via a phonological loop. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=SkFAWax0->.
- Thrun, S. and Pratt, L. Y. (eds.). *Learning to Learn*. Springer, 1998.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 5998–6008, 2017.
- Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., and Wierstra, D. Matching networks for one shot learning. In Lee, D. D., Sugiyama, M., von Luxburg, U., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 3630–3638, 2016.
- Wan, L., Wang, Q., Papir, A., and Lopez-Moreno, I. Generalized end-to-end loss for speaker verification. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2018, Calgary, AB, Canada, April 15-20, 2018*, pp. 4879–4883. IEEE, 2018. doi: 10.1109/ICASSP.2018.8462665.
- Wang, Y., Skerry-Ryan, R., Stanton, D., Wu, Y., Weiss, R. J., Jaitly, N., Yang, Z., Xiao, Y., Chen, Z., Bengio, S., Le, Q. V., Ajiomyriannakis, Y., Clark, R., and Saurous, R. A. Tacotron: Towards end-to-end speech synthesis. In *INTERSPEECH*, 2017.
- Wang, Y., Stanton, D., Zhang, Y., Skerry-Ryan, R. J., Battenberg, E., Shor, J., Xiao, Y., Jia, Y., Ren, F., and Saurous, R. A. Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 5167–5176. PMLR, 2018.
- Yamagishi, J., Veaux, C., and Macdonald, K. Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit (version 0.92). 2019.
- Zen, H., Dang, V., Clark, R., Zhang, Y., Weiss, R. J., Jia, Y., Chen, Z., and Wu, Y. Libritts: A corpus derived from librispeech for text-to-speech. *ArXiv*, abs/1904.02882, 2019.
- Zhang, Z., Tian, Q., Lu, H., Chen, L., and Liu, S. Adadurian: Few-shot adaptation for neural text-to-speech with durian. *ArXiv*, abs/2005.05642, 2020.