



Human-in-the-loop Speaker Adaptation for DNN-based Multi-speaker TTS

Kenta Udagawa, Yuki Saito, Hiroshi Saruwatari

Graduate School of Information Science and Technology, The University of Tokyo, Japan.

udagawa-kenta338@g.ecc.u-tokyo.ac.jp, yuuki.saito@ipc.i.u-tokyo.ac.jp

Abstract

This paper proposes a human-in-the-loop speaker-adaptation method for multi-speaker text-to-speech. With a conventional speaker-adaptation method, a target speaker's embedding vector is extracted from his/her reference speech using a speaker encoder trained on a speaker-discriminative task. However, this method cannot obtain an embedding vector for the target speaker when the reference speech is unavailable. Our method is based on a human-in-the-loop optimization framework, which incorporates a user to explore the speaker-embedding space to find the target speaker's embedding. The proposed method uses a sequential line search algorithm that repeatedly asks a user to select a point on a line segment in the embedding space. To efficiently choose the best speech sample from multiple stimuli, we also developed a system in which a user can switch between multiple speakers' voices for each phoneme. Experimental results indicate that the proposed method can achieve comparable performance to the conventional one in objective and subjective evaluations even if reference speech is not used as the input of a speaker encoder directly.

Index Terms: DNN-based multi-speaker TTS, speaker adaptation, human-computer interaction, Bayesian optimization

1. Introduction

Text-to-speech (TTS) based on deep neural networks (DNNs) [1, 2] has attracted attention as a technique for synthesizing a high-quality speech from text. Multi-speaker TTS, such as Deep Voice 2 [3], aims to synthesize various speaker's voices using a single TTS model. A common method for multi-speaker TTS is to use a lookup table of speaker embeddings, i.e., a trainable matrix representing each speaker's real-valued ID. Because most parameters of the TTS model with this method are shared among different speakers, one can build a multi-speaker TTS system using a small amount of data per speaker. However, with such multi-speaker TTS methods using a speaker-embedding table, it is difficult to synthesize the voice of an unseen speaker who is not included in the training data. Speaker adaptation for multi-speaker TTS [4, 5, 6] is a technique for synthesizing the speech of such unseen speakers.

Speaker adaptation for multi-speaker TTS can be interpreted as the task of searching the speaker-embedding space for a vector that can reproduce the target speaker's voice. With a transfer-learning-based speaker-adaptation method [7], a few seconds of the target speaker's reference speech is prepared, and the target speaker's embedding vector is estimated by feeding the reference speech to the speaker encoder pre-trained on the speaker-verification task. This method can synthesize high-quality speech for the target speaker from only a small amount of reference speech while reducing the cost of collecting high-quality speech and transcription pair data for TTS model training. However, this method cannot be applied when reference speech from the target speaker is unavailable.

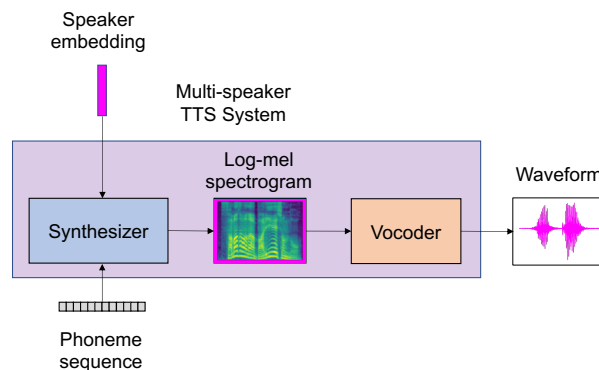


Figure 1: Overview of multi-speaker TTS system

We propose a speaker-adaptation method that incorporates human speech perception to explore the target speaker's embedding. The proposed method is based on a sequential line search (SLS) algorithm [8], in which a user repeatedly selects a point on a presented line segment in the search parameter space by manipulating a single slider. In the speaker-adaptation setting, the user selects the voice closest to the target speaker from voices synthesized from multiple speaker embeddings on a line segment of the speaker-embedding space. The proposed method synthesizes multiple voices in advance from multiple speaker embeddings to improve the exploration efficiency, and the user selects one voice using a system with which the user can switch the played-back voice for each phoneme while looping an utterance. The proposed method can be used without the target speaker's reference speech because it feeds back only human-preference data for searching the target speaker's embedding. Experimental results indicate that the proposed method can synthesize speech with the same or better quality than the conventional method, depending on the user and test speaker.

2. Conventional speaker adaptation

2.1. Multi-speaker text-to-speech synthesis

We first give an overview of the multi-speaker TTS system used in this paper. As shown in Figure 1, this system consists of two DNNs: a synthesizer that generates a mel spectrogram from a phoneme sequence and speaker-embedding vector and a vocoder that synthesizes a speech waveform from the generated mel spectrogram. A synthesizer is trained on pairs of text and audio, and a vocoder is trained separately on pairs of mel spectrograms and speech waveforms. The structure of this TTS system is common to both conventional and proposed methods.

2.2. Speaker adaptation based on transfer learning from speaker verification

With the speaker-adaptation method based on transfer learning from speaker verification [7], the target speaker's embedding

is obtained using a speaker encoder that extracts the speaker-embedding vector from the speech waveform. The speaker encoder is trained to optimize the generalized end-to-end (GE2E) loss for speaker verification [9] separately from the synthesizer and vocoder networks. The training using GE2E loss makes speaker embeddings extracted from the same speaker’s speech waveforms closer together and moves different speakers’ embeddings farther apart. The synthesizer is trained to predict a mel spectrogram from given text with the speaker embedding extracted with the trained speaker encoder. During inference, the target speaker’s embedding is extracted from a few seconds of the target speaker’s reference speech, then the extracted speaker embedding and text are fed into the trained multi-speaker TTS system to synthesize a speech waveform with the voice of the target speaker.

3. Human-in-the-loop speaker adaptation

3.1. Sequential line search

We first describe the SLS algorithm [8], which is the search algorithm used with the proposed method. Let D be the dimensionality of the search parameters, $\mathcal{X} = [0, 1]^D$ is search parameter space, and $g : \mathcal{X} \rightarrow \mathbb{R}$ is user’s perceptual preference function, respectively. The SLS algorithm repeats the process of having the user select a parameter point from a line segment in \mathcal{X} to obtain \mathbf{x}^* , which is defined as:

$$\mathbf{x}^* = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} g(\mathbf{x}). \quad (1)$$

Let S_t be the line segment proposed for the t th step in the SLS algorithm, \mathbf{x}_{t-1}^+ , and $\mathbf{x}_{t-1}^{\text{EI}}$ be the endpoints of S_t , respectively. In the original SLS setting, the initial endpoints \mathbf{x}_0^+ , \mathbf{x}_0^{EI} are two random points in \mathcal{X} . The endpoints \mathbf{x}_t^+ , \mathbf{x}_t^{EI} of the line segment S_{t+1} are determined from the following equations:

$$\mathbf{x}_t^+ = \mathbf{x}_t^{\text{chosen}}, \quad (2)$$

$$\mathbf{x}_t^{\text{EI}} = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} a^{\text{EI}}(\mathbf{x}; D_t). \quad (3)$$

Here, $\mathbf{x}_t^{\text{chosen}}$ represents the point selected by the user from S_t , and this definition of \mathbf{x}_t^+ was proposed in [10], D_t represents the user’s selection up to the t th step, and $a^{\text{EI}}(\cdot)$ is an acquisition function based on expected improvement (EI) [11] and defined as:

$$a^{\text{EI}}(\mathbf{x}; D_t) = \mathbb{E}[\max\{g(\mathbf{x}) - g^+, 0\}], \quad (4)$$

where g^+ is the best preference function among the observation points. To calculate $a^{\text{EI}}(\mathbf{x}; D_t)$, a Gaussian process prior is assumed on $g(\cdot)$ to make the preference-function value $g(\mathbf{x}_*)$ at the unobserved point \mathbf{x}_* follow the normal distribution:

$$g(\mathbf{x}_*) \sim \mathcal{N}(\mu_t(\mathbf{x}_*), \sigma_t(\mathbf{x}_*)). \quad (5)$$

The mean $\mu_t(\mathbf{x}_*)$ and variance $\sigma_t(\mathbf{x}_*)$ of the distribution depend on the preference-function value $\mathbf{g} = (g(\mathbf{x}_1), \dots, g(\mathbf{x}_{N_t}))^\top$ at observation point $\{\mathbf{x}_i\}_{i=1}^{N_t}$ and hyperparameters θ . The preference-function value at the observation point \mathbf{g} and hyperparameters θ are estimated by maximum a posteriori estimation. This allows the probability distribution of the preference-function value $g(\mathbf{x}_*)$ at the unobserved point \mathbf{x}_* to be expressed in closed form and the acquisition function $a^{\text{EI}}(\mathbf{x}; D_t)$ to be calculated. Therefore, the endpoints \mathbf{x}_t^+ , \mathbf{x}_t^{EI} of line segment S_{t+1} can be determined from Equations (2) and (3).

3.2. SLS-based speaker adaptation

Our speaker-adaptation method uses the SLS algorithm to search for the target speaker’s embedding. An overview of the

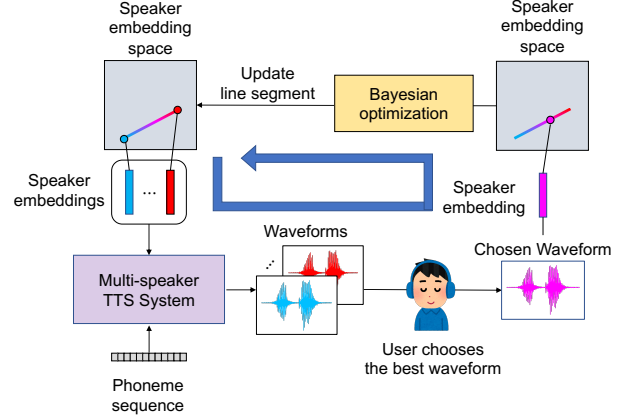


Figure 2: Overview of proposed speaker-adaptation method

proposed method is shown in Figure 2. The user selects the waveform closest to the target speaker’s voice from multiple synthesized speech waveforms. When the user perceives non-time-series data such as images, a single piece of data can be perceived instantly. However, for time-series data such as voice, it takes several seconds to perceive an entire piece of data. For the user to efficiently select the best waveform from multiple waveforms, we developed a system in which the user can switch the played-back voice for each phoneme while looping an utterance. With this system, the user selects the preferred speech waveform from a continuously changing voice by manipulating a single slider.

Since the proposed human-in-the-loop speaker-adaptation method does not use the reference voice as input for the speaker encoder, it has the advantage of synthesizing the voice of the speaker imagined by the user. This means that the method can be used even when reference voices are difficult to obtain, such as for deceased or inarticulate speakers. The proposed method also has possible applications such as cross-lingual speaker adaptation [12] and audiobook speech synthesis [13].

3.3. Initial line segment and search-parameter space

To search for high-quality synthetic speech¹, the endpoints of the initial line segment with the proposed method are set to the mean male-speaker embedding and mean female-speaker embedding. In addition, the search-parameter space \mathcal{X} is set to the quantile Q of the training data. Let q_i be the i th component of the SLS search parameter $\mathbf{q} \in Q = [0, 1]^{16}$, e_i be i th component of the speaker embedding $\mathbf{e} \in [0, 1]^{16}$, and $e_{i,j}^{\text{train}}$ be j th smallest i -component of the N speaker embeddings in the training data for the synthesizer ($e_{i,0}^{\text{train}} = -\infty, e_{i,N+1}^{\text{train}} = \infty$), respectively. The transformation Conv from e_i to q_i and its inverse function IConv q_i to e_i are defined as follows (where $\lfloor \cdot \rfloor$ is a floor function):

$$\text{Conv}(e_i) = j/N, \quad (e_{i,j}^{\text{train}} \leq e_i < e_{i,j+1}^{\text{train}}), \quad (6)$$

$$\text{IConv}(q_i) = e_{i, \lfloor q_i(N-1) \rfloor + 1}^{\text{train}}. \quad (7)$$

¹In our preliminary experiments, we observed that the quality of the synthesized speech is lower when the speaker embedding is far from the distribution of the training data.

4. Experiments

4.1. Experimental conditions

For training the speaker encoder, we used the Japanese speech data included in the Corpus of Spontaneous Japanese (CSJ) [14]. The CSJ contains 660 hours of speech data from 1417 Japanese speakers (947 men and 470 women). The CSJ audio data were resampled to 16 kHz, and the frame shift was set to 10 ms. For training, validating, and testing the synthesizer, we used the Japanese speech data included in the parallel data of the Japanese Versatile Speech (JVS) corpus [15]. The parallel data of the JVS corpus contain 22 hours of speech data from 100 Japanese speakers (49 men and 51 women; 100 sentences per speaker). The validation and test data were 100 utterances of 10 speakers (5 men and 5 women) randomly sampled from the parallel data of the JVS corpus, and the training data were 20 hours of speech data from the remaining 90 speakers. As corner cases in speaker-adaptation difficulty, we selected four speakers for the test data on the basis of the speaker’s average subjective similarity [16] to speakers in the training data: the highest male (“jvs078”) and female (“jvs060”) and the lowest male (“jvs005”) and female (“jvs010”). We excluded data containing annotation errors and recording failure². The JVS speech data were resampled to 22.05 kHz to match the settings of neural vocoding, and the frame shift was set to 12 ms.

For the synthesizer, we used the open source implementation of FastSpeech 2 [17] published by ming024³. FastSpeech 2 predicted the 80-dimensional mel spectrogram from Japanese phonemes with the aid of the variance adaptors that predicted the F0 and energy of synthetic speech. The F0 was estimated using the WORLD vocoder [18, 19]. In the FastSpeech 2 training, we used the Adam optimizer [20] with 4,000 Warmup [21] steps, an initial learning rate of 0.0625, batch size of 8, and 50,000 training steps.

We used the same architecture of a speaker encoder as with the conventional method [7]. However, we reduced the speaker embedding’s dimensionality from 256 to 16 for the proposed method to make the Bayesian-optimization-based SLS algorithm work well [22]. We also changed the speaker-embedding layer’s activation function from a rectified linear unit (ReLU) [23] with L2-normalization to sigmoid. This is because with the SLS algorithm a hypercube of $[0, 1]^D$ as the D -dimensional search parameter space is assumed⁴. We denote **EMB256dim** and **EMB16dim** as the speaker embeddings used for the conventional and proposed methods, respectively. Both types of speaker embedding were projected to the 256-dimensional vector space through the two fully connected layers: 256 hidden units with the ReLU activation function and 256 output units with the linear activation function, then added to the output of text encoder in the synthesizer. In the speaker-encoder training, we used the Adam optimizer with a learning rate of 0.0001, batch size of 8, and 1,000,000 training step. The vocoder for the speech-waveform synthesis was the generator_universal model of HiFi-GAN [24] included in the FastSpeech 2 repository published by ming024.

We used the SLS setting in which the line segment is extended by 1.25 times before presenting it to the user. The number of points that can be selected from the SLS line segment

²The preprocessing of the JVS corpus was on the basis of the repository at https://github.com/Hiroshiba/jvs_hiho

³<https://github.com/ming024/FastSpeech2>

⁴In our preliminary experiments, we observed that these modifications did not affect the quality of synthetic speech with the conventional speaker adaptation method.

was 20. The hyperparameters θ were set to default, i.e., same as the repository at <https://github.com/yuki-koyama/sequential-line-search>.

The mean absolute error (MAE) between the natural and synthetic mel spectrograms was used as an objective evaluation metric regarding the quality of the speech by speaker adaptation. The mel spectrogram MAE was calculated for speech synthesized by giving FastSpeech 2 the phoneme duration of the natural speech. With both the proposed and conventional methods, noise was observed in the pause interval of the speech synthesized by FastSpeech 2. Therefore, the paused part of the mel spectrogram was masked when synthesizing speech and omitted when calculating the mel spectrogram MAE.

4.2. Human-in-the-loop adaptation experiment

We conducted an experiment in which participants used the proposed method to search for the target speakers’ voices. The participants were instructed to select the synthesized voice they thought was closest to the target speaker’s voice at each step. The number of SLS steps was set to 30 so that the experimental time per speaker adaptation would be 15–25 minutes. The participants were eight members of a laboratory specializing in speech and acoustic processing. Although the proposed method can execute speaker adaptation to the target speaker in the user’s mind without the reference voice, in this experiment, to make the participants fully aware of the target speaker’s voice, we allowed them to listen to the reference voice when using the proposed method. However, we prohibited the participants from playing the synthesized voice and reference voice simultaneously and comparing them to simulate a situation in which the target speaker’s voice is in the user’s mind. We used the VOICEACTRESS100.001 utterance in the parallel data of the JVS corpus for the speaker-embedding search with the proposed method. The duration of the speech VOICEACTRESS100.001 in the test speaker was about eight seconds. The following two methods were compared with the proposed method.

1. **TL**: The mean of “EMB256dim” speaker embeddings extracted from all utterances per test speaker. This represents the conventional method [7].
2. **Mean-Speaker**: The mean of “EMB16dim” speaker embedding extracted from all the JVS training data of the same sex as the target speaker. This represents the initial value of the SLS algorithm with the proposed method. The difference between the proposed method and this means a change in a quality of synthesized voice due to human manipulation.

To examine the quality of the voice selected by a participant at each step, we calculated the mel spectrogram MAE of the speech for embedding exploration per step. Figure 3 shows the mel spectrogram MAE of the speech for embedding exploration per step when the proposed method was used up to 30 steps by the eight participants. The light blue line in Figure 3 represents the results from one participant, and the dark blue line represents the average results from all eight participants. The red and green lines represent the mel spectrogram MAE of “TL” and “Mean-Speaker,” respectively. The light blue line with the lowest mel spectrogram MAE shows that the proposed method can synthesize speech comparable to the conventional method in terms of objective evaluation. However, if we focus on the dark blue line (the average results from the eight participants), we can see that the mel spectrogram MAE does not tend to improve by human manipulation for all test speakers. The possible reasons for this are as follows. (1) The manipulation

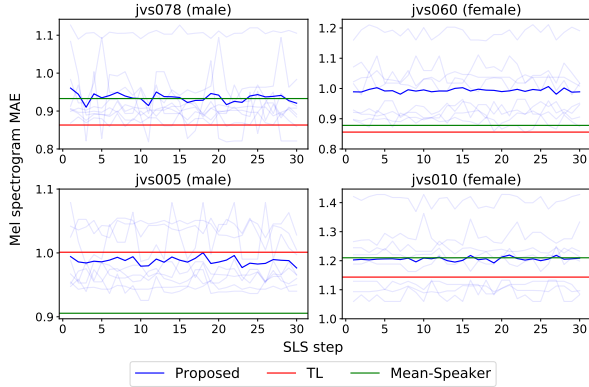


Figure 3: Mel spectrogram MAE per step when eight participants used the proposed method.

Table 1: Mel spectrogram MAE of 100 test utterances. **Bold values are best among 5 methods per test speaker.**

	SLS			TL	Mean-Speaker
	Worst	Mean	Best		
jvs078	1.03	0.85	0.80	0.82	0.89
jvs005	1.06	0.99	0.94	0.96	0.95
jvs060	1.10	0.95	0.90	0.84	0.90
jvs010	1.27	1.14	1.05	1.05	1.14

task was difficult because the search parameters (i.e., speaker embedding) were too abstract. (2) The mel spectrogram MAE with “Mean-Speaker” (i.e., the initial value with the proposed method) was sufficiently low.

4.3. Objective evaluation

As the objective evaluation of the conventional and proposed methods, we computed the mel spectrogram MAE between the natural and synthetic speech by using each test speaker’s 100 utterances. In addition to “TL” and “Mean-Speaker”, we defined **SLS- $\{\text{Best, Mean, Worst}\}$** as the speaker embedding whose the mel spectrogram MAE with the utterance for searching was {minimum, mean, maximum} among speaker embeddings obtained after the eight participants manipulated the proposed method up to 30 steps. Table 1 shows the results⁵. The quality of “SLS-Best” for speakers other than “jvs060” and “SLS-Mean” for speakers of “jvs078” and “jvs005” was comparable to “TL”. However, for all test speakers, “SLS-Worst” was significantly inferior to “TL” and “Mean-Speaker”. The reason for this could be that the manipulation task was difficult, as discussed in Section 4.2.

4.4. Subjective evaluation

The naturalness and speaker similarity of 100 utterances of the four test speakers synthesized with the conventional and proposed methods were evaluated on the basis of the five-level mean opinion score (MOS) and degradation MOS (DMOS) tests, respectively. The synthetic voices to be compared were the same as with the five methods compared in Section 4.3. In the subjective evaluation of naturalness, natural speech ground-truth (GT) was evaluated in addition to speech synthesized with the five methods, and in the subjective evaluation of speaker similarity, natural speech was used as the reference speech. Subjective evaluation was conducted separately for each of the eight tasks (four test speakers \times {MOS or DMOS}). For each task, 50 listeners were recruited by crowdsourcing, and they evaluated the quality of speech samples. The number of speech

⁵The audio samples are available on [this page](#).

Table 2: Subjective evaluation results. *Underlined values are comparable with or better than “TL” with statistical significance at 5-percent level.*

(a) MOS of speech naturalness

	SLS			TL	Mean-Speaker	GT
	Worst	Mean	Best			
jvs078	3.29	3.46	3.48	3.67	<u>3.56</u>	4.17
jvs005	2.81	<u>3.63</u>	<u>3.61</u>	3.56	<u>3.52</u>	4.12
jvs060	3.19	3.13	<u>3.42</u>	3.35	<u>3.34</u>	4.25
jvs010	<u>3.48</u>	<u>3.62</u>	3.13	3.42	<u>3.58</u>	3.75

(b) DMOS of speaker similarity

	SLS			TL	Mean-Speaker
	Worst	Mean	Best		
jvs078	1.90	<u>2.89</u>	<u>3.04</u>	2.71	1.95
jvs005	1.47	2.12	<u>2.54</u>	2.33	<u>2.67</u>
jvs060	2.25	2.19	2.68	3.20	2.98
jvs010	1.54	2.03	<u>3.08</u>	2.98	1.96

samples in the evaluation was five utterances (randomly sampled from 100 utterances) \times the number of comparison methods (six for the MOS test and five for the DMOS test). One of the five utterances was used as a dummy voice for the listeners to define the evaluation criteria.

The results of the subjective evaluation of naturalness are shown in Table 2(a). First, “Mean-Speaker” was about as natural as “TL”. This indicates that the proposed method using the SLS algorithm with “Mean-Speaker” as the initial value can achieve the same level of naturalness as the conventional method. Although “SLS-Worst” of “jvs005” was significantly inferior to the conventional method, the naturalness of the proposed method was comparable to the conventional method in several cases.

The results of the subjective evaluation of speaker similarity are shown in Table 2(b). With the proposed method, “SLS-Best” outperformed the conventional method for all speakers except “jvs060”, and “SLS-Mean” was comparable to the conventional method for “jvs078”. The results suggest that the proposed method can achieve the same quality of speaker adaptation as the conventional method even if reference speech is not used as the input of a speaker encoder. However, there were several cases in which the DMOS of the proposed method was below “Mean-Speaker”. This means that the speaker similarity was lowered by the human search than the speech contained in the initial line. The reason for this could be that the manipulation task was difficult, as discussed in Section 4.2.

5. Conclusions

We proposed a human-in-the-loop speaker-adaptation method that does not input reference speech to the speaker encoder and uses only human preference to find the target speaker’s embedding. The experimental results indicate that the proposed method can synthesize speech with the same quality as the conventional method in both objective and subjective evaluations. To make the user’s manipulation of the parameters more intuitive for non-expert users, we will investigate a method for providing explainability to the speaker embedding.

6. Acknowledgements

This work was supported by JST, Moonshot R&D Grant Number JPMJPS2011.

7. References

- [1] H. Zen, A. Senior, and M. Schuster, "Statistical parametric speech synthesis using deep neural networks," in *Proc. ICASSP*, Vancouver, Canada, May 2013, pp. 7962–7966.
- [2] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan, R. A. Saurous, Y. Agiomyriannakis, and Y. Wu, "Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions," in *Proc. ICASSP*, Calgary, Canada, Apr. 2018, pp. 4779–4783.
- [3] S. O. Arik, G. Damos, A. Gibiansky, J. Miller, K. Peng, W. Ping, J. Raiman, and Y. Zhou, "Deep Voice 2: Multi-speaker neural text-to-speech," in *Proc. NIPS*, Red Hook, New York, USA, May 2017, pp. 2966–2974.
- [4] Z. Wu, P. Swietojanski, C. Veaux, S. Renals, and S. King, "A study of speaker adaptation for DNN-based speech synthesis," in *Proc. INTERSPEECH*, Dresden, Germany, Sep. 2015, pp. 879–883.
- [5] Y. Taigman, L. Wolf, A. Polyak, and E. Nachmani, "Voiceloop: Voice fitting and synthesis via a phonological loop," in *Proc. ICLR*, Vancouver, BC, Canada, May 2018.
- [6] S. O. Arik, J. Chen, K. Peng, W. Ping, and Y. Zhou, "Neural voice cloning with a few samples," in *Proc. NeurIPS*, Montreal, Canada, Dec. 2018, pp. 10 019–10 029.
- [7] Y. Jia, Y. Zhang, R. Weiss, Q. Wang, J. Shen, F. Ren, Z. Chen, P. Nguyen, R. Pang, I. L. Moreno, and Y. Wu, "Transfer learning from speaker verification to multispeaker text-to-speech synthesis," in *Proc. NeurIPS*, Montreal, Canada, Dec. 2018, pp. 4480–4490.
- [8] Y. Koyama, I. Sato, D. Sakamoto, and T. Igarashi, "Sequential line search for efficient visual design optimization by crowds," *ACM Transactions on Graphics*, vol. 36, no. 4, pp. 1–11, Jul. 2017.
- [9] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, "Generalized end-to-end loss for speaker verification," in *Proc. ICASSP*, Alberta, Canada, Apr. 2018, pp. 4879–4883.
- [10] Y. Koyama, I. Sato, and M. Goto, "Sequential gallery for interactive visual design optimization," *ACM Transactions on Graphics*, vol. 39, no. 4, pp. 88:1–88:12, Jul. 2020.
- [11] D. R. Jones, M. Schonlau, and W. J. Welch, "Efficient global optimization of expensive black-box functions," *Journal of Global optimization*, vol. 13, no. 4, pp. 455–492, Dec. 1998.
- [12] D. Xin, Y. Saito, S. Takamichi, T. Koriyama, and H. Saruwatari, "Cross-Lingual Text-To-Speech Synthesis via Domain Adaptation and Perceptual Similarity Regression in Speaker Space," in *Proc. INTERSPEECH*, Shanghai, China, 2020, pp. 2947–2951.
- [13] W. Nakata, T. Koriyama, S. Takamichi, N. Tanji, Y. Ijima, R. Masumura, and H. Saruwatari, "Audiobook speech synthesis conditioned by cross-sentence context-aware word embeddings," in *Proc. SSW*, Budapest, Hungary, Aug. 2021, pp. 211–215.
- [14] K. Maekawa, "Corpus of spontaneous Japanese: Its design and evaluation," in *Proc. SSPR*, Tokyo, Japan, Apr. 2003, pp. 7–12.
- [15] S. Takamichi, R. Sonobe, K. Mitsui, Y. Saito, T. Koriyama, N. Tanji, and H. Saruwatari, "JSUT and JVS: Free Japanese voice corpora for accelerating speech synthesis research," *Acoustical Science and Technology*, vol. 41, no. 5, pp. 761–768, Sep. 2020.
- [16] Y. Saito, S. Takamichi, and H. Saruwatari, "Perceptual-similarity-aware deep speaker representation learning for multi-speaker generative modeling," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1033–1048, Feb. 2021.
- [17] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, "Fastspeech 2: Fast and high-quality end-to-end text to speech," in *Proc. ICLR*, Virtual Conference, Sep. 2021.
- [18] M. Morise, F. Yokomori, and K. Ozawa, "WORLD: a vocoder-based high-quality speech synthesis system for real-time applications," *IEICE Transactions on Information and Systems*, vol. E99-D, no. 7, pp. 1877–1884, 2016.
- [19] M. Morise, "D4C, a band-a-periodicity estimator for high-quality speech synthesis," *Speech Communication*, vol. 84, pp. 57–65, Nov. 2016.
- [20] D. Kingma and B. Jimmy, "Adam: A method for stochastic optimization," in *arXiv preprint arXiv:1412.6980*, 2014.
- [21] P. Goyal, P. Dollár, R. B. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He, "Accurate, large minibatch SGD: Training imagenet in 1 hour," *arXiv*, vol. abs/1706.02677, 2017.
- [22] R. Moriconi, M. P. Deisenroth, and K. S. Kumar, "High-dimensional bayesian optimization using low-dimensional feature spaces," *Machine Learning*, vol. 109, no. 9, pp. 1925–1943, Sep. 2020.
- [23] X. Glorot, A. Bordes, and Y. Bengio, "Deep sparse rectifier neural networks," in *Proc. AISTATS*, Lauderdale, U.S.A., Apr. 2011, pp. 315–323.
- [24] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," in *Proc. NeurIPS*, Virtual Conference, Dec. 2020, pp. 17 022–17 033.