



# Creating Personalized Synthetic Voices from Post-Glossectomy Speech with Guided Diffusion Models

*Yusheng Tian, Guangyan Zhang, Tan Lee*

Department of Electronic Engineering, The Chinese University of Hong Kong, Hong Kong SAR

{ystian0617, gyzhang}@link.cuhk.edu.hk, tanlee@ee.cuhk.edu.hk

## Abstract

This paper is about developing personalized speech synthesis systems with recordings of mildly impaired speech. In particular, we consider consonant and vowel alterations resulted from partial glossectomy, the surgical removal of part of the tongue. The aim is to restore articulation in the synthesized speech and maximally preserve the target speaker's individuality. We propose to tackle the problem with guided diffusion models. Specifically, a diffusion-based speech synthesis model is trained on original recordings, to capture and preserve the target speaker's original articulation style. When using the model for inference, a separately trained phone classifier will guide the synthesis process towards proper articulation. Objective and subjective evaluation results show that the proposed method substantially improves articulation in the synthesized speech over original recordings, and preserves more of the target speaker's individuality than a voice conversion baseline.

**Index Terms:** personalized speech synthesis, post-glossectomy speech, articulation disorder, guided diffusion models

## 1. Introduction

People diagnosed with oral cancer may need to have the entire voice box removed by surgery, losing the ability to speak thereafter. One way to save their voices is to collect audio recordings of their speech before the scheduled operation date and create personalized text-to-speech (TTS) models from those speech data. The TTS systems would allow these individuals to communicate with other people using their own voices. Ideally the speech used to create TTS models should carry accurate and clear pronunciation such that highly intelligible speech can be generated. In reality, some of the patients already suffer speech impairment at the time of recording. For example, tongue cancer, one of the most common sites of oral cancer, is often treated with surgery. People who underwent tongue surgery typically show consonant and vowel alterations [1, 2]. We encountered one such case recently. A young female Cantonese speaker approached us, expressing the hope of saving her voice through speech synthesis technology. She received partial glossectomy six years ago, and about 3/4 of her tongue was removed surgically. This has resulted in significant difficulties in articulating accurately due to the defect in the tongue. Consonant and vowel alterations are reflected noticeably in this young lady's speech recordings. This application context leads to the task committed in the present study: developing a personalized TTS system from post-glossectomy speech. Our goal is to restore articulation in the synthesized speech and maximally preserve the target speaker's individuality.

The problem of voice reconstruction from impaired speech was tackled in a few previous studies. They share the same idea

of substituting impaired speech segments by normal ones. The Voicebank project [3–5] applied HMM-based speech synthesis techniques to create personalized synthetic voices for people with speech disorder. Synthetic voices were repaired by substituting selected acoustic feature parameters with those of an average healthy voice. In [6, 7], neural speech synthesis systems were developed for individuals with dysarthria. A two-step process was developed: first train a neural TTS model on recordings from a healthy speaker, then perform voice conversion on the synthesized speech from the healthy speaker to the impaired speaker.

The above substitution-based approaches are considered suboptimal for the task we are considering, where the major symptoms of speech disorder is alteration of certain phones. On the one hand, the articulation style is part of an individual's personality. Preserving only the voice timbre as in [6, 7] would completely discard the target speaker's articulation style, including both bad and good aspects, leading to undesirable loss of individuality. On the other hand, although HMM-based systems can allow acoustic feature substitution on selected phones [4], the average voice model may not provide a replacement close enough to the target speaker's voice.

We propose to restore articulation in the synthesized speech with guided diffusion models. Specifically, a diffusion-based TTS model [8] is trained on the target speaker's original recordings (with impaired articulation on certain speech sounds), and hence capture and preserve the original articulation style. When using the trained model for inference, the synthesis process is guided by the gradient from a separately trained phone classifier. The output of the phone classifier indicates how precise the articulation is on a continuous scale, which facilitates finer adjustment of articulation than simple substitution.

The proposed system design is inspired by GuidedTTS [9], which employs a phone classifier to guide an unconditional diffusion model for speech generation. However, the motivation of using classifier guidance in GuidedTTS is to develop TTS systems with untranscribed speech data. The phone classifier guidance in the present study is in the spirit more similar to shallow-fusion biasing in end-to-end automatic speech recognition (ASR) [10, 11], in which an external language model steers the recognition process towards a particular context at each decoding step. In the proposed system, the external phone classifier guides the synthesis process towards correct pronunciation at each reverse diffusion step.

In the next section, we outline key concepts in diffusion models that are related to diffusion-based speech synthesis. Section 3 describes the proposed system of guided speech synthesis. Section 4 and 5 present the experimental results on the aforementioned real patient case. Section 6 concludes and discusses the limitations of the proposed approach.

## 2. Background

Diffusion models are a family of probabilistic generative models. The modelling involves a forward process that progressively contaminates the data with random noise, and a reverse process that generates data from random noise. The forward and backward process were originally formulated as Markov chains [12–14]. In [15], the discrete-time Markov process is generalized to a continuous-time stochastic differential equation (SDE). Specifically, the forward process is defined by the following equation:

$$d\mathbf{x}_t = \mathbf{f}(\mathbf{x}_t, t)dt + g(t)d\mathbf{w}_t, \quad (1)$$

where  $t \sim \mathcal{U}(0, T)$ ;  $\mathbf{w}_t$  is a standard Brown motion;  $\mathbf{f}(\cdot)$  and  $g(\cdot)$  are the so-called drift and diffusion coefficient. These two coefficients are constructed such that as  $t$  grows from 0 to  $T$ , the probability distribution of  $\mathbf{x}_t$  would evolve from the original data distribution to a tractable prior, typically a Gaussian distribution with fixed mean and variance.

The corresponding reverse process also forms a SDE as stated in [16]:

$$d\mathbf{x}_t = [\mathbf{f}(\mathbf{x}_t, t) - g^2(t)\nabla_{\mathbf{x}_t} \log P(\mathbf{x}_t)] dt + g(t)d\bar{\mathbf{w}}_t \quad (2)$$

where  $\bar{\mathbf{w}}_t$  is a standard Brown motion running backward in time. The core part of a diffusion model is to train a neural network  $S_\theta$  to estimate the value of  $\nabla_{\mathbf{x}_t} \log P(\mathbf{x}_t)$ , a.k.a., the score. Once the score is known for all  $t$ , we can sample data by solving the reverse SDE using numerical solvers.

It was proved in [15] that the reverse SDE is associated with an ordinary differential equation (ODE) as follows:

$$d\mathbf{x}_t = \left[ \mathbf{f}(\mathbf{x}_t, t) - \frac{1}{2}g^2(t)\nabla_{\mathbf{x}_t} \log P(\mathbf{x}_t) \right] dt, \quad (3)$$

which shares the same marginal distribution for all  $t$ . It is empirically shown in [8] that inference with the ODE formulation requires fewer sampling steps. In the remainder of this paper, we use the ODE formulation to model the reverse process.

## 3. Approach

### 3.1. Diffusion-based TTS

We follow GradTTS [8] and use diffusion models to generate Mel-Spectrograms conditioned on the input text and speaker labels. The forward process SDE is defined as:

$$d\mathbf{x}_t = \frac{1}{2}(\boldsymbol{\mu} - \mathbf{x}_t)\beta_t dt + \sqrt{\beta_t}d\mathbf{w}_t, \quad (4)$$

where  $t \sim \mathcal{U}(0, 1)$ ,  $\beta_t = \beta_0 + (\beta_1 - \beta_0)t$  is a predefined linear noise scale, and  $\boldsymbol{\mu}$  is an average Mel-spectrogram corresponding to the input phone sequence. One important result derived from (4) is the conditional distribution of  $\mathbf{x}_t$  given  $\mathbf{x}_0$ :

$$P(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\boldsymbol{\rho}(\mathbf{x}_0, t), \sigma_t^2 \mathbf{I}), \quad (5)$$

where  $\boldsymbol{\rho}(\mathbf{x}_0, t) = (1 - e^{-\frac{1}{2}\int_0^t \beta_s ds})\boldsymbol{\mu} + e^{-\frac{1}{2}\int_0^t \beta_s ds}\mathbf{x}_0$ , and  $\sigma_t^2 = 1 - e^{-\int_0^t \beta_s ds}$ . If  $\mathbf{x}_0$  is known, we can then draw samples of  $\mathbf{x}_t$  using the reparameterization trick:

$$\mathbf{x}_t = \boldsymbol{\rho}(\mathbf{x}_0, t) + \sigma_t \boldsymbol{\epsilon}_t, \boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (6)$$

The reverse time ODE is given by:

$$d\mathbf{x}_t = \frac{1}{2}\beta_t [\boldsymbol{\mu} - \mathbf{x}_t - \nabla_{\mathbf{x}_t} \log P(\mathbf{x}_t|\boldsymbol{\mu}, s)] dt, \quad (7)$$

where  $s$  stands for the speaker label. Note that unlike the unconditional reverse process given by (3), the reverse process in the context of TTS is conditioned on the input text and the speaker label.

The neural network  $S_\theta$  is trained to predict the conditional score function  $\nabla_{\mathbf{x}_t} \log P(\mathbf{x}_t|\boldsymbol{\mu}, s)$  using a weighted L2 loss:

$$\mathcal{L}(\theta) = \mathbb{E}_t \sigma_t^2 \mathbb{E}_{\mathbf{x}_0} \mathbb{E}_{\boldsymbol{\epsilon}_t} \|S_\theta(\mathbf{x}_t, t, \boldsymbol{\mu}, s) + \sigma_t^{-1} \boldsymbol{\epsilon}_t\|_2^2, \quad (8)$$

where we have made use of the following results:

$$P(\mathbf{x}_t|\mathbf{x}_0, \boldsymbol{\mu}, s) = P(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\boldsymbol{\rho}(\mathbf{x}_0, t), \sigma_t^2 \mathbf{I}), \quad (9)$$

$$\nabla_{\mathbf{x}_t} \log P(\mathbf{x}_t|\mathbf{x}_0) = -\sigma_t^{-1} \boldsymbol{\epsilon}_t. \quad (10)$$

### 3.2. Guided synthesis process

Once the diffusion-based acoustic model is trained, the target speaker’s articulation style will be embedded into the generated Mel-Spectrogram through the speaker condition  $s$ . The synthetic voice will inevitably preserve the problematic articulation pattern. In order to improve articulation in the synthesized speech, another condition  $\mathbf{y}_t$  is provided to the score estimator.  $\mathbf{y}_t$  stands for the phone label sequence of  $\mathbf{x}_t$  predicted by an external frame-level phone classifier, i.e., what the external phone classifier thinks the synthetic voice is saying. Now the reverse time ODE is augmented as

$$d\mathbf{x}_t = \frac{1}{2}\beta_t(\boldsymbol{\mu} - \mathbf{x}_t)dt - \frac{1}{2}\beta_t [\nabla_{\mathbf{x}_t} \log P(\mathbf{x}_t|\boldsymbol{\mu}, s) + \nabla_{\mathbf{x}_t} \log P(\mathbf{y}_t|\mathbf{x}_t)] dt, \quad (11)$$

where we have assumed that the phone classifier is speaker-independent, i.e.  $\log P(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\mu}, s) = \log P(\mathbf{y}_t|\mathbf{x}_t)$ .

To see why incorporating this external phone classifier can improve the articulation in the synthetic voice, consider the following scenarios. First, suppose the articulation is not good, i.e. a low value of  $P(\mathbf{y}_t|\mathbf{x}_t)$ . Thus the additional term  $\nabla_{\mathbf{x}_t} \log P(\mathbf{y}_t|\mathbf{x}_t)$  will bias the synthesis process towards a direction that increases  $\log P(\mathbf{y}_t|\mathbf{x}_t)$ . If the phone classifier is trained on speech from healthy speakers, then the articulation would be improved as  $\log P(\mathbf{y}_t|\mathbf{x}_t)$  grows. Second, suppose the articulation is good, i.e., a high log probability of  $P(\mathbf{y}_t|\mathbf{x}_t)$ . In such case the additional term is close to zero and will have little impact on the synthesized speech. Consequently, the good part of the target speaker’s articulation style will be preserved.

In practice, we usually scale  $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{y}_t|\mathbf{x}_t)$  with a positive coefficient  $\gamma$  in order to obtain better sample quality [9, 17]. We adopt the norm-based scale [9], which is computed as  $\gamma = \alpha \cdot \|\nabla_{\mathbf{x}_t} \log P(\mathbf{x}_t|\boldsymbol{\mu}, s)\|_2 / \|\nabla_{\mathbf{x}_t} \log P(\mathbf{y}_t|\mathbf{x}_t)\|_2$ , where  $\alpha$  is a hyperparameter to be tuned.

Assume that the prediction of phone label for each frame is independent of each other, then we have

$$\nabla_{\mathbf{x}_t} \log P(\mathbf{y}_t|\mathbf{x}_t) = \sum_{i=1}^L \nabla_{\mathbf{x}_t} \log P(\mathbf{y}_t^{(i)}|\mathbf{x}_t), \quad (12)$$

where  $L$  refers to the number of acoustic frames. This property enables flexible guidance weights on different frames. For example, we can put extra weights on frames corresponding to phones that the target speaker tends to make mistakes.




Figure 1: Schematic diagram of the proposed guided speech synthesis process.

### 3.3. System design

Figure 1 gives an overview of the proposed speech synthesis system. The input phone sequence  $p$  is fed into a text encoder. The resulted phone embeddings, as well as the speaker label  $s$ , are sent to a duration predictor to estimate the frame length for each phone. The input phone sequence  $p$  is then expanded to  $\bar{p}$  according to the predicted duration  $\bar{d}$ , similar to the length regulator mechanism in [18–20].

In GradTTS  $\mu$  is designed to be dependent on both the input phone sequence and the speaker label, but we augment it to be conditioned only on the phone sequence, with the hope that a shared phone embedding across speakers will transfer some articulation knowledge from healthy speakers to the target speaker, and therefore alleviate the articulation disorder in the synthetic voice. Specifically,  $\mu$  is obtained by looking up a predefined phone-to-Mel-spectrum dictionary computed over training data, similar to that in [21].

The score estimator  $S_\theta$  is trained on speech data from a multi-speaker speech corpus, plus recordings from the target speaker. In contrast, the phone classifier is trained only on recordings from healthy speakers, to ensure that it is sensitive to abnormal articulation style.

## 4. Experimental Setup

### 4.1. Baseline systems for comparison

We compare the proposed system with two TTS baselines. The first baseline, *DuriTaco*, is adapted from Tacotron [22], with the attention module replaced by a duration-informed aligner as in DurIAN [19]. The second baseline, *DuriTaco+VC*, has exactly the same architecture as *DuriTaco*, but is trained on voice-converted speech: unimpaired speech from a healthy source speaker converted into the target speaker’s voice via a voice conversion (VC) model. We use the recently proposed NANSY [23] for voice conversion, as it shows strong performance in cross-lingual setting. The underlying assumption is that impaired speech can be viewed as a new language, therefore a strong cross-lingual VC model is expected to perform well on voice conversion between normal and impaired speech.

### 4.2. Datasets

We use the following four datasets for different purposes.

- **CUSENT** [24] for multi-speaker TTS pre-training, as well as for training the external phone classifier and many-to-many VC model. It is a Cantonese speech corpus of around 20 hours clean read speech from 80 speakers, sampled at 16 kHz.
- **Recording.T** for personalized TTS fine-tuning. It contains 377 utterances, giving a total of approximately 24 minutes speech from our target speaker, a young Cantonese-speaking female. The recording script is adapted from CUSENT and all Cantonese initials, finals and tones are included. The audio was recorded with a TASCAM DR-44WL at 44.1 kHz in her living place, under reasonably quiet condition.

- **Recording.S** as the source speaker’s data for voice conversion. It contains 377 utterances that share exactly the same content as Recordings.T, recorded by another female Cantonese speaker with no articulation disorder. The recording was carried out in a sound-proof studio.
- **KingASR086** for training a CTC-based ASR model for objective evaluation. It is a commercial Cantonese speech corpus purchased from SpeechOcean<sup>1</sup>, which contains 80-hour read speech from 136 speakers, sampled at 44.1 kHz.

### 4.3. Implementation details

We adopted the official implementation of GradTTS<sup>2</sup> for the proposed model, but augmented the text encoder to be conditioned only on the input phone sequence as mentioned earlier. For simplification, we also used forced alignment [25], instead of monotonic alignment search to obtain the duration labels. We used the Jasper architecture [26] for the frame-level phone classifier, but removed the convolution stride so that it could predict phone labels for each frame. The *DuriTaco* baseline was similar to a public implementation of a baseline DurIAN<sup>3</sup>, except that the duration predictor and the acoustic model were separately trained. The VC model in the *DuriTaco+VC* baseline followed a public implementation of NANSY<sup>4</sup>. The CTC-based ASR model for objective evaluation was implemented following the recipe from SpeechBrain<sup>5</sup>.

All TTS models generate 80-dimensional log Mel-Spectrograms to drive a pre-trained HiFi-GAN vocoder [27]. TTS audio data were resampled to 22.05 kHz for Mel-Spectrogram computation, to be consistent with the settings of HiFi-GAN. The one-hot speaker embedding dimension is 16 in all TTS models. The coefficient  $\alpha$  for classifier guidance in the proposed model is set to 0.3. We also assigned extra guidance weight of 5.0 to several selected phones that the target speaker had difficulties to articulate. All TTS models were pre-trained on CUSENT for 700 epochs at batch size 32, then fine-tuned on the target speaker’s data (original or voice-converted) for 5000 steps at batch size 16. A total of 30 sentences from target speaker’s data were held out for evaluation. The number of reverse steps for the proposed diffusion-based TTS model is 25 using the DPM solver [28].

## 5. Results

### 5.1. Objective evaluation

We resynthesized all 377 sentences of the target speaker’s recordings with the proposed system and the two baselines,

<sup>1</sup><https://en.speechocean.com/datacenter/recognition.html>

<sup>2</sup><https://github.com/huawei-noah/Speech-Backbones/tree/main/Grad-TTS>

<sup>3</sup><https://github.com/ivanvovk/durian-pytorch>

<sup>4</sup><https://github.com/dhchoi99/NANSY>

<sup>5</sup><https://github.com/speechbrain/speechbrain/tree/develop/recipes/TIMIT/ASR/CTC>

Table 1: *PER(%) on resynthesized speech and the original recordings with the separately trained ASR model.*

Case	Audio Source	PER%
Recording_S	Real	11.2
Recording_T	Real	43.1
DuriTaco	TTS	37.6
DuriTaco + VC	TTS	<b>14.9</b>
Diffusion	TTS	36.4
GuidedDiff	TTS	<b>22.1</b>




Figure 2: *MUSHRA scores for overall impression. Mean value and 95% confidence interval are reported at the bottom.*

and used the separately trained CTC-based ASR to evaluate the quality of articulation. The Phone Error Rate (PER%) results in Table 1 show that the proposed system GuidedDiff produces significantly better articulation than the voice in the original recordings and the DuriTaco baseline, though not as good as the VC baseline. By comparing the last two rows we can conclude that the improvement comes from the classifier guidance, rather than the diffusion model.

## 5.2. Subjective evaluation

The subjective evaluation was conducted through a web-based listening test. It consists of three parts. The first two parts was in MUSHRA-style format (no anchor but reference was given) [29]. Specifically, the same 30 sentences from the evaluation set were synthesized by the proposed and baseline voices. For each sentence, the three synthetic voices plus the original recordings from the target and source speaker formed five stimuli in total. In Part1, recordings from the source speaker were provided as the reference, and listeners were asked to rate the *overall impression* of each stimulus on a scale from 0 to 100. We avoided the term *naturalness* as it might be confusing whether disordered speech from a natural person count as natural. Part2 was similar to Part1, except that this time recordings from the target speaker were provided as the reference to evaluate speaker similarity. Part3 used A/B test format to examine whether perceptually the proposed GuidedDiff improves articulation over the DuriTaco baseline. We skipped the comparison between the proposed GuidedDiff and the VC baseline because results from objective test clearly show that the VC baseline produces articulation nearly as good as natural speech. For the A/B test 15 sentences were selected from a script written by the target speaker. These sentences contain at least one word that the target speaker found difficult to articulate. The listening test is no-paid. In order to constrain the test duration to be within 15 minutes, we chose to randomly expose 10 out of 30 questions to the listeners in the first two parts.

We received 87 effective responses after filtering out listeners who failed to spot the hidden reference. For the MUSHRA-




Figure 3: *MUSHRA scores for speaker similarity. Mean value and 95% confidence interval are reported at the bottom.*




Figure 4: *Preference test demonstrating that GuidedDiff produces better articulation than the DuriTaco baseline.*

style tests we performed a two-sided Wilcoxon signed-rank test on all pairs of stimuli, and corrected with Holm-Bonferroni. Results in Figure 2 reveal similar pattern as the objective test: GuidedDiff improves the overall quality over the DuriTaco baseline and the target speaker’s original recordings ( $p < 0.001$ ), though not as good as the VC baseline. One interesting finding is that even the DuriTaco baseline produces perceptually better speech than those in original recordings, indicating that a shared phone embedding across speakers does help to alleviate the articulation disorder in the synthetic voice. Results in Figure 3 demonstrate the advantage of GuidedDiff over the VC baseline in terms of preserving the target speaker’s individuality, with the voice in GuidedDiff and DuriTaco not significantly different from that in original recordings ( $p > 0.09$ ), while the same conclusion does not apply to DuriTaco+VC ( $p < 0.001$ ). The preference test result given in Figure 4 further demonstrates that GuidedDiff synthesizes better articulation than the DuriTaco baseline. Readers are encouraged to visit <https://diffcorrect.github.io/> to listen to audio samples.

## 6. Conclusion and Discussion

We introduced the use of classifier-guided diffusion models for the creation of personalized synthetic voices from post-glossectomy speech. Experimental results on a real patient case show that the synthetic voice can restore articulation for phones that the target speaker had difficulties to articulate, and at the same time maximally preserve the target speaker’s individuality. The proposed model takes around 1.4 seconds to synthesize 30 Cantonese characters on a single 2080Ti GPU. In the future we will extend the current model to take speech as additional condition, aiming at correcting improper articulations in the recordings directly. In this way the modified recordings with correct articulation can be used to train other TTS models to enable faster inference speed.

## 7. Acknowledgements

We thank all the participants in the listening test, for their kindness and contributions. Our deepest gratitude belongs to the young lady, who generously agreed to reveal her synthetic voice on the demo page. The first author is supported by the Hong Kong Ph.D. Fellowship Scheme of the Hong Kong Research Grants Council.

## 8. References

- [1] A. Acher, P. Perrier, C. Savariaux, and C. Fougeron, "Speech production after glossectomy: methodological aspects," *Clinical linguistics & phonetics*, vol. 28, no. 4, pp. 241–256, 2014.
- [2] M. Girod-Roux, T. Hueber, D. Fabre, S. Gerber, M. Canault, N. Bedoin, A. Acher, N. Béziaud, E. Truy, and P. Badin, "Rehabilitation of speech disorders following glossectomy, based on ultrasound visual illustration and feedback," *Clinical Linguistics & Phonetics*, vol. 34, no. 9, pp. 826–843, 2020.
- [3] J. Yamagishi, C. Veaux, S. King, and S. Renals, "Speech synthesis technologies for individuals with vocal disabilities: Voice banking and reconstruction," *Acoustical Science and Technology*, vol. 33, no. 1, pp. 1–5, 2012.
- [4] C. Veaux, J. Yamagishi, and S. King, "Using hmm-based speech synthesis to reconstruct the voice of individuals with degenerative speech disorders," in *Proc. Interspeech – 13<sup>th</sup> Annual Conference of the International Speech Communication Association*, 2012, pp. 967–970.
- [5] S. Creer, S. Cunningham, P. Green, and J. Yamagishi, "Building personalised synthetic voices for individuals with severe speech impairment," *Computer Speech & Language*, vol. 27, no. 6, pp. 1178–1193, 2013.
- [6] R. Nanzaka and T. Takiguchi, "Hybrid text-to-speech for articulation disorders with a small amount of non-parallel data," in *APSIPA ASC 2018 – 10<sup>th</sup> Asia-Pacific Signal and Information Processing Association Annual Summit and Conference*. IEEE, 2018, pp. 1761–1765.
- [7] K. Matsubara, T. Okamoto, R. Takashima, T. Takiguchi, T. Toda, Y. Shiga, and H. Kawai, "High-intelligibility speech synthesis for dysarthric speakers with lpcnet-based tts and cyclevae-based vc," in *ICASSP 2021 – 46<sup>th</sup> IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2021, pp. 7058–7062.
- [8] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, and M. Kudinov, "Grad-tts: A diffusion probabilistic model for text-to-speech," in *ICML 2021 – 38<sup>th</sup> International Conference on Machine Learning*. PMLR, 2021, pp. 8599–8608.
- [9] H. Kim, S. Kim, and S. Yoon, "Guided-tts: A diffusion model for text-to-speech via classifier guidance," in *ICML 2022 – 39<sup>th</sup> International Conference on Machine Learning*. PMLR, 2022, pp. 11 119–11 133.
- [10] I. Williams, A. Kannan, P. S. Aleksic, D. Rybach, and T. N. Sainath, "Contextual speech recognition in end-to-end neural network systems using beam search," in *Proc. Interspeech 2018 – 19<sup>th</sup> Annual Conference of the International Speech Communication Association*, 2018, pp. 2227–2231.
- [11] D. Zhao, T. N. Sainath, D. Rybach, P. Rondon, D. Bhatia, B. Li, and R. Pang, "Shallow-fusion end-to-end contextual biasing," in *Proc. Interspeech 2019 – 20<sup>th</sup> Annual Conference of the International Speech Communication Association*, 2019, pp. 1418–1422.
- [12] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *ICML 2015 – 32<sup>nd</sup> International Conference on Machine Learning*. PMLR, 2015, pp. 2256–2265.
- [13] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," in *NeurIPS 2019 – 33<sup>rd</sup> Advances in neural information processing systems*, 2019, pp. 11 895–11 907.
- [14] J. Ho, A. Jain, and P. Abbeel, "Denosing diffusion probabilistic models," in *NeurIPS 2020 – 34<sup>th</sup> Advances in Neural Information Processing Systems*, 2020, pp. 6840–6851.
- [15] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," in *ICLR 2021 – 9<sup>th</sup> International Conference on Learning Representations*, 2021.
- [16] B. D. Anderson, "Reverse-time diffusion equation models," *Stochastic Processes and their Applications*, vol. 12, no. 3, pp. 313–326, 1982.
- [17] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," in *NeurIPS 2021 – 35<sup>th</sup> Advances in Neural Information Processing Systems*, 2021, pp. 8780–8794.
- [18] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, "FastSpeech: Fast, robust and controllable text to speech," in *NeurIPS 2019 – 33<sup>rd</sup> Advances in neural information processing systems*, 2019, pp. 3165–3174.
- [19] C. Yu, H. Lu, N. Hu, M. Yu, C. Weng, K. Xu, P. Liu, D. Tuo, S. Kang, G. Lei, D. Su, and D. Yu, "DurIAN: Duration Informed Attention Network for Speech Synthesis," in *Proc. Interspeech 2020 – 21<sup>st</sup> Annual Conference of the International Speech Communication Association*, 2020, pp. 2027–2031.
- [20] J. Shen, Y. Jia, M. Chrzanowski, Y. Zhang, I. Elias, H. Zen, and Y. Wu, "Non-attentive tacotron: Robust and controllable neural TTS synthesis including unsupervised duration modeling," *CoRR*, 2020.
- [21] T. Sadekova, V. Gogoryan, I. Vovk, V. Popov, M. A. Kudinov, and J. Wei, "A unified system for voice cloning and voice conversion through diffusion probabilistic modeling," in *Proc. Interspeech 2022 – 23<sup>rd</sup> Annual Conference of the International Speech Communication Association*, 2022, pp. 3003–3007.
- [22] Y. Wang, R. J. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. V. Le, Y. Agiomyriannakis, R. Clark, and R. A. Saurous, "Tacotron: Towards end-to-end speech synthesis," in *Proc. Interspeech 2017 – 18<sup>th</sup> Annual Conference of the International Speech Communication Association*, 2017, pp. 4006–4010.
- [23] H.-S. Choi, J. Lee, W. Kim, J. Lee, H. Heo, and K. Lee, "Neural analysis and synthesis: Reconstructing speech from self-supervised representations," in *NeurIPS 2021 – 35<sup>th</sup> Advances in Neural Information Processing Systems*, 2021, pp. 16 251–16 265.
- [24] T. Lee, W. K. Lo, P.-C. Ching, and H. Meng, "Spoken language resources for cantonese speech processing," *Speech Communication*, vol. 36, no. 3-4, pp. 327–342, 2002.
- [25] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using kaldii," in *Proc. Interspeech 2017 – 18<sup>th</sup> Annual Conference of the International Speech Communication Association*, 2017, pp. 498–502.
- [26] J. Li, V. Lavrukhin, B. Ginsburg, R. Leary, O. Kuchaiev, J. M. Cohen, H. Nguyen, and R. T. Gadde, "Jasper: An end-to-end convolutional neural acoustic model," in *Proc. Interspeech 2019 – 20<sup>th</sup> Annual Conference of the International Speech Communication Association*, 2019, pp. 71–75.
- [27] J. Kong, J. Kim, and J. Bae, "Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis," in *NeurIPS 2020 – 34<sup>th</sup> Advances in Neural Information Processing Systems*, 2020, pp. 17 022–17 033.
- [28] C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, and J. Zhu, "DPM-solver: A fast ODE solver for diffusion probabilistic model sampling in around 10 steps," in *NeurIPS 2022 – 36<sup>th</sup> Advances in Neural Information Processing Systems*, 2022, pp. 5775–5787.
- [29] B. Series, "Method for the subjective assessment of intermediate quality level of audio systems," *International Telecommunication Union Radiocommunication Assembly*, 2014.