

# DrawSpeech: Expressive Speech Synthesis Using Prosodic Sketches as Control Conditions

Weidong Chen

*The Chinese University of Hong Kong*  
Hong Kong SAR, China  
wdchen@se.cuhk.edu.hk

Shan Yang

*Tencent AI Lab*  
Shenzhen, China  
shaanyang@tencent.com

Guangzhi Li

*Tencent AI Lab*  
Shenzhen, China  
guangzhilei@tencent.com

Xixin Wu

*The Chinese University of Hong Kong*  
Hong Kong SAR, China  
wuxx@se.cuhk.edu.hk

**Abstract**—Controlling text-to-speech (TTS) systems to synthesize speech with the prosodic characteristics expected by users has attracted much attention. To achieve controllability, current studies focus on two main directions: (1) using reference speech as prosody prompt to guide speech synthesis, and (2) using natural language descriptions to control the generation process. However, finding reference speech that exactly contains the prosody that users want to synthesize takes a lot of effort. Description-based guidance in TTS systems can only determine the overall prosody, which has difficulty in achieving fine-grained prosody control over the synthesized speech. In this paper, we propose DrawSpeech, a sketch-conditioned diffusion model capable of generating speech based on any prosody sketches drawn by users. Specifically, the prosody sketches are fed to DrawSpeech to provide a rough indication of the expected prosody trends. DrawSpeech then recovers the detailed pitch and energy contours based on the coarse sketches and synthesizes the desired speech. Experimental results show that DrawSpeech can generate speech with a wide variety of prosody and can precisely control the fine-grained prosody in a user-friendly manner. Our implementation and audio samples are publicly available<sup>1</sup>.

**Index Terms**—Speech Synthesis, Prosody Control, Sketch

## I. INTRODUCTION

Given a piece of text, we can express it through various prosodic patterns, each conveying distinct meanings [1]. For example, in the same sentence, speakers may emphasize different words to highlight specific information, leading to entirely different interpretations [2]. Thus, controlling a text-to-speech (TTS) system to synthesize the exact prosody desired by users has attracted considerable attention from researchers [3]–[8].

Current research on prosody-controllable TTS systems has explored two primary directions. One approach involves selecting a reference speech as a prosody prompt and instructing the TTS model to replicate the prosodic pattern on the target text [8]–[15]. Typically, StyleTTS [10] extracted a style vector from the reference speech and adopted a prosody predictor to predict pitch and energy based on the extracted style vector.

NaturalSpeech 3 [11] disentangled prosody attributes from the reference speech, which then served as prosody condition for a prosody diffusion model to generate prosody representations for the input text. However, finding a reference speech that precisely matches the desired prosody pattern can be time-consuming [16]. In addition, transferring the prosody from the reference speech results in reduced prosodic variations [17].

Another mainstream approach for controlling synthesized prosody is to use natural language descriptions to specify the desired prosody [18]–[25]. For example, PromptStyle [20] aligns the description prompt embedding with the style embedding during training, enabling description-based control during inference. PromptTTS [21] leveraged a pre-trained language model BERT [26] to capture the semantic information in the descriptions, which was subsequently used to guide the speech synthesis process. However, description-based guidance can only manage utterance-level prosody with descriptions such as “his voice has a low pitch”, “generate a voice with high volume”, or “please use gentle tone”, leaving precise control at a fine-grained level still an unsolved problem. Moreover, constructing a dataset with descriptions to train such models requires significant effort and a high level of expertise [21].

In this paper, we propose a novel prosody control signal, called prosody sketch, for expressive and controllable TTS. Since pitch and energy are crucial components of prosody [27], [28], our focus is primarily on controlling these two attributes. To be specific, the sketch serves as a guide, indicating the desired prosody trends to the model, such as a rise and fall in pitch or an increase in energy within a specific region. Users can easily sketch the target prosody for each word, making fine-grained prosody control effortless. To achieve the new control paradigm, we propose DrawSpeech, which first reconstructs the pitch and energy contours from the abstract sketches, and then employs a diffusion model conditioned on the sketches, contours, and target text to synthesize speech.

The contributions of this work are summarized as follows:

- We propose using a novel conditioning signal, called the prosody sketch, which can be easily sketched and enables fine-grained prosody control.

This project is partially supported by the General Research Fund from the Research Grants Council of Hong Kong SAR Government (Project No. 14202623).

<sup>1</sup><https://happycolor.github.io/DrawSpeech>

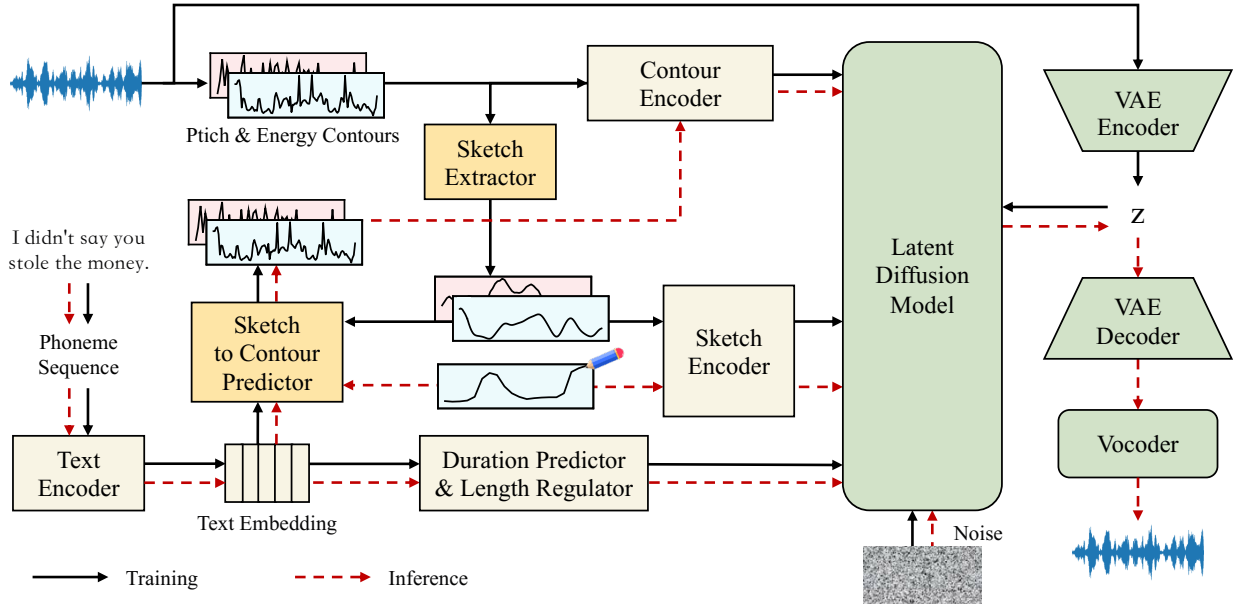


Fig. 1. Overview structure of the proposed DrawSpeech. Paired speech and text data are used for training. User-supplied text and drawn pitch or energy sketch are used as inputs during inference.

- We propose a diffusion-based generative model, named DrawSpeech, which conditions on the sketch and is capable of generating speech with varied prosody.
- Experiments show the effectiveness of the sketch-based control paradigm and the DrawSpeech system, especially on its capability to achieve precise prosody control.

## II. METHODOLOGY

The architecture of DrawSpeech is shown in Fig. 1, where the sketch extractor and sketch-to-contour predictor are used to achieve the transformation between the sketches and contours. Contour encoder and sketch encoder convert the corresponding inputs into different embedding features. A text encoder equipped with a duration predictor and length regulator handles the input transcript. A latent diffusion model with a variational autoencoder are employed to generate log-mel spectrogram and a vocoder is used to reconstruct speech waveform. Details will be provided in the following subsections.

### A. Sketch Extractor

DrawSpeech aims to allow users to easily control the prosody of the synthesized speech. Taking pitch as an example, as shown in Fig. 2(a), the pitch contour is vibrating and changes quickly between the adjacent values. Therefore, controlling prosody based on the contour requires users to provide the detailed fluctuations of pitch or energy, which is inconvenient and sometimes even impractical. A more feasible solution is to simply ask the user to provide the sketch that reflects the pitch or energy trend in the sentence. For example, the pitch sketch shown in Fig. 2(b) tells the model that the sentence should start out at a higher pitch then drop quickly, followed by a gradual rise and arrive at the peaks on certain positions, etc. It's much easier for users to draw a pitch sketch




Fig. 2. Illustrations of (a) pitch contour and (b) pitch sketch. (c) Put the contour and the sketch in the same frame for direct comparison.

as in Fig. 2(b) as a control signal because it relieves users' burden to provide the detailed vibration of the pitch contour as in Fig. 2(a), which will be handled by the model instead.

To obtain the sketch for model training, we first extract the phoneme-level pitch contour  $P \in \mathbb{R}^M$  and energy contour  $E \in \mathbb{R}^M$  of the input speech, where  $M$  denotes the number of phonemes in the sentence.  $P$  and  $E$  are then smoothed using Savitzky-Golay filter [29], which is useful for smoothing data while preserving the shape of the signal. The resulting pitch and energy sketches are denoted as  $P_{ske} \in \mathbb{R}^M$  and  $E_{ske} \in \mathbb{R}^M$ , respectively. As shown in Fig. 2(c), the smoothed sketch effectively captures the variation trends of the contour.

### B. Sketch-to-Contour Predictor

While pitch and energy sketches are intuitive for users to comprehend, they omit too much detailed information, making it challenging for speech synthesis systems to reconstruct natural speech. The sketch-to-contour predictor is designed to leverage text and sketches to restore the detailed contours.

Given the input transcript  $y$ , we first convert it into the phoneme sequence (with an open-sourced grapheme-to-

phoneme toolkit<sup>2</sup>). Then, a text encoder  $f_{text}(\cdot)$  is employed to extract the phoneme embedding  $H^y \in \mathbb{R}^{M \times D}$ , where  $D$  represents the phoneme dimension. The sketch-to-contour predictor can then be written as:

$$H^y = f_{text}(y) \quad (1)$$

$$P_{pred}, E_{pred} = \mathcal{P}(H^y, P_{ske}, E_{ske}) \quad (2)$$

where  $\mathcal{P}$  represents the sketch-to-contour predictor;  $P_{pred} \in \mathbb{R}^M$  and  $E_{pred} \in \mathbb{R}^M$  are the predicted pitch and energy contours, respectively.

### C. Latent Diffusion Model

Firstly, a variational autoencoder (VAE) model [30] is employed to compress the mel-spectrogram  $x \in \mathbb{R}^{T \times F}$  into a small latent space  $z_0 \in \mathbb{R}^{C \times \frac{T}{r} \times \frac{F}{r}}$ , where  $T$  is the number of frames and  $F$  is the frequency dimension,  $C$  is the channel dimension and  $r$  is the compression rate. Simply put, the calculation flow in VAE is as follows:

$$z_0 = Enc(x) \quad (3)$$

$$x_{rec} = Dec(z_0) \quad (4)$$

where  $Enc$  and  $Dec$  are the encoder and decoder part of VAE,  $x_{rec}$  denotes the reconstructed mel-spectrogram.

Secondly, a contour encoder and a sketch encoder are employed to generate conditions for the latent diffusion model (LDM). Specifically, the contour encoder quantizes pitch and energy ( $P$  and  $E$  during training, and  $P_{pred}$  and  $E_{pred}$  during inference) into 256 discrete values, subsequently converting them into pitch embedding  $H^p \in \mathbb{R}^{M \times F}$  and energy embedding  $H^e \in \mathbb{R}^{M \times F}$ . The same process takes place in the sketch encoder, except that it utilizes  $P_{ske}$  and  $E_{ske}$  as inputs and produces pitch sketch embedding  $\hat{H}_{ske}^p \in \mathbb{R}^{M \times F}$  and energy sketch embedding  $\hat{H}_{ske}^e \in \mathbb{R}^{M \times F}$ . To solve the problems of word skipping and repetition, we introduce a duration predictor and a length regulator following [31]. The ground-truth phoneme durations are used to upsample embeddings during training, while the predicted durations are used during inference. The expanded embeddings are denoted as  $\hat{H}^y, \hat{H}^p, \hat{H}^e, \hat{H}_{ske}^p, \hat{H}_{ske}^e \in \mathbb{R}^{T \times F}$ . Note that a 1D convolutional layer is applied to the phoneme embedding  $H^y$  to project the phoneme dimension from  $D$  to  $F$  prior to the expansion.

Given all the conditions, we employ a LDM [32] to generate the latent representation  $z_0$  from Gaussian noise  $z_N \sim \mathcal{N}(0, I)$ , where  $N$  denotes the total sampling steps. For each time step  $n \in [1, N]$ , the reverse process in LDM is:

$$\bar{H}^y = \hat{H}^y + \hat{H}^p + \hat{H}^e \quad (5)$$

$$C_n = Concat(z_n, \hat{H}_{ske}^p, \hat{H}_{ske}^e, \bar{H}^y) \quad (6)$$

$$z_{n-1} = z_n - \epsilon_\theta(C_n, t) \quad (7)$$

where  $Concat(\cdot)$  denotes the concatenation along the channel dimension,  $t$  denotes the timestep,  $\epsilon_\theta$  denotes the diffusion model with parameters  $\theta$ . For further details on the forward and reverse processes, as well as the training techniques of the diffusion model, please refer to [32]–[34].

<sup>2</sup><https://github.com/Kyubyong/g2p>

## III. EXPERIMENTS

### A. Experimental Setup

1) *Dataset*: LJSpeech dataset [35] is a widely used speech corpus containing 13,100 short audio clips from a single speaker, each paired with a transcription. The total recording duration is roughly 24 hours. We randomly divided the dataset into three subsets: 12,500 samples for training, 300 for validation, and 300 for testing, with a sample rate of 22,050 Hz.

2) *Implementation Details*: The text encoder and sketch-to-contour predictor are composed of multiple stacked Transformer [36] blocks (6 blocks for the former and 2 blocks for the latter). In both, the linear layer in the feed-forward network is replaced by a 1D convolutional layer. The duration predictor and length regulator are identical to those described in [31]. The VAE and LDM adopt the structures and configurations used in [37]. The contours are normalized using the mean and variance of the entire dataset. The sketches are normalized to [0, 1]. The dimension of the phoneme embedding  $D$  is 256. The window size and hop size for extracting mel-spectrogram are 1024 and 256, respectively. The number of mel frequency bins  $F$  is 80. The pre-trained HiFi-GAN<sup>3</sup> vocoder [38] is utilized to reconstruct speech waveform.

We first train the VAE on LJSpeech for 80k steps with batchsize of 8. Subsequently, the VAE and vocoder remain frozen, while the other modules are trained alongside the LDM for 80k steps with batchsize of 32. We use the Adam [39] optimizer to update the VAE and the AdamW [40] optimizer to update the LDM along with other modules, and follow the same learning rate schedule in [37]. We set either  $P_{ske}$  or  $E_{ske}$  to an all-zero vector with a probability of 20% during model training, enabling DrawSpeech to require only a single sketch input from the user during the inference phase. The user-drawn sketch is used as the pitch sketch by default.

3) *Baselines*: We adopt two TTS models, FastSpeech 2 [41] and NaturalSpeech 2 [12], as our baseline models. This selection is based on the fact that both models explicitly incorporate pitch or energy information in speech, allowing us to easily insert the sketch conditions into the original models for comparison. However, since FastSpeech 2 and NaturalSpeech 2 do not provide official implementations, we instead use the widely adopted public implementation of FastSpeech 2<sup>4</sup> and utilize the NaturalSpeech 2 checkpoint provided by Amphion<sup>5</sup>.

4) *Evaluation Metrics*: We use mean opinion score (MOS) to evaluate the sound naturalness, with ratings ranging from 1 to 5 (1-bad, 2-poor, 3-fair, 4-good, 5-excellent). To evaluate prosody controllability, we introduce sketch correlation (SC), which measures the alignment between the speech and the given sketch. SC is rated on a scale from 1 to 5, with 5 signifying that the sketch accurately reflects the prosody trend of the speech, and 1 indicating no correlation between the given sketch and the synthesized speech. We invite 22 listeners for subjective tests. Note that listeners are instructed

<sup>3</sup><https://github.com/jik876/hifi-gan>

<sup>4</sup><https://github.com/ming024/FastSpeech2>

<sup>5</sup><https://github.com/open-mmlab/Amphion/tree/main/egs/tts/NaturalSpeech2>

TABLE I

MOS AND SC COMPARISON WITH 95% CONFIDENCE INTERVALS WHEN ADOPTING SKETCHES (SKE) OR CONTOURS (CON) AS CONTROL SIGNALS

| Method                         | Control          | MOS $\uparrow$                    | SC $\uparrow$                     |
|--------------------------------|------------------|-----------------------------------|-----------------------------------|
| FastSpeech 2                   | Ske              | $2.40 \pm 0.10$                   | $3.62 \pm 0.09$                   |
|                                | Con <sup>†</sup> | $3.37 \pm 0.08$                   | $2.82 \pm 0.09$                   |
| FastSpeech 2 (FT) <sup>‡</sup> | Ske              | $3.55 \pm 0.14$                   | $3.98 \pm 0.13$                   |
|                                | Con <sup>†</sup> | $2.51 \pm 0.09$                   | $2.38 \pm 0.10$                   |
| NaturalSpeech 2                | Ske              | $2.51 \pm 0.09$                   | $2.38 \pm 0.10$                   |
|                                | Con <sup>†</sup> | $2.51 \pm 0.09$                   | $2.35 \pm 0.10$                   |
| DrawSpeech                     | Ske              | <b><math>4.49 \pm 0.06</math></b> | <b><math>4.30 \pm 0.07</math></b> |

<sup>†</sup> The contours are adjusted based on the sketches.

<sup>‡</sup> We fine-tune FastSpeech 2 with sketch conditions. The Sketch Extractor and Sketch-to-Contour Predictor are integrated into FastSpeech 2 during fine-tuning for a fair comparison.

TABLE II

MOS WITH 95% CONFIDENCE INTERVALS IN THE ABLATION STUDIES

| Method                            | MOS $\uparrow$                    | p-value   |
|-----------------------------------|-----------------------------------|-----------|
| GroundTruth                       | $4.46 \pm 0.10$                   | -         |
| DrawSpeech with text only         | $3.91 \pm 0.09$                   | $1.3e-24$ |
| DrawSpeech with sketch conditions | <b><math>4.49 \pm 0.06</math></b> | 0.058     |

to disregard the speech quality when assessing SC. For objective metrics, root mean square error (RMSE) is employed to assess the accuracy of the prosody control. We select 10 text samples from the test set, with each sample synthesized using two different sketches as prosody condition, resulting in 20 utterances from each system being used for evaluation.

### B. Experimental Results

In this section, we aim to answer three questions: Can other models without special designs utilize sketches as conditions? Will using sketches reduce sound quality? Can using sketches as conditions enable fine-grained and precise prosody control?

1) *Applying sketches to other models:* To incorporate sketches as conditions into FastSpeech 2, we directly replace the predicted pitch and energy contours in FastSpeech 2 with the pitch and energy sketches. We also try to multiply the sketches with the predicted contours so that the resulting contours retain the overall trends of the sketches while preserving the details of the original contours. In this case, the control signals are the contours adjusted based on the sketches. For NaturalSpeech 2, since it only predicts the pitch contour, we incorporate only the pitch sketch as condition. The results are shown in Table I. From the MOS scores, we can observe that DrawSpeech outperforms its competitors by a substantial margin. We speculate that the abstract sketches provide insufficient information for FastSpeech 2 and NaturalSpeech 2 to recover natural speech, and the adjusted contours are out-of-domain data for the two baselines. From the SC scores, we find that DrawSpeech excels at synthesizing speech that closely matches the desired prosody variations, confirming that DrawSpeech can effectively utilize the sketch conditions. Additionally, when FastSpeech 2 is fine-tuned with specialized modules, the SC score shows a significant improvement compared to the case without fine-tuning, validating the generalization ability of the proposed sketch control method.

TABLE III

RMSE BETWEEN SYNTHESIZED SPEECH AND REFERENCE SPEECH

| Method                            | RMSE $\downarrow$ |             |
|-----------------------------------|-------------------|-------------|
|                                   | Pitch (Hz)        | Energy (dB) |
| DrawSpeech with text only         | 110.65            | 19.04       |
| DrawSpeech with sketch conditions | <b>62.48</b>      | <b>9.48</b> |

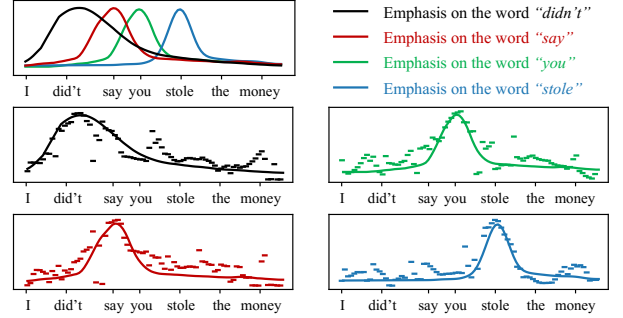


Fig. 3. Drawing different sketches to achieve precise prosody control. The solid line represents the drawn pitch sketch. The points indicate the pitch of each frame in the synthesized speech.

2) *Impact on sound quality:* Table II presents the MOS scores of DrawSpeech with and without the sketch conditions. The samples generated by DrawSpeech under sketch conditions even slightly surpass the ground truth speech, with a Wilcoxon rank-sum test [42] at a p-level of  $p > 0.05$ . We conclude that sketch conditions introduce sufficient prosody variations, enhancing the naturalness of the synthetic speech. The p-value serves as an indicator, demonstrating that DrawSpeech can synthesize speech of high quality or, as defined in [4], that DrawSpeech has achieved human-level quality.

3) *Precise prosody control:* When a reference speech is provided, we can extract the pitch and energy sketches from reference and use them as conditions to guide the generation process. As shown in Table III, DrawSpeech conditioned on the sketches of the reference speech achieves lower pitch and energy RMSE compared to the case without sketch conditions, indicating that DrawSpeech successfully imitate the prosody of the reference speech. To clearly demonstrate the fine-grained prosody control, we draw four different sketches to emphasize four different words in sentence “I didn’t say you stole the money”. The visualization shown in Fig. 3 suggests that the pitch variations in the synthesized speech closely follow the trends of the corresponding sketches. Also, the highest pitch is accurately assigned to the intended emphasis words, indicating that DrawSpeech can achieve prosody control precisely.

## IV. CONCLUSION

In this paper, we propose to use the pitch and energy sketches to control the synthesis process. Accordingly, we propose DrawSpeech to generate speech with rich prosody based on the sketches. Experimental results show the effectiveness of DrawSpeech and its capability to achieve precise prosody control. In the future, we plan to incorporate more control signals, such as jitter and shimmer, in a user-friendly manner.

## REFERENCES

- [1] R. S. Burdin, "Review: Prosodic patterns in english conversation," *Journal of English Linguistics*, vol. 48, no. 2, pp. 208–212, 2020.
- [2] S. Shechtman, R. Fernandez, and D. Haws, "Supervised and unsupervised approaches for controlling narrow lexical focus in sequence-to-sequence speech synthesis," in *IEEE Spoken Language Technology Workshop*, 2021, pp. 431–437.
- [3] X. Tan, T. Qin, F. Soong, and T.-Y. Liu, "A survey on neural speech synthesis," *arXiv preprint arXiv:2106.15561*, 2021.
- [4] X. Tan, J. Chen, H. Liu, J. Cong, C. Zhang, Y. Liu, X. Wang, Y. Leng, Y. Yi, L. He, S. Zhao, T. Qin, F. Soong, and T.-Y. Liu, "NaturalSpeech: End-to-end text-to-speech synthesis with human-level quality," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 6, pp. 4234–4245, 2024.
- [5] R. Liu, Y. Hu, Y. Ren, X. Yin, and H. Li, "Emotion rendering for conversational speech synthesis with heterogeneous graph-based context modeling," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 17, pp. 18 698–18 706, 2024.
- [6] Y. Lei, S. Yang, and L. Xie, "Fine-grained emotion strength transfer, control and prediction for emotional speech synthesis," in *IEEE Spoken Language Technology Workshop*, 2021, pp. 423–430.
- [7] D. Xin, X. Tan, K. Shen, Z. Ju, D. Yang, Y. Wang, S. Takamichi, H. Saruwatari, S. Liu, J. Li, and S. Zhao, "RALL-E: Robust codec language modeling with chain-of-thought prompting for text-to-speech synthesis," *arXiv preprint arXiv:2404.03204*, 2024.
- [8] X. Li, C. Song, J. Li, Z. Wu, J. Jia, and H. Meng, "Towards multi-scale style control for expressive speech synthesis," in *Interspeech*, 2021, pp. 4673–4677.
- [9] Y. Xiao, X. Wang, X. Tan, L. He, X. Zhu, S. Zhao, and T. Lee, "Contrastive context-speech pretraining for expressive text-to-speech synthesis," in *ACM International Conference on Multimedia*, 2024.
- [10] Y. A. Li, C. Han, and N. Mesgarani, "StyleTTS: A style-based generative model for natural and diverse text-to-speech synthesis," *arXiv preprint arXiv:2205.15439*, 2023.
- [11] Z. Ju, Y. Wang, K. Shen, X. Tan, D. Xin, D. Yang, E. Liu, Y. Leng, K. Song, S. Tang, Z. Wu, T. Qin, X. Li, W. Ye, S. Zhang, J. Bian, L. He, J. Li, and sheng zhao, "NaturalSpeech 3: Zero-shot speech synthesis with factorized codec and diffusion models," in *Forty-first International Conference on Machine Learning*, 2024.
- [12] K. Shen, Z. Ju, X. Tan, E. Liu, Y. Leng, L. He, T. Qin, sheng zhao, and J. Bian, "NaturalSpeech 2: Latent diffusion models are natural and zero-shot speech and singing synthesizers," in *International Conference on Learning Representations*, 2024.
- [13] A. Sorin, S. Shechtman, and R. Hoory, "Principal style components: Expressive style control and cross-speaker transfer in neural tts," in *Interspeech*, 2020, pp. 3411–3415.
- [14] X. Zhu, Y. Lei, K. Song, Y. Zhang, T. Li, and L. Xie, "Multi-speaker expressive speech synthesis via multiple factors decoupling," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2023, pp. 1–5.
- [15] S. Chen, S. Liu, L. Zhou, Y. Liu, X. Tan, J. Li, S. Zhao, Y. Qian, and F. Wei, "VALL-E 2: Neural codec language models are human parity zero-shot text to speech synthesizers," *arXiv preprint arXiv:2406.05370*, 2024.
- [16] A. Sigurgeirsson and S. King, "Controllable speaking styles using a large language model," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2024, pp. 10 851–10 855.
- [17] P. Anastassiou, J. Chen, J. Chen, Y. Chen, Z. Chen, Z. Chen, J. Cong, L. Deng, C. Ding, L. Gao, M. Gong, P. Huang, Q. Huang, Z. Huang, Y. Huo, D. Jia, C. Li, F. Li, H. Li, J. Li, X. Li, X. Li, L. Liu, S. Liu, S. Liu, X. Liu, Y. Liu, Z. Liu, L. Lu, J. Pan, X. Wang, Y. Wang, Y. Wang, Z. Wei, J. Wu, C. Yao, Y. Yang, Y. Yi, J. Zhang, Q. Zhang, S. Zhang, W. Zhang, Y. Zhang, Z. Zhao, D. Zhong, and X. Zhuang, "Seed-TTS: A family of high-quality versatile speech generation models," *arXiv preprint arXiv:2406.02430*, 2024.
- [18] D. Yang, S. Liu, R. Huang, C. Weng, and H. Meng, "InstructTTS: Modelling expressive tts in discrete latent space with natural language style prompt," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 32, pp. 2913–2925, 2024.
- [19] T. Bott, F. Lux, and N. T. Vu, "Controlling emotion in text-to-speech with natural language prompts," in *Interspeech*, 2024, pp. 1795–1799.
- [20] G. Liu, Y. Zhang, Y. Lei, Y. Chen, R. Wang, L. Xie, and Z. Li, "PromptStyle: Controllable style transfer for text-to-speech with natural language descriptions," in *Interspeech*, 2023, pp. 4888–4892.
- [21] Z. Guo, Y. Leng, Y. Wu, S. Zhao, and X. Tan, "Prompttts: Controllable text-to-speech with text descriptions," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2023, pp. 1–5.
- [22] D. Yang, J. Tian, X. Tan, R. Huang, S. Liu, H. Guo, X. Chang, J. Shi, sheng zhao, J. Bian, Z. Zhao, X. Wu, and H. M. Meng, "UniAudio: Towards universal audio generation with large language models," in *Forty-first International Conference on Machine Learning*, 2024.
- [23] Z. Du, Q. Chen, S. Zhang, K. Hu, H. Lu, Y. Yang, H. Hu, S. Zheng, Y. Gu, Z. Ma *et al.*, "Cosyvoice: A scalable multilingual zero-shot text-to-speech synthesizer based on supervised semantic tokens," *arXiv preprint arXiv:2407.05407*, 2024.
- [24] T. SpeechTeam, "FunAudioLLM: Voice understanding and generation foundation models for natural interaction between humans and llms," *arXiv preprint arXiv:2407.04051*, 2024.
- [25] C.-Y. Kuan, C.-A. Li, T.-Y. Hsu, T.-Y. Lin, H.-L. Chung, K.-W. Chang, S.-Y. Chang, and H.-Y. Lee, "Towards general-purpose text-instruction-guided voice conversion," in *IEEE Automatic Speech Recognition and Understanding Workshop*, 2023, pp. 1–8.
- [26] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Conference of the North American Chapter of the Association for Computational Linguistics*, 2019, pp. 4171–4186.
- [27] R. W. Frick, "Communicating emotion: The role of prosodic features," *Psychological bulletin*, vol. 97, no. 3, pp. 412–429, 1985.
- [28] G. Degano, P. W. Donhauser, L. Gwilliams, P. Merlo, and N. Golestani, "Speech prosody enhances the neural processing of syntax," *Communications Biology volume*, vol. 7, no. 748, pp. 1–10, 2024.
- [29] A. Savitzky and M. J. Golay, "Smoothing and differentiation of data by simplified least squares procedures," *Analytical chemistry*, vol. 36, no. 8, pp. 1627–1639, 1964.
- [30] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *International Conference on Learning Representations*, 2014.
- [31] 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 *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 2019, pp. 3171–3180.
- [32] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 10 674–10 685.
- [33] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pp. 6840–6851, 2020.
- [34] H. Liu, Z. Chen, Y. Yuan, X. Mei, X. Liu, D. Mandic, W. Wang, and M. D. Plumbley, "AudioLDM: Text-to-audio generation with latent diffusion models," in *Forty International Conference on Machine Learning*, 2023, pp. 21 450–21 474.
- [35] K. Ito and L. Johnson, "The lj speech dataset," <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [36] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [37] H. Liu, Y. Yuan, X. Liu, X. Mei, Q. Kong, Q. Tian, Y. Wang, W. Wang, Y. Wang, and M. D. Plumbley, "AudioLDM 2: Learning holistic audio generation with self-supervised pretraining," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- [38] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pp. 17 022–17 033, 2020.
- [39] P. K. Diederik and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations*, 2015.
- [40] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *international Conference on Learning Representations*, 2019.
- [41] 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 *International Conference on Learning Representations*, 2021.
- [42] F. Wilcoxon, "Individual comparisons by ranking methods," in *Breakthroughs in statistics*, 1992, pp. 196–202.