

Fine-Grained Quantitative Emotion Editing for Speech Generation

Sho Inoue*, Kun Zhou[†], Shuai Wang*[§], and Haizhou Li*,

* School of Data Science, Shenzhen Research Institute of Big Data, The Chinese University of Hong Kong, Shenzhen, China

[†] Speech Lab, Institute for Intelligent Computing, Alibaba Group, Singapore

E-mail: shoinoue@link.cuhk.edu.cn, kun.z@alibaba-inc.com, wangshuai@cuhk.edu.cn, and haizhouli@cuhk.edu.cn

Abstract—It remains a significant challenge how to quantitatively control the expressiveness of speech emotion in speech generation. In this work, we propose an approach for quantitative manipulation of the emotion rendering for emotion editing in speech generation. We apply a hierarchical emotion distribution extractor, i.e. Hierarchical ED, that quantifies the intensity of emotions at different levels of granularity. Hierarchical ED is subsequently integrated into the FastSpeech2 framework, guiding the model to learn emotion intensity at phoneme, word, and utterance levels. During synthesis, users can manually edit the emotional intensity of the generated voices. Both objective and subjective evaluations demonstrate the effectiveness of the proposed network in terms of fine-grained quantitative emotion editing.

I. INTRODUCTION

Speech emotion, characterized by prosodic patterns at phoneme, word, and utterance level [1]–[3], is effectively represented through a hierarchical structure [4], [5]. These patterns, including pitch, tempo, rhythm, stress, lexical emphasis, and speaker individuality, are crucial for accurate emotion rendering [6]–[9]. Existing methods, like style tokens, represent prosody patterns but often lack interpretability, quantitiveness, and fine-grained controllability [10], [11], highlighting the need for more effective models. In neural text-to-speech, speech editing techniques empower models to alter audio segments based on user instructions, enabling fine-grained prosody prediction in synthesized speech [12], [13]. However, this approach is deficient in quantitative control.

The complexity of speech emotions, as evidenced by various studies [14]–[17], indicates that merely adjusting prosodic features is inadequate. Additionally, the hierarchical nature of speech emotion [6], with unique characteristics at utterance [18], word [19], and phoneme [6] levels, calls for a multi-level approach for accurate and human-like emotional expression in speech synthesis. Previous studies on text-to-speech and emotional voice conversion [17], [20], [21] also highlight the importance of multi-level modeling for speech

emotions. Thus, a quantifiable method capable of modeling the hierarchical structure of emotions is imperative for effective and nuanced emotion rendering in speech synthesis.

In our previous study [22], we investigated a fine-grained emotion control approach for text-to-speech (TTS) systems. In this paper, we further extend the approach to enable quantitative emotion editing. This provides an intuitive and quantitative method for manipulating emotion rendering at the phoneme, word, and utterance levels. Compared to our previous work [22], which integrated fine-grained embedding solely into the variance adaptor of FastSpeech2, our proposed emotion editing approach allows for fine-grained and flexible human editing across any TTS model.

Our contributions are summarized as follows:

- We propose a novel fine-grained emotion editing framework that facilitates the emotion rendering in previously unseen audios and allows for the adjustment of emotion intensity for a speech generation framework;
- We thoroughly compare our approach with the baseline method of phoneme-level emotion intensity control, demonstrating our model’s superior emotion expressiveness and controllability.
- At run-time, users have the flexibility to analyze and manipulate the emotion distribution from the audio sample (“emotion editing”). This approach empowers users with a quantitative method to modify emotion rendering of a spoken utterance.

The rest of this paper is organized as follows: In Section 2, we discuss the related works. Section 3 describes our proposed methodology. In Section 4, we report our experiments and results. Section 5 concludes our study.

II. RELATED WORKS

A. Control of Speech Emotion

There has been a growing interest in enabling the control of synthesized emotions [23]–[25]. Emotion-enhanced GST [23] incorporated an emotion recognition task to facilitate the modeling of emotion-related prosody. Some studies explored quantitative methods for controlling emotions through relative attributes [26]–[29]. Another study [30] delved into inter- and intra-class distances to achieve fine-grained control over recognizable intensity differences. EmoQ-TTS [31] adopted

[§]Corresponding Author

The research is supported by National Natural Science Foundation of China (Grant No. 62271432); Internal Project of Shenzhen Research Institute of Big Data (Grant No. T00120220002); Shenzhen Science and Technology Program ZDSYS20230626091302006; Shenzhen Science and Technology Research Fund (Fundamental Research Key Project Grant No. JCYJ20220818103001002); and CCF-NetEase ThunderFire Innovation Research Funding (No. CCF-Netease 202302).

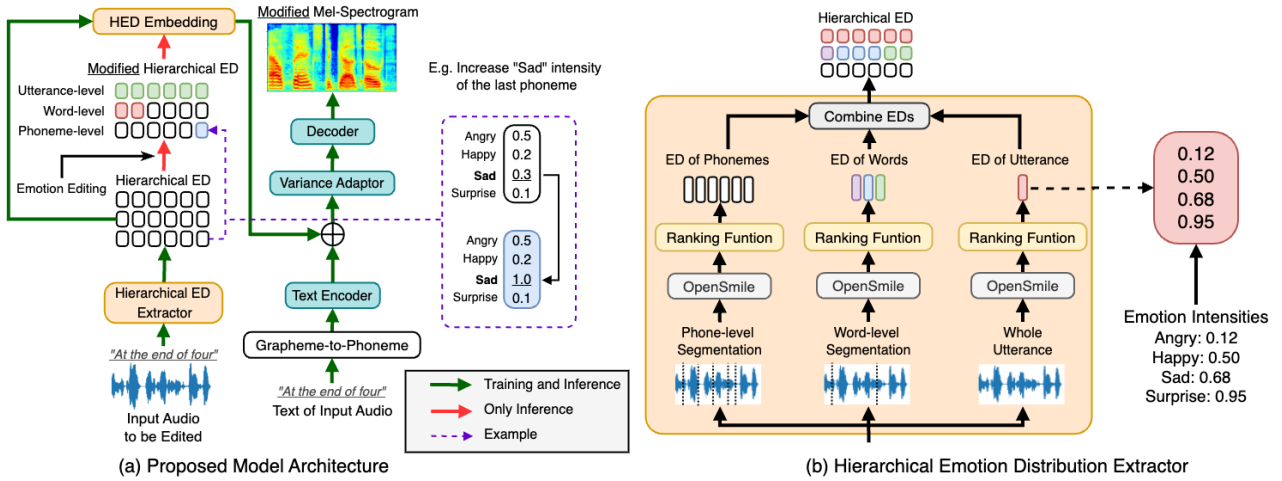


Fig. 1: (a) Model architecture with hierarchical emotion distribution (HED) mechanism in TTS. During inference, the framework extracts hierarchical ED from input audio (“emotion editing”). Users can manually modify hierarchical ED to control emotion intensities at phoneme, word, and utterance levels; (b) Hierarchical ED extraction workflow for emotion distributions at phoneme, word, and utterance levels.

a distance-based intensity quantization approach to capture emotion intensity.

Unlike MsEmoTTS [32], which covers only utterance-level emotional change with fine-grained intensity changes, this work aims to offer quantitative and hierarchical control over emotions at various levels of speech units, e.g. phoneme, word, and utterance. It provides a versatile tool for editing emotion rendering in any speech segment.

B. Speech Editing

Speech editing, particularly semantic editing, focuses on altering the textual information of speech while preserving the naturalness of the synthesized output [12], [13], [33]. EditSpeech [33] employed partial inference and bidirectional fusion techniques to effectively modify speech semantics. Another research, EdiTTS [12], utilized perturbations in the Gaussian space to allow for detailed audio edits, adjusting context and pitch, and enabling user-guided prosody editing. Yet another study [13] presented a method for phoneme-level pitch adjustment and time-stretching, providing users with greater control over prosody.

This work is motivated to provide a systematic approach for precisely controlling high-level prosodic patterns, which differs from the prior studies which primarily focused on modifying the physical speech attributes, e.g. pitch and duration.

In this section, we delve into the details of the proposed fine-grained emotion editing mechanism. We begin by formulating the problems and introducing a quantitative and hierarchical emotion control mechanism. We then detail the design of the hierarchical emotion distribution (ED) extractor, along with its associated training scheme. Lastly, we explain how to render the desired emotions by editing.

C. Problem Formulation

Given an audio input and its text transcript, we would like to develop a model to control the emotion intensity of the speech segments, as illustrated in Fig.1(a). This approach is expected to work for the rendering of a single emotion or a mixture of emotions and to work seamlessly with various text-to-speech frameworks. Typically, we label emotions at the utterance level in most speech databases without paying attention to the nuanced variations in emotion intensity. It remains a significant challenge as to how to achieve effective and quantifiable control over speech emotion. To address this, we aim to automatically generate fine-grained and quantitative intensity labels that act as ‘soft labels’ for speech generation models. In this way, we eliminate the need for manual labeling of emotional intensity. Our approach is effective in emotion intensity control and mixed-emotion rendering, which can be easily adapted to any speech generation frameworks including text-to-speech and voice conversion.

D. Hierarchical Emotion Distribution Extractor

Following [22], we apply a hierarchical emotion distribution extraction module consisting of an OpenSmile feature extractor and a pre-trained ranking function for each segmental level, which automatically quantifies the intensities of each emotion type in an utterance. This approach, grounded in the concept of relative attributes [34], allows us to measure the prominence of specific emotions in speech. By treating emotion style as a speech attribute, we model and rank the presence of various emotions relative to the other emotions.

In particular, we define the ranking function as $f(\mathbf{x}_i) = \mathbf{w}^T \mathbf{x}_i + b$, where \mathbf{x}_i , \mathbf{w} , and b denote the acoustic features of the i -th training sample, the weight vector, and bias, respectively. The parameters (\mathbf{w} and b) are optimized using the support vector machine’s objective function for binary classification (e.g. Angry and Non-angry classification) [35].

To provide a quantifiable measure, we further normalize the values obtained from the ranking function to a range of $[0, 1]$, where a larger value represents a stronger emotion intensity. Our methodology facilitates the labeling of training data with continuous emotion intensities across various emotions. Additionally, it allows for the quantification and regulation of emotion intensity in unseen utterances during run-time.

The hierarchical ED extractor consists of pre-trained ranking functions as illustrated in Fig.1(b). The extractor first segments the input audio signal into phoneme, word, and utterance levels via Montreal Forced Aligner [36]. Then, it extracts an 88-dimensional audio feature set for each temporal segment with OpenSMILE [37]. The pre-trained ranking function automatically estimates an ED vector for each audio segment. Each value in this vector represents the intensity of a specific emotion encoded in the audio segment. In practice, the utterance-level emotion distribution is duplicated across all phonemes, and the word-level emotion distribution is replicated across the corresponding phonemes. Once formulated, the hierarchical ED vectors are integrated with linguistic embeddings in the variance adaptor during training.

E. Hierarchical Emotion Intensity Modeling

We integrate the aforementioned hierarchical ED extractor into the FastSpeech2 framework, which comprises a text encoder, variance adaptor, and decoder, as depicted in Fig.1(a). During training, given an input pair $\langle \text{audio}, \text{text} \rangle$, this extractor predicts the emotion intensity at phoneme, word, and utterance levels from the audio. These predictions form a hierarchical ED as illustrated in Fig.1 (b). The text encoder converts the phoneme sequence into text embeddings. The variance adaptor, utilizing both text and ED embeddings, predicts pitch, duration, and energy. Subsequently, the decoder reconstructs the Mel-spectrogram, guided by an L1 loss function. This process enables the framework to learn the hierarchical nature of emotion intensity and to establish a link between linguistic and hierarchical ED information.

F. Emotion Editing

During run-time, as depicted in Fig. 1(a), our framework enables emotion editing. For unseen $\langle \text{audio}, \text{text} \rangle$ inputs, it derives the hierarchical Emotion Distribution (ED) from the audio signal. Users can then edit the emotion rendering by adjusting emotion distributions at three different levels (Hierarchical ED).

III. EXPERIMENTS AND RESULTS

A. Experimental Setup

We conduct our experiments on two datasets: Blizzard Challenge 2013 dataset, i.e. Blizzard [38], and Emotion Speech Dataset, i.e. ESD [39], [40]. Blizzard, audiobooks read by a single female speaker, offers approximately 42 hours of expressive speech with varied prosody but lacks emotion labels. We randomly select 200 samples from four audiobooks for testing. On the other hand, ESD comprises over 29 hours of emotional speech in five categories—Neutral, Angry, Happy,

TABLE I: MOS with 95% confidence interval

Ground Truth	MsEmoTTS	Hierarchical ED (ours)
4.211±0.142	3.544±0.138	3.596±0.141

Sad, and Surprise—from 20 speakers, including 10 native English and 10 Mandarin speakers. Our study exclusively employs English recordings. For training the hierarchical ED extractor’s ranking functions, we select 20 random samples per speaker and emotion from ESD.

B. Model Architecture

We use FastSpeech2 as the backbone framework [41], which consists of a text encoder, variance adaptor, and decoder. A transformer [42] based text encoder converts an input phoneme sequence into a linguistic embedding. Variance adaptors use feed-forward networks for the duration, pitch, and energy prediction. A transformer-based decoder synthesizes a mel-spectrogram from the variance-adapted features. The loss function combines the L1 loss between the predicted and target mel-spectrogram and the mean squared error loss from predicted prosodic features. We opt for the Adam optimizer [43], setting a batch size of 32 and conducting 800,000 iterations of training over 48 hours on a single GPU. The ED embedding layers comprise fully connected layers with a Tanh activation function. We adopt the text-based emotion intensity predictor from MsEmoTTS [32], which uses two 1D convolution layers, layer normalization, and dropout. HiFiGAN [44] serves as the vocoder, pre-trained on the Blizzard dataset.

C. Results and Analysis

We conduct both objective and subjective evaluations, focusing on speech quality, emotional expressiveness, and emotion controllability. We integrate MsEmoTTS [32] into the FastSpeech2 framework as the baseline. Our subjective evaluation involves listening experiments with 18 participants, each listening to 120 synthesized samples under specific guidelines. We suggest readers to refer to our demo page¹.

1) *Speech Quality*: We conduct a Mean Opinion Score (MOS) test to evaluate the overall speech quality, where a higher MOS represents better speech quality. As shown in Table I, our model demonstrates superior performance over the baseline, as indicated by its consistently higher MOS scores.

2) *Emotion Expressiveness*: We conduct A/B preference tests where participants select the audio that more closely matches the reference in terms of emotional expressiveness. Both our model’s hierarchical ED and the baseline’s emotion intensity are derived from the same reference audio. As indicated in Table II, our model outperforms the baseline with a preference rate of 43.51%.

To further objectively evaluate emotion expressiveness, we calculate various metrics: (1) Mel-Cepstral Distortion

¹Speech Demos: <https://shinshoji01.github.io/Hierarchical-ED-Demo/>

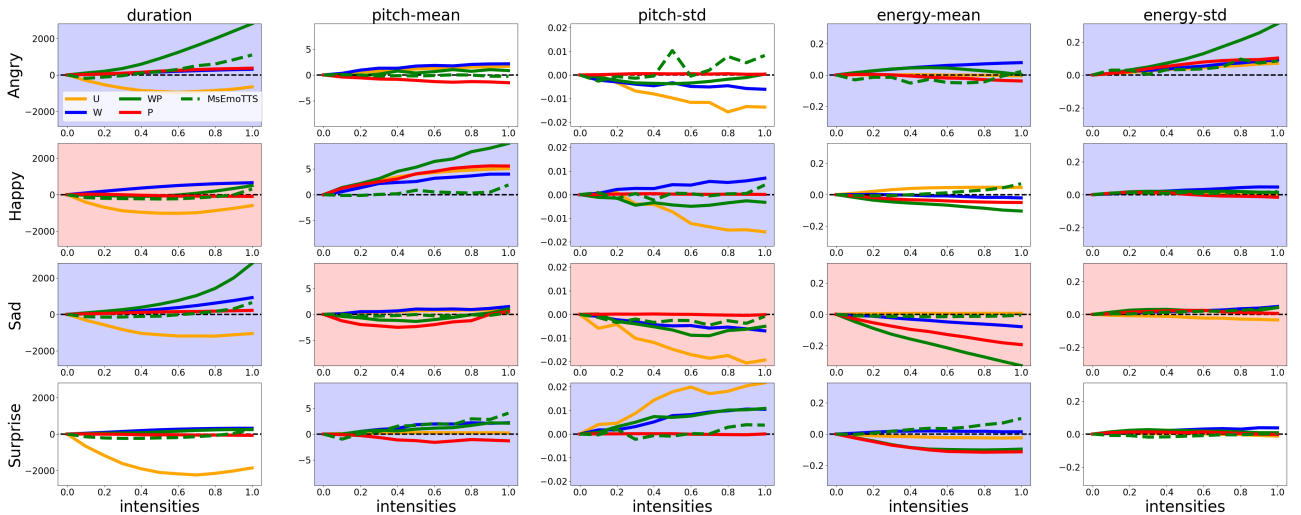


Fig. 2: The illustration of prosodic variants with intensity changes. The red background represents the expected negative trend, the blue indicates the expected positive trend, both summarized from the ESD dataset. ‘U’, ‘W’, ‘P’, and ‘WP’ indicate the controlling segments in Hierarchical ED (our model): Utterance, Word, Phoneme, and Word-and-Phoneme, respectively.

TABLE II: A/B preference test for emotion expressiveness with 95 % confidence interval.

MsEmoTTS	Neutral	Hierarchical ED (ours)
28.77±5.256	27.72±5.197	43.51±5.756

TABLE III: Objective evaluation for emotion expressiveness

	MCD	Pitch (10^1)	Energy (10^{-2})	FD (10^1)
MsEmoTTS	4.783±0.135	1.215±0.064	4.884±0.225	7.841±0.842
HED (ours)	4.348±0.109	1.151±0.052	4.018±0.162	6.881±0.710

(MCD)[45] to measure spectral similarity, (2) Pitch/Energy Distortion for assessing prosody alignment, and (3) Frame Disturbance (FD)[46] to evaluate duration deviation. The results summarized in Table III confirms our framework’s superiority over the baseline MsEmoTTS in replicating target emotions, thereby validating its effectiveness in emotion modeling.

3) *Emotion Controllability*: We conducted best-worst scaling (BWS) tests [47] to compare word-level emotion controllability between our model and the baseline. In these tests, we modified the emotion intensity of three words in each utterance to 0.0, 0.5, and 1.0. Evaluators listened to the audio samples and selected the least and most expressive ones. As shown in Table IV, our model demonstrated a distinct preference for the least expressive sample at the lowest intensity and the most expressive sample at the highest intensity, particularly for Sad and Surprise emotions. This result highlights our model’s capability to effectively differentiate intensity variations.

We further validated the controllability of our model across fine-grained emotional variances, examining them at utterance, word, phoneme, and word-and-phoneme combination levels. We increased the emotion intensity values from 0.0 to 1.0

TABLE IV: BWS Test Result: The value represents evaluator preferences (%), with red and blue indicating the heatmap for the least expressive and most expressive audio, respectively.

	Hierarchical ED (ours)				MsEmoTTS				
	Ang	Hap	Sad	Sur	Ang	Hap	Sad	Sur	
Least	0.0	79	63	67	81	42	32	21	33
0.5	0	0	28	14	14	47	47	54	42
1.0	21	9	9	19	5	11	21	25	25
Most	0.0	11	18	16	7	11	12	30	18
0.5	16	16	7	9	7	26	16	23	30
1.0	74	75	75	75	86	63	72	47	53

and subsequently computed various prosody features, such as duration, mean/standard deviation of pitch, and mean/standard deviation of energy, shown in Fig.2. These prosodic characteristics are known to correlate significantly with emotion intensity, as indicated in the literature [5] For example, sadness often manifests in a slower speaking rate and lower values for pitch and energy mean/standard deviation. We performed an analysis of ESD to explore the relationship between acoustic features and various emotions, using different colors in Figure 2 to indicate positive and negative correlations. A red background represents an expected negative trend, while a blue background signifies a positive trend as intensity increases. We have observed that our proposed model closely aligns with the expected trends. For example, ESD reveals a positive correlation between happiness and mean pitch, as well as a negative correlation between sadness and mean energy. Our models can effectively adjust these features as emotion intensity varies.

Furthermore, we observe significant prosodic changes when editing both word and phoneme-level emotions and the changes in the standard deviation of pitch at the utterance level align with our expectations. The figure also illustrates that

Hierarchical ED outperforms the baseline model, especially when controlling word and phoneme levels simultaneously. Our model exhibits a significantly closer alignment with the expected trend.

IV. CONCLUSION

We introduce a fine-grained emotion editing approach for speech generation tasks. At run-time, users can quantitatively control emotion rendering at phoneme, word, and utterance levels. Both objective and subjective evaluations validate the effectiveness of our proposed idea. As a future work, we aim to enhance paragraph-level information processing and develop emotional text-to-speech models featuring robust emotion consistency. Additionally, the Hierarchical ED framework holds potential for application in diverse speech generation frameworks and contexts, including voice conversion tasks.

REFERENCES

- [1] A. Triantafyllopoulos, B. W. Schuller, G. İymen, *et al.*, “An overview of affective speech synthesis and conversion in the deep learning era,” *Proceedings of the IEEE*, 2023.
- [2] Z. KUN, “Emotion modelling for speech generation,” 2022.
- [3] J. Hirschberg, “Pragmatics and intonation,” *The handbook of pragmatics*, pp. 515–537, 2006.
- [4] M. El Ayadi, M. S. Kamel, and F. Karray, “Survey on speech emotion recognition: Features, classification schemes, and databases,” *Pattern recognition*, vol. 44, no. 3, pp. 572–587, 2011.
- [5] B. W. Schuller, “Speech emotion recognition: Two decades in a nutshell, benchmarks, and ongoing trends,” *Communications of the ACM*, vol. 61, no. 5, pp. 90–99, 2018.
- [6] S. R. Krothapalli and S. G. Koolagudi, “Emotion recognition using prosodic information,” in *Emotion Recognition using Speech Features*. New York, NY: Springer New York, 2013, pp. 79–91, ISBN: 978-1-4614-5143-3. DOI: 10.1007/978-1-4614-5143-3_5. [Online]. Available: https://doi.org/10.1007/978-1-4614-5143-3_5.
- [7] K. Zhou, B. Sisman, M. Zhang, and H. Li, “Converting anyone’s emotion: Towards speaker-independent emotional voice conversion,” *Proc. Interspeech 2020*, pp. 3416–3420, 2020.
- [8] K. Zhou, B. Sisman, and H. Li, “Vaw-gan for disentanglement and recombination of emotional elements in speech,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*, IEEE, 2021, pp. 415–422.
- [9] Z. Du, B. Sisman, K. Zhou, and H. Li, “Disentanglement of emotional style and speaker identity for expressive voice conversion,” *arXiv preprint arXiv:2110.10326*, 2021.
- [10] C. Du and K. Yu, *Rich prosody diversity modelling with phone-level mixture density network*, 2021. arXiv: 2102.00851 [cs.SD].
- [11] N. Tits, *Controlling the emotional expressiveness of synthetic speech: A deep learning approach*, 2022.
- [12] J. Tae, H. Kim, and T. Kim, *Editts: Score-based editing for controllable text-to-speech*, 2022. arXiv: 2110.02584 [cs.SD].
- [13] M. Morrison, L. Rencker, Z. Jin, N. J. Bryan, J.-P. Caceres, and B. Pardo, *Context-aware prosody correction for text-based speech editing*, 2021. arXiv: 2102.08328 [eess.AS].
- [14] K. Zhou, B. Sisman, C. Busso, B. Ma, and H. Li, *Mixed-enc: Mixed emotion synthesis and control in voice conversion*, 2023. arXiv: 2210.13756 [eess.AS].
- [15] Y. Xu, “Speech prosody: A methodological review,” *Journal of Speech Sciences*, vol. 1, no. 1, pp. 85–115, 2011.
- [16] J. Latorre and M. Akamine, “Multilevel parametric-base f0 model for speech synthesis,” in *Ninth Annual Conference of the International Speech Communication Association*, 2008.
- [17] K. Zhou, B. Sisman, and H. Li, “Transforming spectrum and prosody for emotional voice conversion with non-parallel training data,” in *Proc. Odyssey 2020 The Speaker and Language Recognition Workshop*, 2020, pp. 230–237.
- [18] E. Rodero, “Intonation and emotion: Influence of pitch levels and contour type on creating emotions,” *Journal of Voice*, vol. 25, no. 1, e25–e34, 2011, ISSN: 0892-1997. DOI: <https://doi.org/10.1016/j.jvoice.2010.02.002>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0892199710000378>.
- [19] J. Hirschberg and G. Ward, “The influence of pitch range, duration, amplitude and spectral features on the interpretation of the rise-fall-rise intonation contour in english,” *Journal of Phonetics*, vol. 20, no. 2, pp. 241–251, 1992, ISSN: 0095-4470. DOI: [https://doi.org/10.1016/S0095-4470\(19\)30625-4](https://doi.org/10.1016/S0095-4470(19)30625-4). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0095447019306254>.
- [20] H. Ming, D.-Y. Huang, L. Xie, J. Wu, M. Dong, and H. Li, “Deep bidirectional lstm modeling of timbre and prosody for emotional voice conversion,” in *Interspeech*, 2016, pp. 2453–2457.
- [21] S. Lei, Y. Zhou, L. Chen, *et al.*, *Towards multi-scale speaking style modelling with hierarchical context information for mandarin speech synthesis*, 2022. arXiv: 2204.02743 [cs.SD].
- [22] S. Inoue, K. Zhou, S. Wang, and H. Li, “Hierarchical emotion prediction and control in text-to-speech synthesis,” in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 10 601–10 605. DOI: 10.1109/ICASSP48485.2024.10445996.
- [23] X. Cai, D. Dai, Z. Wu, X. Li, J. Li, and H. Meng, *Emotion controllable speech synthesis using emotion-unlabeled dataset with the assistance of cross-domain*

- speech emotion recognition*, 2021. arXiv: 2010.13350 [eess.AS].
- [24] T. Li, S. Yang, L. Xue, and L. Xie, *Controllable emotion transfer for end-to-end speech synthesis*, 2020. arXiv: 2011.08679 [cs.SD].
- [25] A. Rabiee, T.-H. Kim, and S.-Y. Lee, “Adjusting pleasure-arousal-dominance for continuous emotional text-to-speech synthesizer,” in *Interspeech*, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:189762268>.
- [26] X. Zhu, S. Yang, G. Yang, and L. Xie, “Controlling emotion strength with relative attribute for end-to-end speech synthesis,” in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019, pp. 192–199. DOI: 10.1109/ASRU46091.2019.9003829.
- [27] K. Zhou, B. Sisman, R. Rana, B. W. Schuller, and H. Li, “Emotion intensity and its control for emotional voice conversion,” *IEEE Transactions on Affective Computing*, vol. 14, no. 1, pp. 31–48, 2023. DOI: 10.1109/TAFFC.2022.3175578.
- [28] K. Zhou, B. Sisman, R. Rana, B. W. Schuller, and H. Li, “Speech synthesis with mixed emotions,” *IEEE Transactions on Affective Computing*, vol. 14, no. 4, pp. 3120–3134, 2023. DOI: 10.1109/TAFFC.2022.3233324.
- [29] Y. Lei, S. Yang, and L. Xie, *Fine-grained emotion strength transfer, control and prediction for emotional speech synthesis*, 2020. arXiv: 2011.08477 [cs.SD].
- [30] S. Wang, J. Guonason, and D. Borth, “Fine-grained emotional control of text-to-speech: Learning to rank inter- and intra-class emotion intensities,” in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5. DOI: 10.1109/ICASSP49357.2023.10097118.
- [31] C.-B. Im, S.-H. Lee, S.-B. Kim, and S.-W. Lee, “Emoqts: Emotion intensity quantization for fine-grained controllable emotional text-to-speech,” in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6317–6321. DOI: 10.1109/ICASSP43922.2022.9747098.
- [32] Y. Lei, S. Yang, X. Wang, and L. Xie, *Msemotts: Multi-scale emotion transfer, prediction, and control for emotional speech synthesis*, 2022. arXiv: 2201.06460 [cs.SD].
- [33] D. Tan, L. Deng, Y. T. Yeung, X. Jiang, X. Chen, and T. Lee, *Editspeech: A text based speech editing system using partial inference and bidirectional fusion*, 2021. arXiv: 2107.01554 [eess.AS].
- [34] D. Parikh and K. Grauman, “Relative attributes,” in *2011 International Conference on Computer Vision*, IEEE, 2011, pp. 503–510.
- [35] C. Cortes and V. Vapnik, “Support-vector networks,” *Machine learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [36] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, “Montreal Forced Aligner: Trainable Text-Speech Alignment Using Kaldi,” in *Proc. Interspeech 2017*, 2017, pp. 498–502. DOI: 10.21437/Interspeech.2017-1386.
- [37] F. Eyben, M. Wöllmer, and B. Schuller, “Opensmile – the munich versatile and fast open-source audio feature extractor,” Jan. 2010, pp. 1459–1462. DOI: 10.1145/1873951.1874246.
- [38] S. King and V. Karaiskos, “The blizzard challenge 2013,” 2014.
- [39] K. Zhou, B. Sisman, R. Liu, and H. Li, “Emotional voice conversion: Theory, databases and esd,” *Speech Communication*, vol. 137, pp. 1–18, 2022.
- [40] K. Zhou, B. Sisman, R. Liu, and H. Li, “Seen and unseen emotional style transfer for voice conversion with a new emotional speech dataset,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2021, pp. 920–924.
- [41] Y. Ren, C. Hu, X. Tan, *et al.*, *Fastspeech 2: Fast and high-quality end-to-end text to speech*, 2022. arXiv: 2006.04558 [eess.AS].
- [42] A. Vaswani, N. Shazeer, N. Parmar, *et al.*, *Attention is all you need*, 2017. arXiv: 1706.03762 [cs.CL].
- [43] D. P. Kingma and J. Ba, *Adam: A method for stochastic optimization*, 2017. arXiv: 1412.6980 [cs.LG].
- [44] J. Kong, J. Kim, and J. Bae, *Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis*, 2020. arXiv: 2010.05646 [cs.SD].
- [45] R. Kubichek, “Mel-cepstral distance measure for objective speech quality assessment,” in *Proceedings of IEEE Pacific Rim Conference on Communications Computers and Signal Processing*, vol. 1, 1993, 125–128 vol.1. DOI: 10.1109/PACRIM.1993.407206.
- [46] B. Sisman, G. Lee, H. Li, and K. C. Tan, “On the analysis and evaluation of prosody conversion techniques,” in *2017 International Conference on Asian Language Processing (IALP)*, 2017, pp. 44–47. DOI: 10.1109/IALP.2017.8300542.
- [47] S. Kiritchenko and S. Mohammad, “Best-worst scaling more reliable than rating scales: A case study on sentiment intensity annotation,” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, R. Barzilay and M.-Y. Kan, Eds., Vancouver, Canada: Association for Computational Linguistics, Jul. 2017, pp. 465–470. DOI: 10.18653/v1/P17-2074. [Online]. Available: <https://aclanthology.org/P17-2074>.