



Language Model-Based Emotion Prediction Methods for Emotional Speech Synthesis Systems

Hyun-Wook Yoon¹, Ohsung Kwon¹, Hoyeon Lee¹, Ryuichi Yamamoto²,
Eunwoo Song¹, Jae-Min Kim¹, and Min-Jae Hwang¹

¹NAVER Corp., Seongnam, Korea, ²LINE Corp., Tokyo, Japan
hyunwook.yoon@navercorp.com

Abstract

This paper proposes an effective emotional text-to-speech (TTS) system with a pre-trained language model (LM)-based emotion prediction method. Unlike conventional systems that require auxiliary inputs such as manually defined emotion classes, our system directly estimates emotion-related attributes from the input text. Specifically, we utilize generative pre-trained transformer (GPT)-3 to jointly predict both an emotion class and its strength in representing emotions' coarse and fine properties, respectively. Then, these attributes are combined in the emotional embedding space and used as conditional features of the TTS model for generating output speech signals. Consequently, the proposed system can produce emotional speech only from text without any auxiliary inputs. Furthermore, because the GPT-3 enables to capture emotional context among the consecutive sentences, the proposed method can effectively handle the paragraph-level generation of emotional speech.

Index Terms: Text-to-speech (TTS), emotional TTS, emotion modeling from text, language model, GPT-3

1. Introduction

As the quality of the neural text-to-speech (TTS) system has reached natural sound almost indistinguishable from human recordings [1–3], the interest in emotional TTS has also increased [4–8]. One of the general approaches is using reference audio [7–10], where the emotional style representation, i.e., emotion embedding, is extracted from the emotional reference audio and then used as an auxiliary input of the TTS system.

Directly utilizing natural speech's emotion information as the conditional input can effectively deliver the expressiveness to synthetic outputs. However, the overall cost of deciding the proper emotional attributes in each sentence can be high since these methods require a continuous selection process of auxiliary inputs during an inference. For instance, when synthesizing a bulk of sentences, such as audiobook contents [11], the entire TTS process becomes expensive since a large number of emotions need to be annotated properly before synthesizing the corresponding waveforms.

To address the aforementioned problem, we propose an language model (LM)-based automatic emotion prediction method. From the text inputs, a *joint emotion predictor* estimates the emotion class and its strength to model the coarse and fine properties of the target emotion, respectively. Specifically, we adopt a generative pre-trained transformer (GPT)-3 [12] as the backbone network of the emotion predictor that enables to predict the emotional attributes without any auxiliary information. Then, a *joint emotion encoder* converts the emotion class and strength into the embedding space by linearly scaling the embedding vector of the emotion class according to its strength. Finally, the resulting embedding is fed to the TTS model to synthesize emotional output speech. Our contributions are summarized as follows:

- To achieve the emotional TTS system without relying on auxiliary inputs, we propose an LM-based emotion prediction method that estimates both the emotion class and strength of the target speech from input text. The experimental results showed that the proposed method estimated more accurate emotional attributes than the conventional method.
- We tested the proposed emotion prediction method in a neural TTS system based on a Tacotron 2 acoustic model [13–15]

with a Parallel WaveGAN (PWG) vocoder [16–19]. The experimental results verified that the proposed TTS system synthesized more natural and expressive speech than the conventional method, and was competitive with the system that used manually defined emotional attributes. In particular, our TTS system achieved results of 3.92 and 3.94 in the mean opinion score (MOS) tests, which evaluated the naturalness and emotional expressiveness of synthetic speech.

- Under the paragraph-level speech synthesis scenario, we showed that the proposed system was capable of replacing the emotional TTS system with human-annotated emotional attributes. A subjective preference test verified that these TTS systems were perceptually indistinguishable.

2. Related Works

In the emotional TTS tasks, several studies have proposed to predict emotional attributes from the text input. For instance, Cui et al. [20] proposed to predict an emotion class from the text by using a convolutional neural network (CNN) block. Lei et al. [21] proposed to predict emotion strength from the text by using the text encoder module of Tacotron [22]. Although these methods provided highly emotional speech, their applications were limited because they required manual definitions of the desired emotion class [20] or strength [21] during the inference phase, which makes the cost of generating emotional speech much higher.

Unlike those studies, our method differs in that it predicts both the emotion class and the emotion strength by using pre-trained LM. Note that Lei et al. [23] adopted a pre-trained BERT model [24] to predict emotional attributes from text, which is mostly similar to our work. However, the difference in our work is that we adopt GPT-3 instead of BERT, and further propose a joint emotion encoder that maximizes the merits in the usage of emotion strength.

3. Proposed Emotional TTS System

The proposed emotional TTS system is illustrated in Figure 1. First, the joint emotion predictor estimates the emotion class and strength from the input text. Then, the joint emotion encoder mixes them to reflect the joint representation of emotions' coarse and fine characteristics. Finally, this joint emotion embedding is fed to the TTS model to synthesize output emotional speech. The details of this process are described in the following sections.

3.1. Emotion class and strength annotation

Before training the joint emotion predictor, it is necessary to prepare the annotations of emotion class and strength from the recorded data. In the case of the emotion class annotation, we guide the speaker to act on the specific emotion during the recording process by assuming the speaker's acting is enough to appropriately represent the target emotion.

However, annotating the emotion strength is a relatively difficult task, since its subjective and relative properties make it hard to quantify. To overcome this problem, we adopt an approach similar to Lei et al. [25] based on a ranking support vector machine (RankSVM) [26], which is a well-known method for measuring relative attributes between binary classes. Specifically, the emotion features consisting of Mel-frequency cepstral coefficients (MFCC),

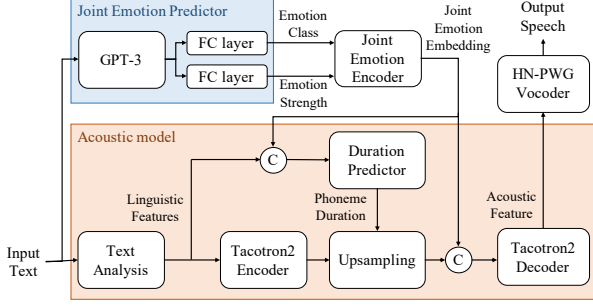


Figure 1: Block diagram of the proposed system; © denotes the concatenation process between input tensors.

fundamental frequency (F0), and so on are first extracted from the neutral and emotional speech. Then, a ranking function between these features is trained with the RankSVM objective [26, 27]. After computing a rank for the emotion of each speech signal, it is normalized to have a value between 0 (weakest) and 1 (strongest); we then define it as annotated emotion strength. Note that, based on the assumption that a neutral emotion has no strength, we fixed the strength of neutral emotion at zero. Consequently, it is possible to quantify the relative strength that represents the intensity of emotion compared to the neutral data.

3.2. Emotion class and strength prediction from text

To construct an efficient emotional TTS system, we propose a *joint emotion predictor* that estimates both the emotion class and its strength from input text. To accurately predict these attributes, we utilize GPT-3 as a feature extractor network of the emotion predictor [12, 28]. In particular, the proposed emotion predictor consists of the GPT-3 followed by two individual fully connected (FC) layers, where one FC layer is used to estimate the one-hot encoded emotion class and the other is used to estimate the scalar-valued emotion strength.

The advantages of the proposed emotion predictor are clear. First, due to the GPT-3 enables the model to accurately predict emotional attributes, it is capable of getting rid of the handwork for emotion annotation of the speech signal which makes the emotional TTS expensive. Additionally, thanks to the GPT-3’s ability to identify the contextual information of input sentences, it effectively considers the emotional context among multiple sentences.

3.3. Joint embedding of emotion class and strength

Once the emotion’s class and strength are generated, the *joint emotion encoder* blends them to compose their joint emotion embedding vector. First, the one-hot encoded emotion class is converted to a corresponding emotion class embedding through a look-up table (LUT). Then, the joint emotion embedding, \mathbf{h}_e , is obtained by mixing emotion class embedding, \mathbf{e}_{emb} , and the emotion strength, e_{str} , as follows:

$$\mathbf{h}_e = \text{softplus}(\mathbf{W}_{emb}\mathbf{e}_{emb} \cdot (1 + w_{str}e_{str})), \quad (1)$$

where \mathbf{W}_{emb} and w_{str} denote the projection matrix and scalar variable, respectively; $\text{softplus}(\cdot)$ denotes the softplus activation function [29]. Finally, the joint emotion embedding is fed to the TTS model for the synthesis of emotional speech.

By utilizing both emotion class and strength, the proposed method can effectively reflect both the coarse and fine structures of emotion properties. To show this, we drew the t-distributed stochastic neighbor embedding (t-SNE) [30] plot of joint emotion embedding, \mathbf{h}_e , which is illustrated in Figure 2. The result shows that each joint emotion embedding converges to the single vector point, which corresponds to the embedding of neutral emotion as the emotion strength weakens; whereas it diverges in a particular direction as the emotion strength increases. This indicates that each emotion class composes its own cluster in the emotional embedding space, which corresponds to the emotion’s coarse property; whereas emo-

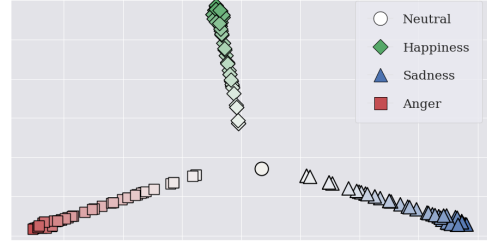


Figure 2: t-SNE plot of joint emotion embeddings generated from randomly sampled emotion strengths with various emotions. The brighter color indicates that the emotion strength becomes stronger, and vice versa.

tion strength can adjust the emotion embedding to finely represent the emotional properties.

3.4. TTS model

As an acoustic model of our TTS system, we adopt a duration-informed Tacotron 2 [13–15], which has the capacity to accurately align the phoneme sequence with the acoustic features. In this method, the linguistic features extracted from input text are fed to the Tacotron 2 encoder to extract high-level linguistic embeddings. Then, the external duration predictor estimates emotion-dependent phoneme duration by receiving the concatenated vectors of linguistic features and joint emotion embeddings in Eq. (1) as inputs. For given phoneme durations, the high-level linguistic embedding is up-sampled through a repetition process to match its temporal resolution to that of the acoustic features. Finally, the Tacotron 2 decoder [1] autoregressively generates the output acoustic features by receiving concatenated vectors of up-sampled linguistic embedding and joint emotion embedding.

After generating acoustic features, the multi-band harmonic-plus-noise PWG (MB-HN-PWG) vocoder [19] synthesizes the output speech. In this vocoder, two separated harmonic and noise WaveNets jointly capture the harmonic and noise components of speech within the PWG framework to generate high-quality speech waveforms. As a result, it can stably generate qualified speech signals with desired emotion.

3.5. Training criteria

Our models consist of two separated parts: the joint emotion predictor and the acoustic model. The loss function of each part is described as follows:

$$\mathcal{L}_a = \mathcal{L}_2(\hat{\mathbf{y}}, \mathbf{y}) + \mathcal{L}_2(\hat{\mathbf{d}}, \mathbf{d}) \quad (2)$$

$$\mathcal{L}_e = \mathcal{L}_2(\hat{e}_{str}, e_{str}) + \lambda_{cls} \cdot CE(\hat{\mathbf{e}}_{cls}, \mathbf{e}_{cls}), \quad (3)$$

where \mathcal{L}_a^1 and \mathcal{L}_e^2 denote the loss functions of the acoustic model and emotion predictor, respectively; \mathbf{y} , \mathbf{d} , $\hat{\mathbf{y}}$, and $\hat{\mathbf{d}}$ denote acoustic features, phoneme duration, and their generated counterparts, respectively; $\mathcal{L}_2(\cdot, \cdot)$ and $CE(\cdot, \cdot)$ denote L-2 loss and cross entropy loss, respectively; λ_{cls} denotes the hyperparameter balancing the two loss terms, which was chosen to be 0.01 based on our preliminary experiments.

4. Experiments

4.1. Database and features

For the experiments, a phonetically balanced emotional TTS corpus recorded by a female Korean professional speaker was used. The speech corpus consisted of four different emotions: neutral,

¹During a training stage, the target emotion class and strength are used in the TTS model instead of predicted values to avoid the effect of the prediction error of emotion predictor.

²For the computational efficiency, we froze the parameters of the GPT-3 during the TTS’s model training.

happiness, sadness, and anger. Specifically, 2,000, 100, and 30 utterances were used as each emotion’s training, validation, and test sets, respectively. In total, 8,000 utterances (8.72 h), 400 utterances (0.43 h), and 120 utterances (0.16 h) were used for the training, validation, and test sets, respectively. The speech signals were sampled at 24-kHz with 16-bit quantization.

As the input of the TTS acoustic model, we used 354-dimensional linguistic feature vectors. Specifically, 330-dimensional vectors including phoneme, accent, break, and punctuation were used as categorical linguistic features; whereas 24-dimensional vectors including positional information of each phoneme were used as numerical features.

As the output vectors of the TTS acoustic model, we used 79-dimensional feature vectors extracted by an improved time-frequency trajectory excitation vocoder [31], which consisted of 40-dimensional line spectral frequencies, F0, energy, binary voicing flag, a 32-dimensional slowly evolving waveform, and a 4-dimensional rapidly evolving waveform. The frame length and the shifting size were set to 20 ms and 5 ms, respectively.

As emotion features for training rankSVM, the 384-dim feature vectors were extracted from the speech signal by using the openSMILE toolkit [32], which consisted of a 12-dimensional MFCC, harmonic-to-noise ratio, root mean square frame energy, zero-crossing rate, and F0 with their statistical parameters [33].

4.2. Model configuration

4.2.1. Emotion predictor and encoder

As a baseline emotion predictor, we adopted a system proposed in Cui et al. [20], which predicted emotion class by using a 2-dimensional CNN model. In particular, the baseline emotion predictor consisted of three convolution blocks composed of two 2-dimensional convolution layers with a 3×3 kernel size followed by a 2-dimensional max pooling operation layer. In each convolution block, the input/output dimensions were set to 1/64, 64/128, and 128/256, respectively. Finally, two FC layers with output dimensions of 128 and four were used as the last layers for predicting the emotion class. After predicting the emotion class, emotion embedding was composed by passing it to the LUT layer.

For the configuration of the proposed joint emotion predictor, we adopted the HyperClova [28], which is a Korean variant of GPT-3 trained on a Korean-centric text corpus, as a feature extraction network for predicting emotional attributes. In detail, the model consisted of 12 transformer decoder layers [34] with 768 hidden units and 16 attention heads and was trained by using a large-scale text corpus consisting of 560 billion tokens. The resulting network parameter was in a total of 137 million³. For the two FC layers, one hidden layer with 256 hidden units followed by ReLU activation was used. Each FC layer had softmax and linear output layers with the dimensions of 4 and 1, respectively.

To configure the joint emotion encoder, we used the LUT to convert the predicted emotion class to one of four feature vectors with a 32-dimensional size. The input and output dimensions of projection matrix were set to 32.

4.2.2. TTS model

For the acoustic model, the duration predictor consisted of three FC layers and a uni-directional long short-term memory (LSTM) layer with 256 units. Each FC layer had 1024, 1024, and 512 units, respectively. The Tacotron 2 encoder consisted of three convolution layers with 10×1 kernels and 512 channels, bidirectional LSTM with 512 memory blocks, and FC layers with 512 units. The Tacotron 2 decoder was composed of PreNet, PostNet, and the main LSTM block. The PreNet consisted of two FC layers with 256 units. The main LSTM block consisted of the two unidirectional LSTM layers with 1,024 memory blocks, followed by two FC layers with 79 units. Finally, the PostNet consisted of five convolution

³This model corresponded to the smallest type in the original HyperClova paper. Noted that the model size was relatively smaller compared to the largest HyperClova with 82 billion parameters [28]. In our preliminary experiments, however, we confirmed that 137 million parameters were enough to effectively train the emotion predictor while maintaining a low computational cost.

True label \ Predicted label	Neutral	Happiness	Sadness	Anger
Neutral	0.43	0.07	0.20	0.30
Happiness	0.03	0.40	0.23	0.33
Sadness	0.00	0.00	0.47	0.53
Anger	0.13	0.00	0.30	0.57

True label \ Predicted label	Neutral	Happiness	Sadness	Anger
Neutral	0.63	0.27	0.10	0.00
Happiness	0.10	0.77	0.07	0.07
Sadness	0.03	0.07	0.67	0.23
Anger	0.03	0.00	0.37	0.60

Figure 3: Classification accuracy of emotion predictors: (a) conventional CNN-based and (b) proposed LM-based methods

layers with 5×1 kernels and 512 channels, followed by a residual connection to its input.

For the configuration of the MB-HN-PWG vocoder, the harmonic WaveNet consisted of 20 dilated residual blocks with two exponentially increasing dilation cycles, whereas the noise WaveNet consisted of 10 residual blocks with one exponentially increasing dilation cycle. The number of residual and skip channels was set to 64, and the convolution filter size was five. More setup details for the neural vocoder can be found in our previous work [19].

4.3. Evaluations

4.3.1. Effectiveness of LM for emotion prediction

To verify the advantages of the LM-based emotion prediction, we compared the classification error of the emotion class predicted by conventional CNN-based and LM-based emotion predictors. As shown in Figure 3, our proposed emotion predictor achieved significantly higher classification accuracy than the conventional CNN-based method. For instance, the proposed LM-based emotion predictor scored 76.7% on classification accuracy when predicting happiness emotion, which is 36.7% higher than that of the conventional one. Overall, the classification accuracy of the proposed emotion predictor averaged by all emotions was 66.7%, which was 20.0% higher than the conventional one’s 46.7% prediction accuracy.

4.3.2. Subjective listening tests

For the subjective evaluation, all TTS systems had the same structures as the acoustic model and the neural vocoder, as described in Section 4.2.2, except for the method of deciding emotional attributes. First, we evaluated the system without (S1) or with (S2) the proposed joint emotion encoder by using manually annotated emotional attributes as described in Section 3.1. We then evaluated the accuracy of predicting emotional attributes by using conventional CNN-based (S3) and proposed LM-based (S4) emotion predictors. By following the method of Cui et al. [20] described in Section 4.2.1, the CNN-based emotion predictor estimated emotion class only. Noted that when the emotion encoder was not used, a simple LUT layer was applied to the emotion class to compute emotion embedding.

We conducted two MOS listening tests to evaluate the (1) naturalness and (2) emotional expressiveness of synthetic speech⁴. During the tests, a total of 18 native Korean listeners were asked to rate the score of speech samples using the following 5-point responses: 1 = Bad, 2 = Poor, 3 = Fair, 4 = Good, and 5 = Excellent. We randomly selected 15 sentences per emotion from the test set and then synthesized speech signals using the different models; thus, a total of 1,080 scores were collected to evaluate each system.

For the naturalness MOS test, the listeners were asked to rate how close the perceptual quality of the synthetic speech was to that of the recordings. The trends of the test results in Table 1 can be analyzed as follows: (1) Utilizing the emotion strength by using

⁴Generated audio samples are available at the following URL: <https://christophyoon.github.io/lmemotiontts/>

Table 1: The naturalness MOS test results with a 95% confidence interval. Note that in the test, the listeners were asked to evaluate the quality of each audio sample.

System Index	Type of Emotion predictor	Use of joint Emotion encoder	Emotion Class	Emotion Strength	Neutral	Happiness	Sadness	Anger	Average
S1	–	–	Manual	–	4.12±0.22	4.20±0.24	3.85±0.30	3.46±0.23	3.90±0.17
S2	–	Yes	Manual	Manual	4.17±0.19	4.23±0.21	3.83±0.28	3.59±0.21	3.96±0.15
S3	CNN [20]	–	Predicted	–	3.96±0.15	3.75±0.24	3.91±0.23	3.56±0.15	3.79±0.09
S4 (Proposed)	LM [28]	Yes	Predicted	Predicted	4.12±0.22	4.19±0.24	3.84±0.32	3.52±0.23	3.92±0.15
Recordings	–	–	–	–	4.64±0.17	4.29±0.30	4.63±0.20	4.69±0.20	4.56±0.09

Table 2: Emotional expressiveness MOS test results with a 95% confidence interval. Note that in the test, the listeners were asked to evaluate how the emotions of synthetic speech matched the content of the text.

System Index	Type of Emotion predictor	Use of joint Emotion encoder	Emotion Class	Emotion Strength	Neutral	Happiness	Sadness	Anger	Average
S1	–	–	Manual	–	4.00±0.24	4.07±0.23	3.87±0.32	3.61±0.33	3.89±0.10
S2	–	Yes	Manual	Manual	4.09±0.23	4.14±0.23	3.86±0.25	3.77±0.29	3.96±0.09
S3	CNN [20]	–	Predicted	–	3.64±0.15	2.52±0.21	3.34±0.16	3.53±0.24	3.26±0.25
S4 (Proposed)	LM [28]	Yes	Predicted	Predicted	4.22±0.25	4.01±0.28	3.84±0.34	3.67±0.30	3.94±0.12
Recordings	–	–	–	–	4.83±0.13	4.46±0.32	4.69±0.21	4.65±0.20	4.66±0.08

Table 3: Subjective preference test results (%) between two different types of generation cases. The systems that achieved significantly better preference at the $p < 0.01$ level are in boldface.

Index	S2	S4 _{single}	S4 _{multi}	Neutral	p -value
Test 1	40.6	–	40.0	19.4	0.47
Test 2	–	24.4	55.0	20.6	< 10⁻³

the proposed joint emotion encoder in addition to the emotion class improved the perceptual quality of the emotional TTS system (S1 vs. S2). (2) When the emotional attributes were predicted from text, the TTS system with the proposed LM-based emotion predictor provided a significantly higher quality of synthetic speech than the system with conventional CNN-based one (S3 vs. S4). (3) When the emotional attributes were predicted from the text, there was a quality degradation due to the prediction error of the emotion predictor (S1 vs. S3 and S2 vs. S4). (4) However, the TTS system with the proposed LM-based emotion predictor showed a competitive quality of synthetic speech to the system with manually decided emotional attributes by achieving only a 0.04 lower MOS result on average (S2 vs. S4).

For the emotion expressiveness MOS test, listeners were asked to score how the emotion of synthetic speech matched the context of the input text. The trends of the test results in Table 2 can be analyzed as follows: (1) The usage of both emotion class and strength improved the emotional expressiveness compared to the system using only emotion class (S1 vs. S2). This also verified the advantages of the proposed emotion encoder for better emotion expressiveness. (2) When the emotional attributes were predicted through the emotion predictor, there was a quality degradation in terms of emotion expressiveness (S1 vs. S3 and S2 vs. S4). This was mainly caused by misclassification in emotion class, which often generated utterances that did not match the emotional context. (3) Nonetheless, the proposed method still outperformed the conventional method significantly (S3 vs. S4). In particular, there was a large gap in the expressiveness score of happiness, where the classification error gap between conventional and proposed methods was the largest as explained in Section 4.3.1. This indicated that a more accurate prediction reduced the peculiarity that occurred in the emotion misclassification. (4) Finally, the gap between the LM-based emotion predictor and the manually decided system was only 0.02 on average (S2 vs. S4), indicating that the predicted emotional attributes were reliable.

4.3.3. Paragraph-level emotional speech synthesis

The merits of the proposed emotion prediction method could also be found in paragraph-level speech synthesis because the LM had the capacity to capture the emotional context among the multiple sentences. To verify this, we conducted two A-B preference tests

by using various audiobook scripts. In these tests, 18 native Korean listeners were asked to select a more preferable one between two generated speech samples. We randomly selected 10 paragraphs consisting of 6–7 sentences, and then synthesized corresponding speech corpora using three different synthesis methods as follows:

- S2: The emotional TTS system that *manually* annotated the emotional attributes. In particular, the class and strength were decided by considering the context of the input.
- S4_{single}: The emotional TTS system with the proposed emotion predictor that received only a single utterance as an input at a time. Note that this model did not consider emotional context among consecutive sentences.
- S4_{multi}: The emotional TTS system with the proposed emotion predictor that received *multiple utterances* at once. In this model, the context of consecutive sentences was considered for predicting emotional attributes.

Table 3 shows the results of the preference tests, whose trend can be summarized as follows: (1) The TTS system with proposed emotion modeling methods provided statistically indistinguishable synthetic speech compared to the TTS system using manually annotated emotional attributes (Test 1; S2 vs. S4_{multi}). This indicated that the proposed emotion prediction method had the capability of replacing the human-annotating process. (2) The listeners preferred using the multiple sentences to using only a single sentence for the input of the emotion predictor (Test 2; S4_{single} vs. S4_{multi}). This verified that the usage of LM was effective in generating paragraph-level content requiring consideration of the input sentences’ context.

5. Conclusion

We proposed LM-based effective emotion prediction methods for an emotional TTS system. As our method employed the LM to predict the emotion class and strength from the text input, the proposed system could generate the corresponding emotional voice without relying on auxiliary inputs. The experimental results verified that the proposed method provided a superior quality over the conventional methods and a comparable quality to a system requiring human labeling. Future works include utilizing more complex emotional representation method such as global style token [35] or variational auto-encoder [36] that can further improve the expressiveness of the synthetic speech.

6. Acknowledgements

This work was supported by Clova Voice, NAVER Corp., Seongnam, Korea.

7. References

- [1] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan, R. A. Saurous, Y. Agiomvrgiannakis, and Y. Wu, "Natural TTS synthesis by conditioning WaveNet on Mel spectrogram predictions," in *Proc. ICASSP*, 2018, pp. 4779–4783.
- [2] S.-H. Lee, H.-W. Yoon, H.-R. Noh, J.-H. Kim, and S.-W. Lee, "Multi-spectrograms: High-diversity and high-fidelity spectrogram generation with adversarial style combination for speech synthesis," in *Proc. AAAI*, 2020.
- [3] Y. Ren, C. Hu, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, "FastSpeech 2: Fast and high-quality end-to-end text-to-speech," in *Proc. ICLR*, 2021.
- [4] O. Kwon, I. Jang, C. Ahn, and H.-G. Kang, "An effective style token weight control technique for end-to-end emotional speech synthesis," *IEEE Signal Process. Letters*, vol. 26, no. 9, pp. 1383–1387, 2019.
- [5] 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," in *Proc. ICASSP*, 2021, pp. 5734–5738.
- [6] C. Lu, X. Wen, R. Liu, and X. Chen, "Multi-speaker emotional speech synthesis with fine-grained prosody modeling," in *Proc. ICASSP*, 2021, pp. 5729–5733.
- [7] Y. Wang, D. Stanton, Y. Zhang, R.-S. Ryan, E. Battenberg, J. Shor, Y. Xiao, Y. Jia, F. Ren, and R. A. Saurous, "Style Tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis," in *Proc. ICML*, 2018, pp. 5180–5189.
- [8] P. Wu, Z. Ling, L. Liu, Y. Jiang, H. Wu, and L. Dai, "End-to-end emotional speech synthesis using style tokens and semi-supervised training," in *Proc. APSIPA*, 2019, pp. 623–627.
- [9] G. E. Henter, J. Lorenzo-Trueba, X. Wang, and J. Yamagishi, "Deep encoder-decoder models for unsupervised learning of controllable speech synthesis," *arXiv preprint arXiv:1807.11470*, 2018.
- [10] Y.-J. Zhang, S. Pan, L. He, and Z.-H. Ling, "Learning latent representations for style control and transfer in end-to-end speech synthesis," in *Proc. ICASSP*, 2019, pp. 6945–6949.
- [11] J. Cambre, J. Colnago, J. Maddock, J. Tsai, and J. Kaye, "Choice of voices: A large-scale evaluation of text-to-speech voice quality for long-form content," in *Proc. CHI*, 2020, pp. 1–13.
- [12] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," in *Proc. NeurIPS*, vol. 33, pp. 1877–1901, 2020.
- [13] T. Okamoto, T. Toda, Y. Shiga, and H. Kawai, "Tacotron-based acoustic model using phoneme alignment for practical neural text-to-speech systems," in *Proc. ASRU*, 2019, pp. 214–221.
- [14] E. Song, M.-J. Hwang, R. Yamamoto, J.-S. Kim, O. Kwon, and J.-M. Kim, "Neural text-to-speech with a modeling-by-generation excitation vocoder," in *Proc. Interspeech*, 2020, pp. 3570–3574.
- [15] M.-J. Hwang, R. Yamamoto, E. Song, and J.-M. Kim, "TTS-by-TTS: TTS-driven data augmentation for fast and high-quality speech synthesis," in *Proc. ICASSP*, 2021, pp. 6598–6602.
- [16] R. Yamamoto, E. Song, and J. Kim, "Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram," in *Proc. ICASSP*, 2020, pp. 6199–6203.
- [17] E. Song, R. Yamamoto, M.-J. Hwang, J.-S. Kim, O. Kwon, and J.-M. Kim, "Improved Parallel WaveGAN vocoder with perceptually weighted spectrogram loss," in *Proc. SLT*, 2021, pp. 470–476.
- [18] R. Yamamoto, E. Song, M.-J. Hwang, and J.-M. Kim, "Parallel waveform synthesis based on generative adversarial networks with voicing-aware conditional discriminators," in *Proc. ICASSP*, 2021, pp. 6039–6043.
- [19] M.-J. Hwang, R. Yamamoto, E. Song, and J.-M. Kim, "High-fidelity Parallel WaveGAN with multi-band harmonic-plus-noise model," in *Proc. Interspeech*, 2021, pp. 2227–2231.
- [20] C. Cui, Y. Ren, J. Liu, F. Chen, R. Huang, M. Lei, and Z. Zhao, "EMOVIE: A mandarin emotion speech dataset with a simple emotional text-to-speech model," in *Proc. Interspeech*, 2021.
- [21] Y. Lei, S. Yang, and L. Xie, "Fine-grained emotion strength transfer, control and prediction for emotional speech synthesis," in *Proc. SLT*, 2021, pp. 423–430.
- [22] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio *et al.*, "Tacotron: Towards end-to-end speech synthesis," in *Proc. Interspeech*, 2017, pp. 4006–4010.
- [23] Y. Lei, S. Yang, X. Wang, and L. Xie, "MsEmoTTS: Multi-scale emotion transfer, prediction, and control for emotional speech synthesis," *IEEE/ACM Trans. Audio, Speech, and Lang. Process.*, vol. 30, pp. 853–864, 2022.
- [24] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. NAACL*, 2019.
- [25] X. Zhu, S. Yang, G. Yang, and L. Xie, "Controlling emotion strength with relative attribute for end-to-end speech synthesis," in *Proc. ASRU*, 2019, pp. 192–199.
- [26] D. Parikh and K. Grauman, "Relative attributes," in *Proc. ICCV*, 2011, pp. 503–510.
- [27] T. Joachims, "Optimizing search engines using clickthrough data," in *Proc. ACM SIGKDD*, 2002, pp. 133–142.
- [28] B. Kim, H. Kim, S.-W. Lee, G. Lee, D. Kwak, D. H. Jeon, S. Park, S. Kim, S. Kim, D. Seo *et al.*, "What changes can large-scale language models bring? Intensive study on HyperCLOVA: Billions-scale korean generative pretrained transformers," in *Proc. EMNLP*, 2021.
- [29] H. Zheng, Z. Yang, W. Liu, J. Liang, and Y. Li, "Improving deep neural networks using softplus units," in *Proc. IJCNN*, 2015, pp. 1–4.
- [30] L. van der Maaten and G. Hinton, "Visualizing data using t-SNE," *Journal of Machine Learning Research*, vol. 9, pp. 2579–2605, 2008.
- [31] E. Song, F. K. Soong, and H.-G. Kang, "Effective spectral and excitation modeling techniques for LSTM-RNN-based speech synthesis systems," *IEEE/ACM Trans. Audio, Speech, and Lang. Process.*, vol. 25, no. 11, pp. 2152–2161, 2017.
- [32] F. Eyben, F. Wenyinger, F. Gross, and B. Schuller, "Recent developments in openSMILE, the munich open-source multimedia feature extractor," in *Proc. ACM*, 2013, pp. 835–838.
- [33] B. Schuller, S. Steidl, and A. Batliner, "The INTERSPEECH 2009 emotion challenge," in *Proc. Interspeech*, 2009, pp. 312–315.
- [34] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. NeurIPS*, 2017, pp. 5998–6008.
- [35] D. Stanton, Y. Wang, and R. Skerry-Ryan, "Predicting expressive speaking style from text in end-to-end speech synthesis," in *Proc. SLT*, 2018, pp. 595–602.
- [36] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *Proc. ICLR*, 2014.