



A Study of Modeling Rising Intonation in Cantonese Neural Speech Synthesis

Qibing Bai¹, Tom Ko^{2,*}, Yu Zhang¹

¹Department of Computer Science and Engineering, Southern University of Science and Technology

²ByteDance AI Lab

Abstract

In human speech, the attitude of a speaker cannot be fully expressed only by the textual content. It has to come along with the intonation. Declarative questions are commonly used in daily Cantonese conversations, and they are usually uttered with rising intonation. Vanilla neural text-to-speech (TTS) systems are not capable of synthesizing rising intonation for these sentences due to the loss of semantic information. Though it has become more common to complement the systems with extra language models, their performance in modeling rising intonation is not well studied. In this paper, we propose to complement the Cantonese TTS model with a BERT-based statement/question classifier. We design different training strategies and compare their performance. We conduct our experiments on a Cantonese corpus named CanTTS. Empirical results show that the separate training approach obtains the best generalization performance and feasibility.

Index Terms: text-to-speech, prosody, intonation, declarative questions

1. Introduction

Significant developments of neural text-to-speech (TTS) synthesis have achieved realistic speech generation [1, 2, 3]. Neural TTS models are based on deep neural networks and trained with large corpora. Using an encoder-decoder architecture, neural TTS models map input characters or phonemes to acoustic features (e.g., mel-spectrograms) or directly to the waveform. The acoustic features can be converted to waveforms with vocoders [4, 5]. Though synthesized speech is close to human speech in naturalness, its prosody is not always suitable given the text.

For prosody modeling, traditional parametric synthesis explicitly represents prosody components with models such as ToBI [6] via text analysis. Neural TTS methods usually learn prosody implicitly, and the generated prosody often sounds flattened. In order to improve the generated prosody, several variational [7, 8] and non-variational [9, 10] methods have been proposed to learn latent prosodic representations. Some methods [11, 12] are proposed for low-level prosody control. In inference, though some prosody attributes like emotions can be captured and transferred through reference audios, other prosody attributes related to the context could be inappropriate. Several works exploit various syntactic and semantic features in neural TTS [13, 14, 15, 16] to generate prosody suitable for the context. A pre-trained BERT-like speech module [17] learns prosodic patterns from unlabeled speech data to improve synthesized prosody. [18] leverages ToBI labels into neural TTS to improve the prosody. Recent works [19, 20] attempt to sample from the learned prosodic distribution using contextual

information. However, sentence-level prosody related to the speaker's intention draws much less attention.

A speaker's attitude can not be fully expressed only by phoneme segments in speech. It has to come along with the intonation. For example, a speaker can change a statement into a declarative question¹ with rising intonation. A declarative question usually corresponds to a yes/no answer. Consider the following sentences²:

- 他去学校。 He goes to school. (A statement)
- 他去学校? He goes to school? (A declarative question)
- 他去不去学校? Does he go to school? (A normal question)

Rising intonation is not necessary for a normal question as listeners can understand that it is a question from spoken phonemes. Therefore, speakers tend to utter normal questions with intonation similar to statements, especially in languages like Mandarin and Cantonese, where question particles play the main role. Although declarative questions are ended with the question mark in their written form, if they are uttered with a normal intonation, the listeners may not perceive the speaker's intention correctly. Thus, rising intonation is necessary for intention clarification.

Basic neural TTS models that solely take phonemes as input are not capable of synthesizing rising intonation for declarative questions. It is because of two reasons. Firstly, questions constitute only a small portion of available data corpora, not to mention declarative questions. Thus, samples with rising intonation are insufficient for a model to learn implicitly. Secondly, the grapheme-to-phoneme (g2p) conversion causes a loss of semantic and syntactic information that identifies the sentence type. g2p helps generalize on rare/unseen characters and is necessary for languages with a large number of characters, e.g., Chinese and Korean. Besides, it is hard for TTS models to capture the sentence type without powerful language models. Thus the TTS model alone cannot map the input text to its associated intonation.

The judgment of declarative questions associated with rising intonation depends on the sentence semantics. However, various methods [7, 8, 9, 10] based on sampling or reference audios do not consider textual semantic information, nor does the pre-trained prosody encoder [17]. Recent attempts on injecting various linguistic features [13, 14, 15, 16] and context-based prosody sampling [19, 20] can compensate for the information loss, but they do not provide discriminative information about sentence types. Therefore, the performance on rising intonation featured by declarative questions is not well studied yet.

In this paper, we propose to complement the TTS model with a BERT-based statement/question classifier. We design

¹In Chinese literature, declarative questions are also referred to as unmarked questions.

²English translations may have different intonation.

* corresponding author

three training approaches: two are named “explicit” because they use sentence type labels in training, and the other requires no label and is called “implicit”. We utilize a BERT-based sentence classifier to provide the sentence type information to the TTS model. We conduct our experiments on a Cantonese corpus named CanTTS, which explicitly labels the sentence type and intonation. It contains three types of sentences: statements, normal questions, and declarative questions. The statements and normal questions are all uttered with non-rising intonation; the declarative questions are all uttered with rising intonation. Objective and subjective evaluations show that all models augmented with the sentence classifier outperform the baseline Tacotron2. A separately trained classifier achieves the best performance while enjoying the feasibility of leveraging extra text data.

2. Related Work

BERT [21] is a language model that learns to capture contextualized text information and generate general-purpose text representations. It is trained on large text corpora with masked language modeling. The generated text representations carry both syntactic and semantic information [22] that is useful in prosody generation. Though the [CLS] token is supposed to capture sentence-level information, BERT usually needs extra discriminative guidance when applied to classification tasks.

In [15], an additional attention module is used to extract context vector for the Tacotron [23] decoder. [15] also directly concatenates the [CLS] token with phoneme representations and feeds them to the decoder. [24] uses BERT to provide semantic features in comparison with syntactic features. They average representations from multiple BERT layers, which are used to select a sentence-level acoustic embedding. [16] uses BERT-derived representations to predict breaks as extra inputs to the encoder. They also test directly concatenating the representations to phoneme embeddings. [25] uses contextual representations of neighboring sentences for prosody improvement. They use a cross-utterance encoder to process BERT-derived representations and concatenate them with phoneme representations. PnG BERT [26] is a TTS encoder that takes characters and phonemes as its input. It is pre-trained on text data with phoneme transcriptions, therefore able to provide both pronunciation and semantic information.

All the above approaches are shown to improve the naturalness of synthesized speech. However, they do not provide adequate information about declarative questions, which features rising intonation in Cantonese.

3. Our Approach

We design a BERT-based statement/question classifier to provide the sentence type information associated with non-rising/rising intonation. The sentence classifier consists of a sentence encoder and a classification layer. All the systems in this work are based on Tacotron2 and the sentence encoder/classifier. The backbone of Tacotron2 is shown in the left part of Fig. 1. We additionally use the MAE loss and the guided attention loss [27] for faster convergence. We also adopt Multi-band MelGAN [5] as the vocoder.

We examine an implicit approach and two explicit approaches for learning the intonation from the data. The implicit approach does not require sentence labels in training, while the

explicit ones do.

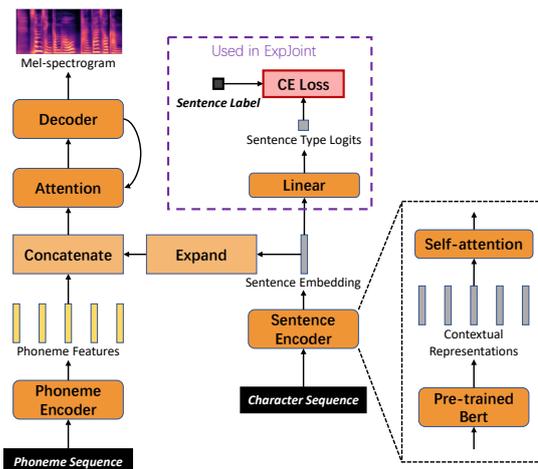


Figure 1: Structure of ImpJoint and ExpJoint. The modules in the black dashed box make up the sentence encoder. The modules in the purple dashed box are used in ExpJoint training.

3.1. Sentence encoder/classifier

The sentence encoder is shown in the black dashed box in Fig. 1. It consists of a pre-trained BERT model and a self-attention module. The pre-trained BERT model, which is powerful in extracting semantic information, outputs token-level representations. The self-attention module extracts a sentence-level embedding from a linear transformation of the token-level representations [28]. The formulations of the self-attention module are as follows. We first compute a scalar score e_t for each input representation \mathbf{h}_t , normalize the scores to obtain weights α_t , and finally sum up the representations:

$$e_t = \mathbf{v}^T \tanh(\mathbf{W}\mathbf{h}_t + \mathbf{b}) \quad (1)$$

$$\alpha_t = \frac{\exp(e_t)}{\sum_{i=1}^T \exp(e_i)} \quad (2)$$

$$\mathbf{s} = \sum_{t=1}^T \alpha_t \mathbf{h}_t, \quad (3)$$

where \mathbf{s} is the sentence embedding, \mathbf{h}_t is the t -th output of the self-attention layers, and transformation parameters \mathbf{W} , \mathbf{b} as well as the query vector \mathbf{v} are learnable.

For sentence-level tasks like sentence classification, as suggested in [21], it is common to use the classifier embedding [CLS] alone with a classification layer. In our work, we empirically find that a linear combination of the [CLS] and all other contextual representations can lead to better overall results. Thus, we apply this setup in all the systems. In our work, the [CLS] token and the contextual representations are derived from a pre-trained Cantonese BERT.

The sentence encoder can be followed by an additional classification layer for statement/question classification, as shown in the purple dashed box in Fig. 1. There are three sentence categories: statement, normal question, and declarative question. If we remove end punctuation marks, the problem becomes binary classification (normal question vs. statement/declarative question). We still adopt the three classes since it is trivial for BERT to consider end punctuation marks.

3.2. Implicit approach – ImpJoint

In this approach, it is assumed that there are certain utterances with rising intonation, but no sentence labels are available. This simulates the real situation where labels on prosody are difficult to acquire. We intend to investigate whether the TTS model can implicitly learn contextual information of rising intonation.

The model architecture is shown in Fig. 1. The input text is processed by the sentence encoder to generate a sentence embedding, which is then expanded and concatenated to encoded phoneme features. The Tacotron2 and the sentence encoder are jointly trained.

3.3. Explicit approach

In this approach, it is assumed that we can make use of the sentence and intonation labels in training. We design and examine two explicit approaches: a joint training approach and a separate training approach.

3.3.1. Joint training – ExpJoint

We utilize the sentence labels in the joint training framework by incorporating the sentence classifier to perform an extra classification task. The sentence embedding is fed to both the TTS model and the classification layer (shown in the purple dashed box in Fig. 1). The cross-entropy (CE) loss is used.

We expect that the sentence embeddings can carry more discriminative information about the sentence type. This helps the TTS model to learn the relationship between the sentence type and the intonation in the spectrogram.

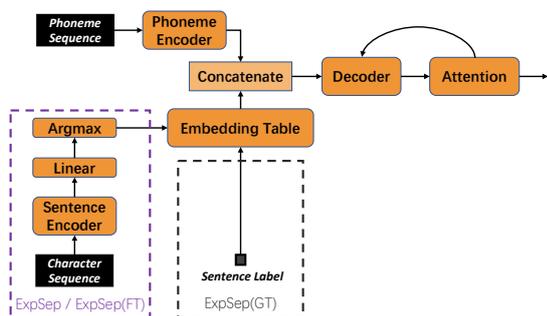


Figure 2: Structure of ExpSep, ExpSep(FT) and ExpSep(GT).

3.3.2. Separate training – ExpSep

As the sentence and intonation labels are available, it is feasible to train the TTS model and the sentence classifier separately (i.e., Fig. 2). The TTS model and the sentence classifier are connected with an embedding table where each entry corresponds to a sentence type.

A major advantage of the separate training approach is that the sentence classifier can be trained with extra text data without paired speech. It is much easier to collect text than to prepare paired speech data. We basically freeze the BERT parameters for comparison with joint training approaches. We also test fine-tuning the BERT parameters while training the classifier and name this model ExpSep(FT).

If we allow the user to hint at the sentence type, the TTS model separately trained in this approach can actually be used alone. We name this TTS model ExpSep(GT) as we assume that it is always fed with ground-truth sentence labels.

4. Experimental Setup

4.1. Dataset

We conduct our experiments on CanTTS³, a single-speaker Cantonese corpus consisting of 12,010 utterances. The corpus contains $\sim 10,000$ statements and ~ 2000 questions. The question set can be further divided into $\sim 1,000$ normal questions and $\sim 1,000$ declarative questions. There is a total of 20 hours of speech recorded from a female speaker aged 20. All the statements are ended up with either a comma or a full stop, and all the questions are ended up with a question mark in their transcription. In our work, all the punctuation marks are processed as part of the input text. All the declarative questions are uttered with rising intonation. All the statements and normal questions are uttered with a non-rising intonation.

Some of the utterances contain English words, so we exclude these code-mixed samples. We further divide the corpus into a training set and a test set. The proportions of sentence types are the same across the two subsets. The test set contains 448 statements, 50 normal questions, and 50 declarative questions.

4.2. Model setting

We use the pre-trained Cantonese BERT-Base from Huggingface [29] and conduct our experiments on ESPnet [30]. As explained in Section 3.1, the sentence classifier has 3 output types, and the sentence embedding table in ExpSep, ExpSep(FT) and ExpSep(GT) has 3 entries. The dimension of the embedding table is 512.

All TTS models are trained with the Adam optimizer [31], a batch size of 135, and a learning rate of 10^{-4} . In training ExpJoint, the weight of sentence classification loss is 0.1. The class weights of statements, normal questions, and declarative questions within the classification loss are 1, 10, and 20. In training the sentence classifier in ExpSep and ExpSep(FT), the learning rate is 10^{-5} and the batch size is 512. We expand the text training set by removing the end punctuation marks of all sentences; the declarative questions without end punctuation are taken as statements.

5. Results and Discussion

All designed models, with the help of the sentence encoder/classifier, can synthesize rising intonation for declarative questions. To further compare the model performance, we use both objective and subjective evaluations. For ease of presentation, statements, normal questions, and declarative questions are abbreviated Sta, Que, and DecQue, respectively. On the testing set, ExpSep(FT) predicts labels for all the sentences. Therefore, we present performance of the two systems with the name ExpSep(GT/FT).

A declarative question example that contains rising intonation is shown in Fig. 3. In the ground-truth mel-spectrogram, the pitch contour rises obviously at the end. Tacotron2 fails to render such a pattern, while other models perform much better. We also find that ExpSep can render rising intonation for long sentences. This indicates its ability in generalization.

5.1. Objective evaluation

Since intonation is mainly reflected by pitch contours, we apply the F0 Frame Error (FFE) [32] for objective evaluation. We

³<https://github.com/parami-ai/CanTTS>

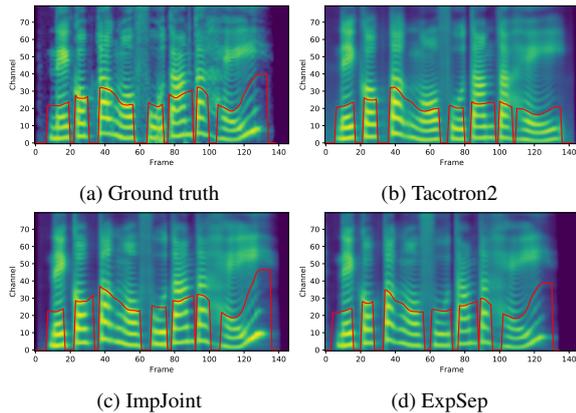


Figure 3: Mel-spectrograms (with F0 contours) of a declarative question. (你觉得我负担得起?, You think I can afford it?)

align each synthesized mel-spectrogram and the ground truth using dynamic time wrapping (DTW) with the Mel-cestral distortion (MCD) cost metric [33]. Then we synthesize waveforms, extract pitch contours, and compute the FFE. We calculate averaged FFE for each model on the test set. The result is shown in Table 1.

Table 1: FFE of five systems on different subsets.

System	Sta	Que	DecQue	All
Tacotron2	13.53%	25.52%	32.76%	16.38%
ImpJoint	13.62%	20.43%	20.77%	14.89%
ExpJoint	12.7%	18.86%	18.107%	13.76%
ExpSep	14.03%	20.08%	19.71%	15.10%
ExpSep(GT/FT)	14.03%	15.93%	14.13%	14.21%

Among the designed systems, ExpSep(GT/FT) generates declarative questions with the most similar pitch contours to the ground truth, and it performs well on normal questions. All proposed systems perform significantly better than Tacotron2 on declarative questions. We attribute this to the help of the sentence encoder/classifier.

5.2. Subjective evaluation

We first employ the MOS test for subjective evaluation.⁴ Each model is evaluated on 30 statements, 30 normal questions, and 30 declarative questions that are randomly selected from the test set. The raters know the sentence types of all test samples, and they need to score based on both the quality and perception of the corresponding sentence type. For example, raters are expected to lower the score when a synthesized declarative question sounds like a statement. Each audio is rated by 10 native Cantonese speakers.

As Table 2 shows, all designed systems with the sentence encoder/classifier have better overall performance than the baseline Tacotron2. It is because the Tacotron2 does not have rich contextual information for deciding whether to synthesize rising intonation. The ImpJoint model, trained without sentence labels, also synthesizes most declarative questions with rising intonation. This suggests that the mel-spectrogram

⁴Audio samples: <https://p1ping.github.io/RI-TTS>

Table 2: The MOS results of five systems on three test sets.

System	Sta	Que	DecQue
Tacotron2	4.22±0.11	4.27±0.12	2.88±0.18
ImpJoint	4.24±0.11	4.21±0.11	4.20±0.12
ExpJoint	4.26±0.11	4.23±0.12	4.21±0.13
ExpSep	4.21±0.11	4.32±0.12	4.33±0.12
ExpSep(GT/FT)	4.26±0.10	4.41±0.10	4.46±0.09

loss also guides the model to learn certain contextual information associated with rising intonation.

We also prepare another test set to evaluate the perception accuracy. This set contains 28 statements and 28 declarative questions. The sentence types are unknown to the raters. They are asked to judge whether each audio is a statement or a declarative question. We then compute the percentage of correctly perceived sentences. The results are summarized in Table 3, which are consistent with the MOS results.

Table 3: Perception accuracy.

System	Sta	DecQue	All
Tacotron2	96.43%	21.83%	59.13%
ImpJoint	94.44%	85.32%	89.88%
ExpJoint	95.63%	87.70%	91.67%
ExpSep	95.63%	90.87%	93.25%
ExpSep(GT/FT)	96.03%	98.41%	97.22%

Even though sentence labels enrich discriminative information of sentence types, ExpJoint only wins a little against ImpJoint. This suggests that the TTS model does not make full use of the discriminative information in the joint training approach.

ExpSep has better subjective evaluation results than ExpJoint. This is because the embedding table stably catches different intonation patterns, and the TTS model can synthesize rising intonation once given the label. Note that ExpJoint outperforms ExpSep objectively in FFE, which measures the entire pitch contour. This suggests that the sentence embedding before the classification layer also helps with the overall intonation.

As expected, ExpSep(GT) performs the best as we assume that the users can hint at the sentence type. Fine-tuning BERT parameters proves to be effective as ExpSep(FT) predicts all testing sentences correctly. When samples with rising intonation are insufficient to train the classifier, ExpSep can be trained with extra text data, which shows the feasibility of ExpSep.

6. Conclusions

In this paper, we propose to complement the TTS model with a BERT-based statement/question classifier. We design different training strategies, conduct our experiments on the CanTTS Cantonese corpus and compare their performance. All evaluated models outperform the vanilla Tacotron2 in synthesizing appropriate rising intonation. Empirical results show that the separate training approach obtains the best performance, generalization ability, and feasibility. In the future work, we will explore transferring rising intonation to other speakers/languages and extend the methods to model more fine-grained prosody by considering contextual information other than the intention.

7. References

- [1] Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerry-Ryan, et al., “Natural tts synthesis by conditioning wavenet on mel spectrogram predictions,” in *ICASSP*. IEEE, 2018, pp. 4779–4783.
- [2] Yi Ren, Yangjun Ruan, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu, “FastSpeech: Fast, robust and controllable text to speech,” in *Advances in Neural Information Processing Systems*, 2019, pp. 3171–3180.
- [3] Jaehyeon Kim, Jungil Kong, and Juhee Son, “Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 5530–5540.
- [4] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu, “Wavenet: A generative model for raw audio,” in *9th ISCA Speech Synthesis Workshop*, pp. 125–125.
- [5] Geng Yang, Shan Yang, Kai Liu, Peng Fang, Wei Chen, and Lei Xie, “Multi-band melgan: Faster waveform generation for high-quality text-to-speech,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 492–498.
- [6] Kim EA Silverman, Mary E Beckman, John F Pitrelli, Mari Ostendorf, Colin W Wightman, Patti Price, Janet B Pierrehumbert, and Julia Hirschberg, “Tobi: A standard for labeling english prosody,” in *ICSLP*, 1992, vol. 2, pp. 867–870.
- [7] Wei-Ning Hsu, Yu Zhang, Ron J Weiss, Heiga Zen, Yonghui Wu, Yuxuan Wang, Yuan Cao, Ye Jia, Zhifeng Chen, Jonathan Shen, et al., “Hierarchical generative modeling for controllable speech synthesis,” in *International conference on learning representations*, 2019.
- [8] Ya-Jie Zhang, Shifeng Pan, Lei He, and Zhen-Hua Ling, “Learning latent representations for style control and transfer in end-to-end speech synthesis,” in *ICASSP*. IEEE, 2019, pp. 6945–6949.
- [9] RJ Skerry-Ryan, Eric Battenberg, Ying Xiao, Yuxuan Wang, Daisy Stanton, Joel Shor, Ron Weiss, Rob Clark, and Rif A Saurous, “Towards end-to-end prosody transfer for expressive speech synthesis with tacotron,” in *international conference on machine learning*. PMLR, 2018, pp. 4693–4702.
- [10] Yuxuan Wang, Daisy Stanton, Yu Zhang, RJ-Skerry Ryan, Eric Battenberg, Joel Shor, Ying Xiao, Ye Jia, Fei Ren, and Rif A Saurous, “Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 5180–5189.
- [11] Alexandra Vioni, Myrsini Christidou, Nikolaos Ellinas, Georgios Vamvoukakis, Panos Kakoulidis, Taehoon Kim, June Sig Sung, Hyoungmin Park, Aimilios Chalamandaris, and Pirros Tsiakoulis, “Prosodic clustering for phoneme-level prosody control in end-to-end speech synthesis,” in *ICASSP*. IEEE, 2021, pp. 5719–5723.
- [12] Cheng Gong, Longbiao Wang, Zhenhua Ling, Shaotong Guo, Ju Zhang, and Jianwu Dang, “Improving naturalness and controllability of sequence-to-sequence speech synthesis by learning local prosody representations,” in *ICASSP*. IEEE, 2021, pp. 5724–5728.
- [13] Haohan Guo, Frank K Soong, Lei He, and Lei Xie, “Exploiting syntactic features in a parsed tree to improve end-to-end tts,” *Proc. Interspeech 2019*, pp. 4460–4464, 2019.
- [14] Changhe Song, Jingbei Li, Yixuan Zhou, Zhiyong Wu, and Helen Meng, “Syntactic representation learning for neural network based tts with syntactic parse tree traversal,” in *ICASSP*. IEEE, 2021, pp. 6064–6068.
- [15] Tomoki Hayashi, Shinji Watanabe, Tomoki Toda, Kazuya Takeda, Shubham Toshniwal, and Karen Livescu, “Pre-trained text embeddings for enhanced text-to-speech synthesis,” in *INTER-SPEECH*, 2019, pp. 4430–4434.
- [16] Yujia Xiao, Lei He, Huaiping Ming, and Frank K Soong, “Improving prosody with linguistic and bert derived features in multi-speaker based mandarin chinese neural tts,” in *ICASSP*. IEEE, 2020, pp. 6704–6708.
- [17] Liping Chen, Yan Deng, Xi Wang, Frank K Soong, and Lei He, “Speech bert embedding for improving prosody in neural tts,” in *ICASSP*. IEEE, 2021, pp. 6563–6567.
- [18] Yuxiang Zou, Shichao Liu, Xiang Yin, Haopeng Lin, Chunfeng Wang, Haoyu Zhang, and Zejun Ma, “Fine-grained prosody modeling in neural speech synthesis using tobi representation,” *Proc. Interspeech 2021*, pp. 3146–3150, 2021.
- [19] Zack Hodari, Alexis Moinet, Sri Karlapati, Jaime Lorenzo-Trueba, Thomas Merritt, Arnaud Joly, Ammar Abbas, Penny Karanasou, and Thomas Drugman, “Camp: a two-stage approach to modelling prosody in context,” in *ICASSP*. IEEE, 2021, pp. 6578–6582.
- [20] Sri Karlapati, Ammar Abbas, Zack Hodari, Alexis Moinet, Arnaud Joly, Penny Karanasou, and Thomas Drugman, “Prosodic representation learning and contextual sampling for neural text-to-speech,” in *ICASSP*. IEEE, 2021, pp. 6573–6577.
- [21] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of NAACL-HLT*, 2019, pp. 4171–4186.
- [22] Ganesh Jawahar, Benoît Sagot, and Djamé Seddah, “What does bert learn about the structure of language?,” in *The 57th Annual Meeting of the Association for Computational Linguistics*, 2019.
- [23] Yuxuan Wang, RJ Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, et al., “Tacotron: Towards end-to-end speech synthesis,” *Proc. Interspeech 2017*, pp. 4006–4010, 2017.
- [24] Shubhi Tyagi, Marco Nicolis, Jonas Rohnke, Thomas Drugman, and Jaime Lorenzo-Trueba, “Dynamic prosody generation for speech synthesis using linguistics-driven acoustic embedding selection,” *Proc. Interspeech 2020*, pp. 4407–4411, 2020.
- [25] Guanghui Xu, Wei Song, Zhengchen Zhang, Chao Zhang, Xiaodong He, and Bowen Zhou, “Improving prosody modelling with cross-utterance bert embeddings for end-to-end speech synthesis,” in *ICASSP*. IEEE, 2021, pp. 6079–6083.
- [26] Ye Jia, Heiga Zen, Jonathan Shen, Yu Zhang, and Yonghui Wu, “Png bert: Augmented bert on phonemes and graphemes for neural tts,” *Proc. Interspeech 2021*, pp. 151–155, 2021.
- [27] Hideyuki Tachibana, Katsuya Uenoyama, and Shunsuke Aihara, “Efficiently trainable text-to-speech system based on deep convolutional networks with guided attention,” in *ICASSP*, 2018.
- [28] Yingke Zhu, Tom Ko, David Snyder, Brian Mak, and Daniel Povey, “Self-attentive speaker embeddings for text-independent speaker verification,” in *Interspeech*, 2018, pp. 3573–3577.
- [29] Thomas Wolf, Julien Chaumond, Lysandre Debut, Victor Sanh, Clement Delangue, Anthony Moi, Pierric Cistac, Morgan Funtowicz, Joe Davison, Sam Shleifer, et al., “Transformers: State-of-the-art natural language processing,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2020, pp. 38–45.
- [30] Tomoki Hayashi, Ryuichi Yamamoto, Katsuki Inoue, Takenori Yoshimura, Shinji Watanabe, Tomoki Toda, Kazuya Takeda, Yu Zhang, and Xu Tan, “Espnet-TTS: Unified, reproducible, and integratable open source end-to-end text-to-speech toolkit,” in *ICASSP*. IEEE, 2020, pp. 7654–7658.
- [31] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” in *International conference on learning representations*, 2015.
- [32] Wei Chu and Abeer Alwan, “Reducing f0 frame error of f0 tracking algorithms under noisy conditions with an unvoiced/voiced classification frontend,” in *ICASSP*. IEEE, 2009, pp. 3969–3972.
- [33] Robert Kubichek, “Mel-cepstral distance measure for objective speech quality assessment,” in *Proceedings of IEEE Pacific Rim Conference on Communications Computers and Signal Processing*. IEEE, 1993, vol. 1, pp. 125–128.