

# PROSODIC REPRESENTATION LEARNING AND CONTEXTUAL SAMPLING FOR NEURAL TEXT-TO-SPEECH

*Sri Karlapati, Ammar Abbas, Zack Hodari\*, Alexis Moinet, Arnaud Joly, Penny Karanasou, Thomas Drugman*

Amazon Research, Cambridge, United Kingdom

\*The Centre for Speech Technology Research, University of Edinburgh, United Kingdom

## ABSTRACT

In this paper, we introduce Kathaka, a model trained with a novel two-stage training process for neural speech synthesis with contextually appropriate prosody. In Stage I, we learn a prosodic distribution at the sentence level from mel-spectrograms available during training. In Stage II, we propose a novel method to sample from this learnt prosodic distribution using the contextual information available in text. To do this, we use BERT on text, and graph-attention networks on parse trees extracted from text. We show a statistically significant relative improvement of 13.2% in naturalness over a strong baseline when compared to recordings. We also conduct an ablation study on variations of our sampling technique, and show a statistically significant improvement over the baseline in each case.

*Index Terms*— TTS, prosody modelling, contextual prosody

## 1. INTRODUCTION

Neural text-to-speech (NTTS) techniques have significantly improved the naturalness of speech produced by TTS systems [1–6]. We refer to NTTS systems as a subset of TTS systems that use neural networks to predict mel-spectrograms from phonemes, followed by the use of neural vocoder to generate audio from mel-spectrograms.

In order to improve the prosody<sup>1</sup> of speech obtained from NTTS systems, there has been considerable work in learning prosodic latent representations from ground truth speech [7, 8]. These methods use the target mel-spectrograms as input to an encoder which learns latent prosodic representations. These representations are used by the decoder in addition to the input phonemes, to generate mel-spectrograms. The latent representations obtained by encoding a target mel-spectrogram at the sentence level will have information that is not directly available from the phonemes, and by the subtractive definition of prosody, we may claim that these representations capture prosodic information. Several variational [8–10] and non-variational [7, 11] methods have been proposed for

learning prosodic latent representations. While these methods improve the prosody of synthesised speech, they need an input mel-spectrogram which is not available while running inference on unseen text. This gives rise to the problem of sampling from the learnt prosodic space. Sampling at random from the prior [10] may result in the synthesised speech not having contextually appropriate prosody, as it has no relationship with the text being synthesised. In order to improve the contextual appropriateness of prosody in synthesised speech, there has been work on using textual features like contextual word embeddings and other grammatical information to directly condition NTTS systems [12–14]. These methods require the NTTS model to learn an implicit correlation between the given textual features and the prosody of the sentence. One work also poses this sampling problem as a selection problem and uses both syntactic distance and BERT embeddings to select a latent prosodic representation from the ones seen at training time [15].

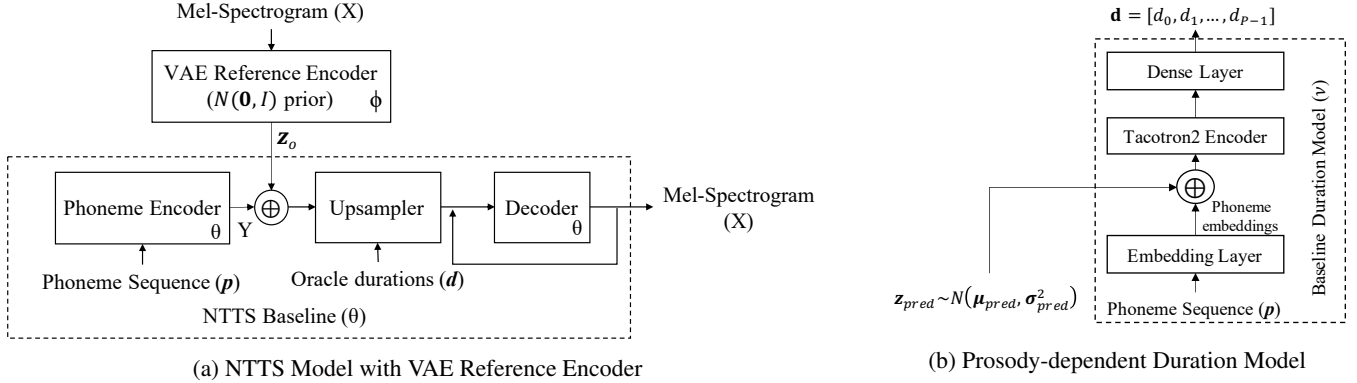
Bringing both the aforementioned ideas of using ground truth speech to learn prosodic latent representations and using textual information, we build Kathaka, a model trained using a two-stage training process to generate speech with contextually-appropriate prosody. In Stage I, we learn the distribution of sentence-level prosodic representations from ground truth speech using a VAE [8]. In Stage II, we learn to sample from the learnt distribution using text. In this work, we introduce the BERT+Graph sampler, a novel sampling mechanism which uses both contextual word-piece embeddings from BERT [16] and the syntactic structure of constituency parse trees through graph attention networks [17]. We then compare Kathaka against a strong baseline and show that it obtains a relative improvement of 13.2% in naturalness.

## 2. NTTS BASELINE WITH DURATION MODELLING

We use a modified version of DurIAN [3] as our baseline. The baseline learns to predict  $T$  mel-spectrogram frames  $X = [x_0, x_1, \dots, x_{T-1}]$ , given  $P$  phonemes  $p = [p_0, p_1, \dots, p_{P-1}]$ . The baseline implicitly models the prosody of synthesised speech without any additional context or reference prosody. We have a phoneme encoder, which

\*Work done as an intern at Amazon Research, Cambridge, UK.

<sup>1</sup>We use the subtractive definition of prosody from [7].



**Fig. 1:** Fig. (a) shows the architecture of the NTTS Baseline system ( $\theta$ ) within the bounding box and the added VAE reference encoder ( $\phi$ ) used in **Stage I** of Kathaka’s training procedure. Fig. (b) shows the architecture of the baseline duration model within the bounding box and the prosody-dependent duration model proposed in Section 3.3.

takes phonemes as input and returns phoneme encodings  $Y = [\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_{P-1}]$ . At training, we use forced alignment to extract durations,  $\mathbf{d} = [d_0, d_1, \dots, d_{P-1}]$ , where  $\sum_{d_i \in \mathbf{d}} d_i = T$ . The phoneme encodings are upsampled by replication as per the durations  $\mathbf{d}$  to obtain upsampled phoneme encodings  $Y^\uparrow = [\mathbf{y}_0^\uparrow, \mathbf{y}_1^\uparrow, \dots, \mathbf{y}_{T-1}^\uparrow]$ . The decoder auto-regressively learns to predict the mel-spectrogram  $X$  from  $Y^\uparrow$ . Unlike DurIAN, we do not use a post-net as this resulted in instabilities when training with a reduction factor of 1. Our baseline is shown within the bounding box in Fig. 1a, and the set  $\{\theta\}$  represents its parameters.

We train a duration model  $\hat{\mathbf{d}} = f_\nu(\mathbf{p})$ , to predict the duration  $\hat{d}_p$  in frames for each phoneme  $p$ , as shown by the blocks within the bounding box in Fig. 1b. We note that the distribution of phoneme durations contains multiple modes  $M$ . We normalize durations for a group of tokens  $G_m$ , for each mode  $m \in M$  separately using mean  $\mu_m$  and standard deviation  $\sigma_m$ . We use L2 loss to train the model as shown below, where  $\llbracket \cdot \rrbracket$  represent Iverson brackets.

$$\mu_m = \frac{1}{|G_m|} \sum_{p \in G_m} d_p, \sigma_m^2 = \frac{1}{|G_m|} \sum_{p \in G_m} (d_p - \mu_m)^2$$

$$L_\nu = \frac{1}{P} \sum_{p \in \mathcal{P}} \left\| \hat{d}_p - \prod_{m \in M} \left( \frac{d_p - \mu_m}{\sigma_m} \right)^{\llbracket p \in G_m \rrbracket} \right\|_2 \quad (1)$$

### 3. KATHAKA

Here, we introduce the two-stage approach to train Kathaka. In Section 3.1, we describe Stage I where the model learns a distribution over the prosodic latent space. In Section 3.2 we introduce our sampling mechanisms and finally discuss prosody dependent duration modelling in Section 3.3.

#### 3.1. Addition of a Variational Reference Encoder

We add a variational reference encoder  $q_\phi$ , with parameters  $\{\phi\}$ , to the NTTS architecture  $\{\theta\}$  as shown in Fig. 1a, to

learn prosodic representations from speech [8], through the subtractive definition of prosody [7]. This encoder  $q_\phi(z_o | X)$  takes a mel-spectrogram  $X$  as input and predicts the parameters of a Gaussian distribution  $\mathcal{N}(\mu_o, \sigma_o^2)$ , from which we sample a prosodic latent representation  $z_o$ . We assume a prior distribution  $p_\theta(z_o) = \mathcal{N}(\mathbf{0}, I)$  and train the model to maximize the evidence lower bound (ELBO) defined in Eq. 2, where  $\alpha$  is used as the annealing factor to avoid posterior collapse [18].

$$L_{\phi, \theta} = \mathbb{E}_{z_o \sim q_\phi} [\log(p_\theta(X | z_o, Y^\uparrow))] - \alpha D_{KL}(q_\phi(z_o | X) \parallel p_\theta(z_o)) \quad (2)$$

#### 3.2. Sampling Using Text

During inference, we need to generate the mel-spectrogram  $X$ , and therefore will not have access to it. We note that prosody is driven by the contextual information available in a sentence [19]. Therefore, we propose the usage of text or features derived from it to learn to predict a sentence-level prosodic latent representation  $z_{pred}$ . This is used in place of a mel-spectrogram based sentence-level prosodic latent representation  $z_o$ . We define samplers  $s_\chi(z_{pred} | W)$ , which take text or textual features  $W$  as input, and learn to predict a sentence-level prosodic latent representation  $z_{pred}$ . Since the reference encoder  $q_\phi$  predicts the parameters of a Gaussian distribution with a diagonal covariance matrix  $\mathcal{N}(\mu_o, \sigma_o^2)$ , we treat this as a distribution matching problem, and train samplers to predict parameters of a Gaussian distribution  $\mathcal{N}(\mu_{pred}, \sigma_{pred}^2)$ . We match these two distributions by minimising the KL Divergence:

$$D_{KL}(s_\chi(z_{pred} | W) \parallel q_\phi(z_o | X)) = \frac{1}{2} \left( \sum_{i=0}^{D-1} \left( \log(\sigma_{o_i}^2) - \log(\sigma_{pred_i}^2) + \frac{\sigma_{pred_i}^2}{\sigma_{o_i}^2} + \frac{(\mu_{o_i} - \mu_{pred_i})^2}{\sigma_{o_i}^2} \right) - D \right), \quad (3)$$

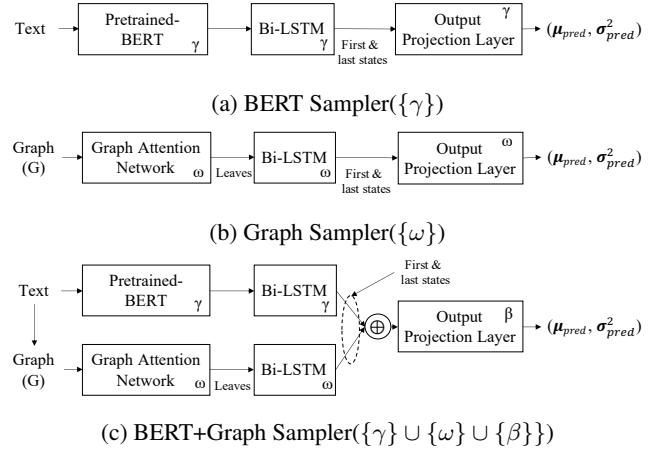
where  $D$  is the number of dimensions of the latent representation. During inference, we replace  $z_o$  from the reference encoder  $q_\phi$ , by  $z_{pred}$  obtained from the sampler  $s_\chi$ . Now we dive into various sampler architectures.

### 3.2.1. BERT Sampler

BERT is a masked language model known to capture contextual information in a given sentence [16]. This contextual information is captured in the word-piece embeddings provided by BERT for a given text. Since there has been a lot of work showing correlations between the contextual information captured by BERT and the semantics of the sentence [20], we use BERT as a semantic sampler. As shown in Fig. 2a, we use a BERT model pre-trained on long-form text to get contextual word-piece embeddings from the text. These word-piece embeddings are passed through a bidirectional LSTM. We concatenate the first and last hidden states to get a single sentence-level representation which is then projected to obtain  $\mu_{pred}$  and  $\sigma_{pred}^2$ . The sampler’s parameters are denoted as  $\{\gamma\} \subseteq \{\chi\}$ , and are trained using the loss in Eq. 3, while also fine-tuning BERT at a low learning rate.

### 3.2.2. Graph Sampler

Constituency parse trees have long been used to represent the grammatical structure of a piece of text [14], and their correlation to prosody is well known [21]. As shown in Fig. 1 from [15], these parse trees  $Tr$  have the words in a sentence as their leaves, and each leaf has only one parent node which represents the part-of-speech of that word. Upon removing the word nodes, the tree  $Tr'$  with its internal nodes can be considered as a representation of the syntax of the sentence. Since all trees  $Tr$  can be represented as undirected acyclic graphs  $G$ , we propose to use the tree as input to a Graph Neural Network. Graph Neural Networks have been used to learn representations at the node level, while exploiting the structure of the input [17, 22, 23]. We train a Message-Passing based Graph Attention Network (MPGAT) [17] with one attention head, which generates a node level representation based on the structure of the tree. In one pass of messages between connected nodes, a representation is learnt at every node depending on the connected nodes and itself. We pass  $N$  such messages, where  $N$  is the 75th percentile of the distribution of graph diameters, so that every node also has information pertaining to nodes that are not its immediate neighbours. As shown in Fig. 2b, we extract  $G$  from the text, and pass it through MPGAT. Upon obtaining a representation at each node, we extract the representations at the leaves in a depth-first order in order to obtain the representations at the parts-of-speech nodes in relation to the structure of the tree. We pass them through a bidirectional LSTM, concatenate the first and last hidden states to get a sentence-level representation, and project it to obtain  $\mu_{pred}$  and  $\sigma_{pred}^2$ . The sampler’s parameters are denoted as  $\{\omega\} \subseteq \{\chi\}$ , and are trained using the loss in Eq. 3.



**Fig. 2:** Fig. (a) and (b) show the architectures of the BERT and Graph Samplers. Fig. (c) shows the architecture of the BERT+Graph Sampler used in **Stage II** of training Kathaka.

### 3.2.3. BERT+Graph Sampler

In Fig. 2c, we combine the BERT and Graph samplers by concatenating the sentence-level representations obtained from each of the methods, and projecting them to obtain  $\mu_{pred}$  and  $\sigma_{pred}^2$ . We denote the set of all parameters in this sampler as  $\{\gamma\} \cup \{\omega\} \cup \{\beta\} \subseteq \{\chi\}$ , where  $\{\beta\}$  is the set of projection parameters, and we train using Eq. 3.

## 3.3. Prosody-dependent duration modelling

We condition our duration model on the phoneme embeddings and the sampled prosodic latent representation,  $z_{pred} \sim \mathcal{N}(\mu_{pred}, \sigma_{pred}^2)$ . Thus, as shown in Fig. 1b, we are modelling  $d = f_\nu(\mathbf{p}, z_{pred})$ .

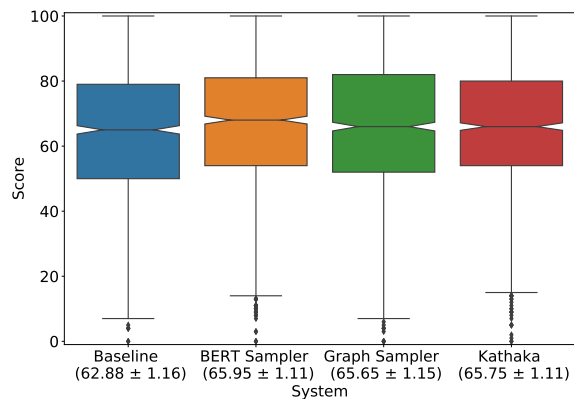
## 3.4. Training & Inference

Our model is trained in 3 steps: 1) we train the model with a variational reference encoder and durations obtained from forced alignment as in Section 3.1; 2) we train a sampler on text as in Section 3.2; 3) we train a prosody-dependent duration model as in Section 3.3 using the prosodic latent space  $\mathcal{N}(\mu_{pred}, \sigma_{pred}^2)$  learnt in Step 2. During inference, we perform 3 steps: a) we use the sampler to predict the latent vector  $z_{pred}$ ; b) we concatenate  $z_{pred}$  obtained in Step a, with the phoneme embeddings  $Y$ , to predict durations; c) we use upsampled phoneme encodings  $Y^\uparrow$ ,  $z_{pred}$ , and the predicted durations, to get mel-spectrograms using the model trained in Step 1 of training. Speech is synthesised from mel-spectrograms by using a WaveNet vocoder [5].

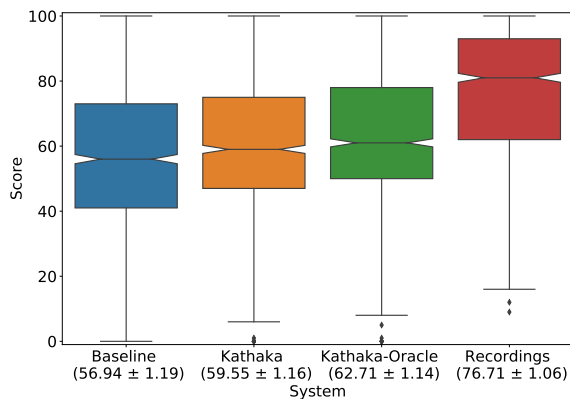
# 4. EXPERIMENTS

## 4.1. Data

We used 38.5 hours of an internal long-form US English dataset recorded by a female narrator. 33 hours of this dataset was used as the training set and the remainder as the test set.



(a) MUSHRA Scores from Ablation Study



(b) MUSHRA Scores from Gap Reduction Study

**Fig. 3:** In both figures, we show the box plots of the scores obtained by each of the systems in the test. We show the mean score with 95% confidence interval range next to the system names. In Fig. (a), all the samplers are statistically significantly better than the baseline. In Fig. (b), Kathaka obtains a statistically significant relative improvement of 13.2% over the baseline.

## 4.2. Evaluation

We conducted 2 separate MUSHRA [24] tests for evaluating the naturalness of our system and the impact of different sampling techniques. Both MUSHRA tests were taken by 25 native US English listeners. Each test consisted of 50 samples which were 15 seconds in length on average. The listeners rated each system on a scale from 0 to 100 in terms of naturalness. We evaluated statistical significance using pairwise two-sided Wilcoxon signed-rank tests.

### 4.2.1. Ablation Study

We compared each of the samplers mentioned in Section 3.2, with the baseline in a 4-system MUSHRA. Each of the samplers showed a statistically significant (all systems had a  $p$ -value  $< 10^{-3}$ , when compared to the baseline) improvement over the baseline, as seen from Fig. 3a. There is no statistical significance in the differences in the scores between the samplers themselves. This shows that both the BERT and Graph samplers may be equally capable at the task of sampling from the learnt latent prosodic space, therefore, even upon combining them, there is no statistically significant difference in the MUSHRA scores of the samplers. We selected the BERT+Graph sampler as the sampler for the Kathaka model, and used it to measure the reduction in gap as it had the lowest minimum score among the three samplers.

### 4.2.2. Gap Reduction Study

In this MUSHRA test, we evaluated 4 systems, namely: 1) the baseline, 2) Kathaka, 3) Kathaka-Oracle, the VAE based NTTS model with oracle prosodic embeddings obtained from recordings, and 4) the original recordings from the narrator. Kathaka obtained a relative improvement over the baseline by

a statistically significant 13.2% ( $p$ -value  $< 10^{-3}$ ), as shown in Fig. 3b. The model using oracle prosodic representations, Kathaka-Oracle, reduces the gap to recordings by 29.19% ( $p$ -value  $< 10^{-4}$ ). We hypothesise that the gap between Kathaka and Kathaka-Oracle is due to the sampler being trained at the sentence level. A sentence may not contain all the context required to determine the appropriate prosody with which a piece of text should be rendered, especially in long-form speech synthesis. While using a sampler trained at the sentence level shows a significant relative improvement, we speculate that exploiting the context available beyond the sentence, multi-speaker data, and other prosodically relevant features, could help to further reduce this gap.

## 5. CONCLUSION

We presented Kathaka, an NTTS model trained using a novel two-stage training approach for generating speech with contextually appropriate prosody. In the first stage of training, we learnt a distribution of sentence-level prosodic representations. We then introduced a novel sampling mechanism of using trained samplers to sample from the learnt sentence-level prosodic distribution. We introduced two samplers, 1) the BERT sampler which uses contextual word-piece embeddings from BERT and 2) the Graph sampler where we interpret constituency parse trees as graphs and use a Message Passing based Graph Attention Network on them. We then combine both these samplers as the BERT+Graph sampler, which is used in Kathaka. We also modify the baseline duration model to incorporate the latent prosodic information. We conducted an ablation study of the samplers and showed a statistically significant improvement over the baseline in each case. Finally, we compared Kathaka against a baseline, and showed a statistically significant relative improvement of 13.2%.

## 6. REFERENCES

- [1] Yuxuan Wang, RJ Skerry-Ryan, Daisy Stanton, et al., “Tacotron: Towards End-to-End Speech Synthesis,” *Proc. Interspeech 2017*, pp. 4006–4010, 2017.
- [2] Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, et al., “Natural TTS synthesis by conditioning wavenet on mel spectrogram predictions,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [3] Chengzhu Yu, Heng Lu, Na Hu, et al., “Durian: Duration informed attention network for multimodal synthesis,” *arXiv preprint arXiv:1909.01700*, 2019.
- [4] Naihan Li, Shujie Liu, Yanqing Liu, Sheng Zhao, and Ming Liu, “Neural speech synthesis with transformer network,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019, vol. 33, pp. 6706–6713.
- [5] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, et al., “WaveNet: A Generative Model for Raw Audio,” in *9th ISCA Speech Synthesis Workshop*, 2016, pp. 125–125.
- [6] Nal Kalchbrenner, Erich Elsen, Karen Simonyan, et al., “Efficient Neural Audio Synthesis,” in *Proceedings of the 35th International Conference on Machine Learning*, 2018, pp. 2410–2419.
- [7] RJ Skerry-Ryan, Eric Battenberg, Ying Xiao, et al., “Towards End-to-End Prosody Transfer for Expressive Speech Synthesis with Tacotron,” in *International Conference on Machine Learning*, 2018, pp. 4693–4702.
- [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 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6945–6949.
- [9] Sri Karlapati, Alexis Moinet, Arnaud Joly, Viacheslav Klimkov, Daniel Sáez-Trigueros, and Thomas Drugman, “CopyCat: Many-to-Many Fine-Grained Prosody Transfer for Neural Text-to-Speech,” *arXiv preprint arXiv:2004.14617*, 2020.
- [10] Kei Akuzawa, Yusuke Iwasawa, and Yutaka Matsuo, “Expressive speech synthesis via modeling expressions with variational autoencoder,” *arXiv preprint arXiv:1804.02135*, 2018.
- [11] Younggun Lee and Taesu Kim, “Robust and fine-grained prosody control of end-to-end speech synthesis,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5911–5915.
- [12] Tomoki Hayashi, Shinji Watanabe, et al., “Pre-Trained Text Embeddings for Enhanced Text-to-Speech Synthesis,” *Proc. Interspeech 2019*, pp. 4430–4434, 2019.
- [13] Huaiping Ming, Lei He, Haohan Guo, and Frank K Soong, “Feature reinforcement with word embedding and parsing information in neural TTS,” *arXiv preprint arXiv:1901.00707*, 2019.
- [14] 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.
- [15] Shubhi Tyagi, Marco Nicolis, Jonas Rohnke, Thomas Drugman, and Jaime Lorenzo-Trueba, “Dynamic prosody generation for speech synthesis using linguistics-driven acoustic embedding selection,” *arXiv preprint arXiv:1912.00955*, 2019.
- [16] Jacob Devlin, Ming-Wei Chang, et al., “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in *NAACL-HLT (1)*, 2019.
- [17] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio, “Graph Attention Networks,” in *International Conference on Learning Representations*, 2018.
- [18] Casper Kaae Sønderby, Tapani Raiko, et al., “How to Train Deep Variational Autoencoders and Probabilistic Ladder Networks,” in *33rd International Conference on Machine Learning (ICML 2016) International Conference on Machine Learning*, 2016.
- [19] Michael Wagner and Duane G Watson, “Experimental and theoretical advances in prosody: A review,” *Language and cognitive processes*, vol. 25, no. 7-9, pp. 905–945, 2010.
- [20] Anna Rogers, Olga Kovaleva, and Anna Rumshisky, “A primer in bertology: What we know about how BERT works,” *arXiv preprint arXiv:2002.12327*, 2020.
- [21] Arne Köhn, Timo Baumann, and Oskar Dörfler, “An empirical analysis of the correlation of syntax and prosody,” *arXiv preprint arXiv:1806.05900*, 2018.
- [22] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel, “Gated graph sequence neural networks,” *arXiv preprint arXiv:1511.05493*, 2015.
- [23] Jie Zhou, Ganqu Cui, Zhengyan Zhang, Cheng Yang, et al., “Graph neural networks: A review of methods and applications,” *arXiv preprint arXiv:1812.08434*, 2018.
- [24] Recommendation BS ITU-R, “1534-1, “Method for the subjective assessment of intermediate quality levels of coding systems (MUSHRA),” *International Telecommunication Union*, 2003.