

CAMP: A TWO-STAGE APPROACH TO MODELLING PROSODY IN CONTEXT

Zack Hodari*, Alexis Moinet, Sri Karlapati, Jaime Lorenzo-Trueba, Thomas Merritt, Arnaud Joly, Ammar Abbas, Penny Karanasou, Thomas Drugman

Amazon Research, Cambridge, United Kingdom

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

zack.hodari@ed.ac.uk, {amoinet, srikarla, drugman}@amazon.com

ABSTRACT

Prosody is an integral part of communication, but remains an open problem in state-of-the-art speech synthesis. There are two major issues faced when modelling prosody: (1) prosody varies at a slower rate compared with other content in the acoustic signal (e.g. segmental information and background noise); (2) determining appropriate prosody without sufficient context is an ill-posed problem. In this paper, we propose solutions to both these issues. To mitigate the challenge of modelling a slow-varying signal, we learn to disentangle prosodic information using a word level representation. To alleviate the ill-posed nature of prosody modelling, we use syntactic and semantic information derived from text to learn a context-dependent prior over our prosodic space. Our context-aware model of prosody (CAMP) outperforms the state-of-the-art technique, closing the gap with natural speech by 26%. We also find that replacing attention with a jointly-trained duration model improves prosody significantly.

Index Terms— speech synthesis, TTS, prosody, duration modelling, representation learning

1. INTRODUCTION

Producing realistic speech with text-to-speech (TTS) synthesis can be split into two goals: naturalness and appropriateness. We define naturalness as similarity of acoustic quality to human speech, and appropriateness as a measure of how well a prosodic rendition fits the given context – where context refers to any information used in prosody planning.

While the naturalness of state-of-the-art TTS is near identical to human speech [1], the prosody may not always be realistic [2]. Overall, prosody in TTS might be described as boring or flat [3], and can become fatiguing for listeners [4].

Unrealistic prosody modelling is often linked to a lack of contextual information. Indeed, without sufficient context, predicting prosody is an ill-posed problem [5], as any number of prosodies could be deemed appropriate for a given text. We believe that this is the biggest limitation of current state-of-the-art TTS models. While it may be possible to learn prosody implicitly, generating appropriate prosody without additional context information is challenging. In this paper, we consider suprasegmental prosody: prosodic phenomena

above the phonetic level [6, 7]. Suprasegmental prosody is influenced by a broad range of context information, from syntax and semantics to affect, pragmatics, and setting [8, 9].

Modelling suprasegmental prosody is made more difficult due to its medium; speech carries segmental information, background noise, and prosody in parallel [6]. Suprasegmental prosody exists at a different resolution to these as it varies more slowly [7], making it challenging to model explicitly.

To deconstruct these problems, we propose a novel two-stage approach: context-aware modelling of prosody (CAMP). In *stage-1* a prosody representation is learnt, and in *stage-2* that representation is predicted using context. We also improve Tacotron using a jointly-trained duration model.

Representation learning for speech has been investigated using unsupervised methods [10, 11, 12]. While there are many methods to define prosodic correlates [13, 14, 15], there is less work on unsupervised representation learning for prosody specifically. Most prosody representations are learnt at the sentence-level [16, 17]. However, these are too coarse and are not able to perfectly reconstruct prosody. To accurately capture prosody we need a sequence of representations, e.g. word or phrase level [18, 19]. Following CopyCat [20], we learn a sequence of representations from the mel-spectrogram, but we do so at the word-level (*stage-1*).

As demonstrated by VQ-VAE [10], a learnt prior is capable of unsupervised phonetic discovery when using a high bit rate representation. Reducing the bit rate of the representation, by using a longer unit length, should allow the learnt prior to capture longer-range effects, such as prosodic patterns. The linguistic linker introduced by Wang et al. [18] experiments with different unit lengths when modelling F_0 , but not other aspects of prosody such as rhythm and intensity. Representations extracted from a spectrogram can capture these aspects of prosody [21, 17]. TP-GST [21] predicts embeddings, but uses segmental information. Tyagi et al. [17] use additional context information, but don't train a prediction model. We argue that additional context information is vital to improving prosody quality in TTS. As such, we propose *stage-2*: use suprasegmental context information to predict the sequence of prosody representations from *stage-1*.

In Section 2, we present our baselines. In Section 3, we describe the two-stage training of CAMP.

*Work performed during an internship at Amazon TTS Research.

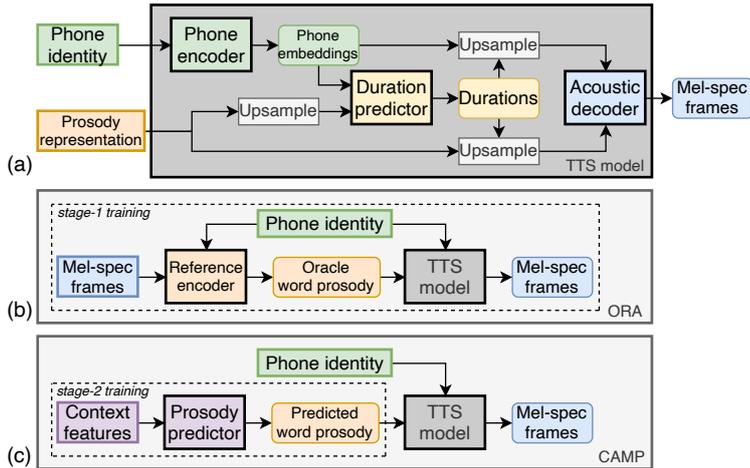


Fig. 1: (a) TTS model – synthesis driven by a prosody representation. (b) ORA – uses prosody copied from speech. (c) CAMP – uses prosody predicted from context. Dashed boxes in (b) and (c) correspond to two training stages in Section 3.

2. BASELINE SEQUENCE-TO-SEQUENCE TTS

Attention-based model — We use Tacotron-2 [22] as our attention-based baseline, s2s. Tacotron-2 consists of a phoneme encoder, attended over by an auto-regressive acoustic decoder using location-sensitive attention.

Duration-based model — Attention serves two purposes in TTS: prediction of rhythm and summarisation of relevant phonetic context. FastSpeech-2 [23] and DurIAN [24] demonstrated that: a duration prediction model can handle rhythm; and convolutions, self-attention, or bi-directional recurrent layers can summarise local context. We introduce a second baseline, *DURIAN+*, that uses a phone duration model instead of attention. *DURIAN+* is similar to DurIAN but, like FastSpeech-2, our duration model is trained jointly and shares a phone encoder with the acoustic decoder. Thus, the phone embeddings are influenced by both acoustic and duration losses, in turn providing richer inputs to the duration model at inference time. Unlike DurIAN, *DURIAN+* does not use style control or prosodic boundary tokens.¹ As shown in Section 4.3.2, it is a competitive state-of-the-art baseline.

3. TWO-STAGE PROSODY MODELLING

Our proposed approach, *CAMP* (Figure 1c), predicts the prosody representation used to drive a “*TTS model*”, detailed in Figure 1a. In section 3.1 we discuss *stage-1* of training, where this *TTS model* is trained using a word-level reference encoder, with the goal of perfectly reconstructing the original prosody. ORA, the oracle TTS, is the *TTS model* when driven by this reference encoder (Figure 1b). Section 3.2 describes *stage-2*, where we train a prosody predictor that aims to predict the oracle prosody representations, using

¹The comparison with DurIAN is intended to link *DURIAN+* to duration-based models in the literature. In reality, *DURIAN+* is very similar to s2s.

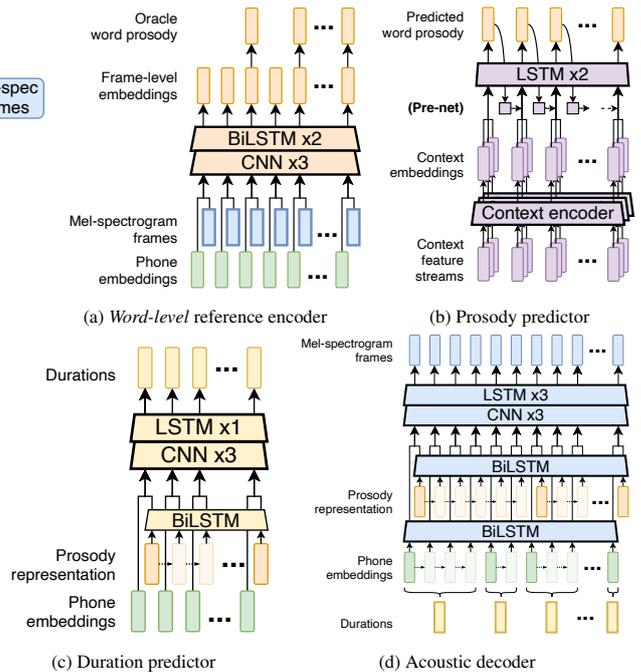


Fig. 2: Details of modules in Figure 1. (a) Represents prosody at the word level. (b) Auto-regressively predicts prosody representations using context features. (c) & (d) predict their targets using either oracle or predicted prosody representations.

suprasegmental text-derived context features. Training these stages separately ensures that the context is used to predict prosody, not lower-level acoustic information. It is essential to context-aware prosody modelling that context information feeds into a loss that explicitly focuses on prosody.

3.1. Word-level prosodic representation learning

To learn a disentangled representation of suprasegmental prosody, mel-spectrograms are auto-encoded with a temporal bottleneck [18, 20] (Figure 1b). The temporal bottleneck limits the model to only select information not predictable from other inputs, i.e. phone identity. However, determining the unit length for the bottleneck involves a trade-off: too short a unit length and segmental and background information will remain entangled [10]; too long and we will lose descriptive power [16, 25]. Since sentences contain an arbitrary amount of information [19], we use words as the unit length in the temporal bottleneck. Specifically, we use the final embedding for each word and each pause (Figure 2a) as the oracle representation to be learned by *stage-2*.

The duration model trained in *stage-1* (Figure-2c) uses phone embeddings *and* our disentangled prosody representation (upsampled to phone-level) as input. This ensures our representation can capture duration-related information.

The acoustic decoder (Figure 2d) is the same as CopyCat’s decoder [20]. This decoder uses no auto-regressive feedback of predictions, unlike the auto-regressive decoders used by the models discussed in Section 2.

3.2. Context-aware prosody prediction

We train a prosody predictor (Figures 1c and 2b) to predict the prosody representations described in Section 3.1, using text-derived features. This task has been referred to as “linguistic linking” [18], or “text-predicted tokens” [21]. We build upon this idea, emphasising the need for suprasegmental features to see improvement on this task.

To take advantage of the correlation between syntax and prosody [9], we experiment with four syntax-related context features. Each syntax feature uses a separate *context encoder*, consisting of 2 CNN layers and a BiLSTM.

- **Part-of-speech (POS)** – syntactic role of a word. We use the LAPOS tagger [26] with the Penn-Treebank tagset.
- **Word-class** – open or closed class (content or function).
- **Compound-noun structure** – flag indicating if a word is part of a compound noun. Extracted using POS tags.
- **Punctuation structure** – flag indicating punctuation.

We use one semantic *context encoder*: a pre-trained BERT_{BASE} model [27], fine-tuned on our prosody prediction task. BERT’s word-piece embeddings are aligned to word-level by average pooling.

Our prosody predictor is auto-regressive (illustrated by the curved downwards arrows in Figure 2b). This allows the model to learn how prosody transitions from one word to the next. Adding attention or self-attention might improve modelling of distant information, however, we believe this is less important as our sentences are only on average 16 words.

Our prosody predictor can be interpreted as a learnt conditional prior, where the latent space – our prosody representation – has unknown variance. Unlike VQ-VAE [10], we use a coarser-grained representation (i.e. word-level), allowing us to focus on prosody, as opposed to frame-level information.

4. EXPERIMENTS

4.1. Data

We train our models on prosodically rich data – an expressive single-speaker internal dataset. The data contains about 25,000 utterances of professionally recorded speech by a native US English female speaker. The training, validation, and test sets are approximately 30 hours, 2 hours, and 6 hours, respectively. Phoneme features are one-hot encodings of phone identity and service tokens, such as word-boundaries. We extracted 80-band mel-spectrograms with a 12.5 ms frame-shift. Durations were extracted using forced alignment with Kaldi.

4.2. Systems

We use the two single-stage baselines described in Section 2: S2S (Tacotron-2 [22]); and DURIAN+ (S2S with a jointly-trained duration model). Table 1 details DURIAN+, along with our two-stage models; CAMP, and ORA. Our proposed model, CAMP, uses prosody predicted from context features (Figure 1c). ORA represents the top-line performance of our

two-stage approach – using oracle prosody copied from human speech (Figure 1b). All models use the same 24kHz mixture of logistics WaveNet vocoder [28]. NAT is natural 24kHz speech with no vocoding, used in Section 4.3.3.

All models were trained using the Adam optimiser. An L1 loss was used for acoustic and duration losses. S2S, DURIAN+, and ORA were trained for 300,000 steps, with a learning rate of 0.001 and a decay factor of 0.98. Prosody prediction used a Huber loss with $\rho = 1$, this produced marginally better prosody than L1 or L2 during informal listening. Prosody predictors were trained for 100,000 steps, with a learning rate of 0.0001 and a decay factor of 0.98.

4.3. Subjective evaluation

For our listening tests we use 96 sentences, chosen randomly from the test set. In Section 4.3.1, we evaluate the efficacy of our syntactic and semantic features. Following this, we measure the contribution of duration modelling (Section 4.3.2). This allows us to control for the effect of duration modelling when evaluating our core contribution in Section 4.3.3.

4.3.1. Context features

In Section 3.2, we described 5 different context features. To determine the contribution of each to prosody performance we conduct an ablation test using a MUSHRA-like format (no anchor and no reference) [29]. We consider three separately trained models: CAMP_{Syntax} – using 4 syntax context encoders, CAMP_{BERT} – using the BERT context encoder, and CAMP_{BERT+Syntax} – using all 5 context encoders. 20 listeners compared these three CAMP models on a scale from 0 to 100.

We perform a two-sided Wilcoxon signed-rank test on all 3 pairs of systems, and correct with Holm-Bonferroni. The results in Table 2 show that CAMP_{Syntax} is significantly worse than both other systems ($p < 0.001$). No statistically significant difference is found between CAMP_{BERT} and CAMP_{BERT+Syntax} ($p = 0.83$). We conclude that our syntactic features provide no additional information compared to BERT’s contextualised representations. This is in agreement with findings that BERT can capture both semantic *and* syntactic information [30]. As such we use CAMP_{BERT} as our proposed system (CAMP for brevity).

Table 1: System comparison. DURIAN+ is a competitive baseline. ORA represents our approaches’ best-case performance. CAMP is our proposed model using context information.

	DURIAN+	ORA	CAMP
Phone encoder	Tacotron-2 encoder		
Prosody representation	—	Reference encoder (Fig. 2a)	Prosody predictor (Fig. 2b)
Duration inputs	Phone embeddings	Phone embeddings and prosody representation	
Acoustic Decoder	Tacotron-2 decoder	CopyCat decoder (Fig. 2d)	

Table 2: MUSHRA ablation of context features in CAMP. Mean rating and 95% confidence intervals are reported. Question asked: “Rate the systems based on your preference”.

	Syntax	BERT	BERT+syntax
Mean score	63.9 \pm 1.21	66.5 \pm 1.21	66.3 \pm 1.19

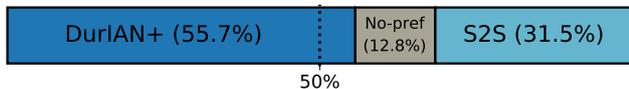


Fig. 3: Preference test demonstrating that a jointly-trained duration model improves prosody significantly over Tacotron-2. Question asked: “Choose which version you prefer”.

4.3.2. Benchmark model

We compare the attention-based S2S with the duration-based DURIAN+ to measure what improvement our duration model provides. We perform a preference test with 15 listeners.

To show that DURIAN+ improves prosody, not just robustness, we remove the 20 sentences where S2S had attention instabilities from the results presented in Figure 3. There were no sentences where DURIAN+ had stability issues. A binomial significance test demonstrates that DURIAN+ is significantly preferred over S2S ($p < 10^{-15}$).

We tried training DURIAN+ with a fine-tuned BERT as an additional encoder. However, there was little or no improvement in prosody – similar to previous findings [31, 32]. We believe this is due to the single-stage model predicting low-level acoustics, as opposed to a prosodic representation.

4.3.3. Prosody prediction

Finally, we evaluate CAMP against natural speech. We use DURIAN+ as a strong baseline following the result in Section 4.3.2. This makes the presence of a duration model a control variable, allowing us to attribute any improvements to the two-stage structure. We perform a MUSHRA-like test [29] with 4 systems: DURIAN+, CAMP, ORA, and NAT.

Since we provide isolated sentences to listeners, there is no concept of correct prosody [5], as such we do not provide a visible reference and do not require listeners to rate two systems as 0 and 100. 25 listeners completed this test.

We present the results in Figure 4, along with 95% confidence intervals, computed using percentile bootstrap. Confidences reported below are computed in the same way. We perform a two-sided Wilcoxon signed-rank test on all 6 pairs of systems, and correct with Holm-Bonferroni. We find all systems are significantly different from each other ($p \ll 10^{-10}$).

From this test we can conclude that: (1) CAMP reduces the gap between natural speech and DURIAN+ by a relative $26\% \pm 7\%$, a significant improvement in prosody quality attributed to our two-stage approach. (2) Given the best-case prediction of the prosody representation (i.e. ORA), our model could close this same gap by a relative $69\% \pm 6\%$. This suggests that, as

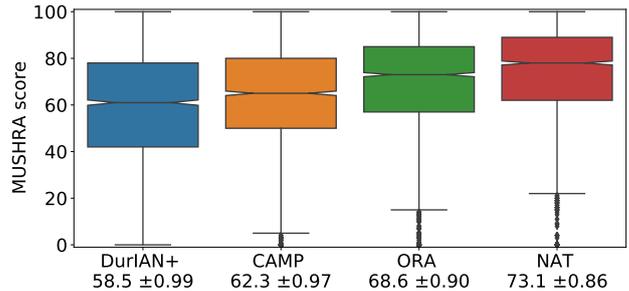


Fig. 4: MUSHRA results with CAMP. Mean rating and 95% confidence intervals are reported below system names. Question asked: “Rate the systems based on your preference”.

expected, our representation is lossy – likely due, in part, to the trade-off between disentanglement and descriptive power. (3) Using BERT as a context encoder, CAMP gets $38\% \pm 10\%$ of the way to achieving this top-line performance (i.e. the gap between DURIAN+ and ORA).

As discussed in Section 4.3.2, the use of additional context information is not enough to improve prosody modelling. We attribute the performance of CAMP to the prosodically relevant loss. The prosody predictor’s loss doesn’t focus on lower-level acoustics, but directly on prosody. We chose to use a learnt prosody representation, enabling our model to determine which information is important. Explicit prosody features, such as F_0 , could be used instead. However, these may limit the range of prosody, and are susceptible to errors.

The results show a significant improvement in prosody. They also suggest we can improve further within this two-stage paradigm: through improved representation learning; or improved prosody prediction – using more context features or increased context width (e.g. neighbouring sentences).

Our experiments used single-speaker data. Learning a multi-speaker representation would likely improve disentanglement. An investigation of a multi-speaker prosody predictor would be interesting, as this may provide insights into the relationship between different speakers’ prosodic patterns.

5. CONCLUSION

We introduced CAMP, a two-stage approach for prosody modelling. In *stage-1*, we learn a prosodic representation using an encoder-decoder model with a new word-level reference encoder. In *stage-2*, we train a prosody predictor which uses prosodically-relevant context to predict the representations learnt in *stage-1*. We stress the importance for using additional context information when predicting prosody.

An intermediate result showed that replacing attention with a jointly-trained duration model improves significantly over the state-of-the-art, both in terms of stability and prosody. Finally, we show that our proposed CAMP model, using BERT to guide prosody prediction, improves further upon this intermediate result, closing the gap between our stronger baseline and natural speech by 26%.

6. REFERENCES

- [1] Naihan Li, Shujie Liu, Yanqing Liu, Sheng Zhao, and Ming Liu, “Neural speech synthesis with Transformer network,” in *Proc. AAAI*, 2019, pp. 6706–6713.
- [2] Srikanth Ronanki, “Prosody generation for text-to-speech synthesis,” 2019.
- [3] Zack Hodari, Oliver Watts, and Simon King, “Using generative modelling to produce varied intonation for speech synthesis,” in *Proc. Speech Synthesis Workshop*, 2019, pp. 239–244.
- [4] Avashna Govender and Simon King, “Using pupillometry to measure the cognitive load of synthetic speech,” in *Proc. Interspeech*, 2018, pp. 2838–2842.
- [5] Rob Clark, Hanna Silen, Tom Kenter, and Ralph Leith, “Evaluating long-form text-to-speech: Comparing the ratings of sentences and paragraphs,” in *Proc. Speech Synthesis Workshop*, Vienna, Austria, 2019, pp. 99–104.
- [6] D Robert Ladd, *Intonational phonology*, Cambridge University Press, 2008.
- [7] Nigel G Ward, *Prosodic patterns in English conversatio*, Cambridge University Press, 2019.
- [8] Daniel Goodhue, Lyana Harrison, Yuen Tung Clémentine Siu, and Michael Wagner, “Toward a bestiary of English intonational contours,” in *Proc. NELS*, 2016, pp. 311–320.
- [9] Arne Köhn, Timo Baumann, and Oskar Dörfler, “An empirical analysis of the correlation of syntax and prosody,” in *Proc. Interspeech*, 2018, pp. 2157–2161.
- [10] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu, “Neural discrete representation learning,” in *Proc. NeurIPS*, Long Beach, USA, 2017, pp. 6306–6315.
- [11] Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli, “wav2vec: Unsupervised pre-training for speech recognition,” *arXiv:1904.05862*, 2019.
- [12] Ewan Dunbar, Robin Algayres, Julien Karadayi, Mathieu Bernard, Juan Benjumea, Xuan-Nga Cao, Lucie Miskic, Charlotte Dugrain, Lucas Ondel, Alan W Black, et al., “The zero resource speech challenge 2019: TTS without T,” in *Proc. Interspeech*, Graz, Austria, 2019.
- [13] Antti Suni, Daniel Aalto, and Martti Vainio, “Hierarchical representation of prosody for statistical speech synthesis,” *arXiv:1510.01949*, 2015.
- [14] Viacheslav Klimkov, Srikanth Ronanki, Jonas Rohnke, and Thomas Drugman, “Fine-grained robust prosody transfer for single-speaker neural text-to-speech,” in *Proc. Interspeech*, Graz, Austria, 2019, pp. 4440–4444.
- [15] Manuel Sam Ribeiro and Robert AJ Clark, “A multi-level representation of F0 using the continuous wavelet transform and the discrete cosine transform,” in *Proc. ICASSP*, Brisbane, Australia, 2015, pp. 4909–4913.
- [16] Yuxuan Wang, Daisy Stanton, Yu Zhang, RJ Skerry-Ryan, Eric Battenberg, Joel Shor, Ying Xiao, Fei Ren, Ye Jia, and Rif A Saurous, “Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis,” *arXiv:1803.09017*, 2018.
- [17] Shubhi Tyagi, Marco Nicolis, Jonas Rohnke, Thomas Drugman, and Jaime Lorenzo-Trueba, “Dynamic prosody generation for speech synthesis using linguistics-driven acoustic embedding selection,” in *Proc. Interspeech*, 2020.
- [18] Xin Wang, Shinji Takaki, Junichi Yamagishi, Simon King, and Keiichi Tokuda, “A vector quantized variational autoencoder (VQ-VAE) autoregressive neural F0 model for statistical parametric speech synthesis,” *IEEE TASLP*, pp. 157–170, 2019.
- [19] Zack Hodari, Catherine Lai, and Simon King, “Perception of prosodic variation for speech synthesis using an unsupervised discrete representation of F0,” in *Proc. Speech Prosody*, Tokyo, Japan, 2020, pp. 965–969.
- [20] 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,” in *Proc. Interspeech*, 2020.
- [21] Daisy Stanton, Yuxuan Wang, and RJ Skerry-Ryan, “Predicting expressive speaking style from text in end-to-end speech synthesis,” *arXiv:1808.01410*, 2018.
- [22] Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerrv-Ryan, et al., “Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions,” in *Proc. ICASSP*, Calgary, Canada, 2018, pp. 4779–4783.
- [23] Yi Ren, Chenxu Hu, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu, “FastSpeech 2: Fast and high-quality end-to-end text-to-speech,” *arXiv:2006.04558*, 2020.
- [24] Chengzhu Yu, Heng Lu, Na Hu, Meng Yu, Chao Weng, Kun Xu, Peng Liu, Deyi Tuo, Shiyin Kang, Guangzhi Lei, et al., “DurIAN: Duration informed attention network for multimodal synthesis,” *arXiv:1909.01700*, 2019.
- [25] Vincent Wan, Chun-an Chan, Tom Kenter, Jakub Vit, and Rob Clark, “CHiVE: Varying prosody in speech synthesis with a linguistically driven dynamic hierarchical conditional variational network,” in *Proc. ICML*, Long Beach, USA, 2019.
- [26] Yoshimasa Tsuruoka, Yusuke Miyao, and Kazama Jun’ichi, “Learning with lookahead: Can history-based models rival globally optimized models?,” in *Proc. Computational Natural Language Learning*, 2011, pp. 238–246.
- [27] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proc. NAACL-HLT*, 2019, pp. 4171–4186.
- [28] Aaron Oord, Yazhe Li, Igor Babuschkin, Karen Simonyan, Oriol Vinyals, Koray Kavukcuoglu, George Driessche, Edward Lockhart, Luis Cobo, Florian Stimberg, et al., “Parallel WaveNet: Fast high-fidelity speech synthesis,” in *Proc. ICML*, Stockholm, Sweden, 2018, pp. 3918–3926.
- [29] BS Series, “Method for the subjective assessment of intermediate quality level of audio systems,” *International Telecommunication Union Radiocommunication Assembly*, 2014.
- [30] Anna Rogers, Olga Kovaleva, and Anna Rumshisky, “A primer in BERTology: What we know about how BERT works,” *arXiv:2002.12327*, 2020.
- [31] Tomoki Hayashi, Shinji Watanabe, Tomoki Toda, Kazuya Takeda, Shubham Toshniwal, and Karen Livescu, “Pre-trained text embeddings for enhanced text-to-speech synthesis,” in *Proc. Interspeech*, Graz, Austria, 2019, pp. 4430–4434.
- [32] Wei Fang, Yu-An Chung, and James Glass, “Towards transfer learning for end-to-end speech synthesis from deep pre-trained language models,” *arXiv:1906.07307*, 2019.