



# Controllable neural text-to-speech synthesis using intuitive prosodic features

*Tuomo Raitio, Ramya Rasipuram, Dan Castellani*

Apple

## Abstract

Modern neural text-to-speech (TTS) synthesis can generate speech that is indistinguishable from natural speech. However, the prosody of generated utterances often represents the average prosodic style of the database instead of having wide prosodic variation. Moreover, the generated prosody is solely defined by the input text, which does not allow for different styles for the same sentence. In this work, we train a sequence-to-sequence neural network conditioned on acoustic speech features to learn a latent prosody space with intuitive and meaningful dimensions. Experiments show that a model conditioned on sentence-wise pitch, pitch range, phone duration, energy, and spectral tilt can effectively control each prosodic dimension and generate a wide variety of speaking styles, while maintaining similar mean opinion score (4.23) to our Tacotron baseline (4.26).

**Index Terms:** Prosody control, end-to-end neural speech synthesis, sequence-to-sequence, attention, encoder-decoder

## 1. Introduction

State-of-the-art neural text-to-speech (TTS) technology commonly consists of a sequence-to-sequence (seq2seq) neural network with attention [1, 2, 3] that maps text or phone input to a Mel-spectrogram, and a neural back-end [4, 5, 6] that converts the Mel-spectrogram into a sequence of speech samples. These two networks are trained with a large amount of speech data, enabling the generation of speech that can be indistinguishable from natural speech [3].

The prosody of the generated speech is largely defined by the seq2seq encoder–decoder neural network, in TTS commonly referred to as Tacotron [2, 3]. However, the prosody of the generated utterances often represents the average style of the database instead of having a wide range of prosodic variation. Moreover, the generated prosody is defined by the input text, which does not allow for different speaking styles for the same sentence. Therefore, the generated prosody may not always correspond to the intended message for the listener. To generate a wide range of prosodic output that is appropriate for the context, the Tacotron model must learn to factorize the various prosodic styles in the training data, and also provide a meaningful control over the prosodic output independent of the input text.

Traditionally, prosody modeling is based on schematizing and labeling prosodic phenomena and developing rule-based systems or statistical models from the derived data [7]. However, the prosodic attributes are difficult and time-consuming to annotate, and therefore such an approach is impractical for modern TTS with large databases. In contrast, end-to-end neural TTS systems learn prosody as an inherent part of the model, which makes the generation of speech with the correct linguistic content and the appropriate prosody a consistent modeling task. However, such prosodically unsupervised models lack the explicit control over the output prosody.

Recent developments in end-to-end neural TTS have provided some solutions to prosody modeling. In [8], an extension to the Tacotron architecture is introduced to capture the residual attributes that are not specified by the linguistic input.

The model learns a latent embedding space of prosody through conditional auto-encoding, derived from the target Mel-spectrograms, to enable prosody transfer.

As an improvement to [8], [9] proposed global style tokens (GSTs) to learn a clustered latent space of prosody. The same prosody encoder as in [8] is used, but a style token layer is added where the prosody embedding is used as a query to an attention layer that maps the embedding to the GSTs using combination weights. The output of the style token layer is a style embedding that describes the prosodic style. The architecture can be used for either prosody transfer using a reference audio or for directly generating a specific prosodic style by choosing a style token.

The aforementioned unsupervised methods for learning a latent space of prosody have many advantages. Unsupervised methods do not need any extra information or labeling of prosody, and all the components are jointly trained to achieve a consistent model. However, there are some drawbacks. First, the latent space represents all the residual acoustic differences in the reference Mel-spectrograms after linguistic content has been accounted for. As a result, any acoustic differences due to channel (recording setup, noise type and level) and (undesirable) variability in speaker’s voice from one day to another are also learned. Therefore, instead of focusing on meaningful prosodic aspects of speech, the model learns an entangled representation of prosody and unknown underlying acoustic factors. Second, manual listening of the tokens is required to determine what sort of style tokens the model has learned from the data, and how to utilize these for reproducing a desired prosody. Third, the number of style tokens must be set heuristically, and the model is not guaranteed to find perceptually distinct and prosodically meaningful clusters for each of the tokens. The individual style tokens are often either similar to each other, or it may be hard to make practical use of the learned style tokens as they may not represent any desired styles.

The prosody of language covers all aspects of speech that are not related directly to the articulation for the linguistic expression. Although there is no agreed number of prosodic variables, there is a set of variables that is widely agreed to have a major contribution to prosody. These variables are 1) pitch, 2) length of the speech sounds, 3) loudness, and 4) timbre or voice quality [10, 11]. These four subjective auditory variables can be accurately measured using the following acoustic variables: fundamental frequency, phone duration, speech energy, and spectral tilt, respectively. Assuming that the four acoustic features cover the prosodic space, then there are certain benefits in using them for prosody modeling and control: 1) they are disentangled so that they can be independently varied, 2) they are intuitive so that it is easy to understand their effect on prosody, 3) they are independent from other acoustic factors, such as background noise or other recordings conditions, thus making them robust for prosody modeling.

In this work, we take a similar, unified end-to-end approach

for prosody modeling as in [8, 9], but instead of performing conditional auto-encoding using the Mel-spectrograms, we use prosodically meaningful acoustic features derived from the speech signal, similar to [12]. This approach has the benefit of guiding the model to learn only perceptually relevant acoustic differences that contribute to prosody. We predict the acoustic features directly from the encoder outputs, similar to [13], but with a more focused set of prosodic features, which enables direct and intuitive control over each prosodic dimension.

Our main contribution in this work is a unified and simple neural network architecture for prosody modeling and control using intuitive prosodic features. In contrast to [8, 9, 13], we use well-defined acoustic features for prosody modeling instead of Mel-spectrograms. Also, in contrast to [12], we use a simple unified encoder-decoder architecture that is trained as a single model. While prosodic features such as fundamental frequency, phone duration, and speech energy, have seen increased interest for prosody modeling in neural TTS [12, 14, 15, 16], we extend our model to utilize the fourth dimension of prosody: spectral tilt [10, 11]. We show that each feature has a specific and effective control over the prosodic space, and that the proposed method can generate various speaking styles while maintaining high quality<sup>1</sup>.

## 2. Technical Overview

The baseline sequence-to-sequence with attention model is similar to the Tacotron 2 model [3]. The input is a phoneme sequence with punctuation and word boundaries, and the output is a Mel-spectrogram. Our experiments show that using location-sensitive monotonic attention [17] or location-sensitive step-wise monotonic attention [18] results in a more robust alignments than global alignment. For consistency, the former is used throughout this study. We use similar mechanism to [3] for utterance end-point prediction. We also use a streaming architecture to reduce the lag in the output. We are able to fit three Tacotron models in a single GPU, each of them generating Mel-spectrograms 5x real-time, while our on-device implementation runs 8.5x real-time on a mobile device.

To generate a sequence of speech samples from the Mel-spectrogram, we use an autoregressive recurrent neural network (RNN) based architecture, similar to WaveRNN [5]. The model consists of a single RNN layer with 512 hidden units, conditioned on Mel-spectrogram, followed by two fully-connected layers ( $512 \times 256$ ,  $256 \times 256$ ), with single soft-max sampling at the output. The model is trained with pre-emphasized speech sampled at 24 kHz and  $\mu$ -law quantized to 8 bits for efficiency. We can run three of these models in parallel on a single GPU, each of them generating speech 7.7x real-time, while our on-device implementation runs 3.3x real-time on a mobile device.

More information about the architecture and the on-device implementation of the baseline system can be found in [19].

### 2.1. Proposed Prosody Control Architecture

We extend the baseline architecture by introducing a prosody encoder [8] that predicts prosodic features from the encoder outputs. The prosody encoder consists of a stacked 3-layer LSTM with hidden state size of 128 at each layer. The last state of the LSTM is fed to a fully connected projection layer with a tanh non-linearity. The input to the prosody encoder is the phone embedding, and the output is the prosody feature vector that is con-

<sup>1</sup>Speech samples can be found at <https://apple.github.io/neural-tts-with-prosody-control/>.

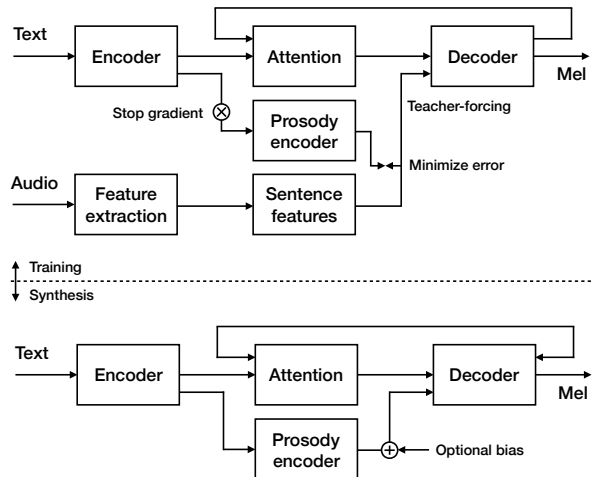


Figure 1: An overview of the proposed prosody modeling encoder-decoder with attention architecture. Top part describes the training phase where the prosody encoder learns to predict the sentence-wise prosodic features, and the decoder is conditioned on the ground-truth features (teacher-forcing). The bottom part describes the inference phase where prosody encoder predicts prosodic features to condition the decoder, with an additional bias option for prosody control.

catenated with the decoder input for prosody control. To jointly train the model with both phone and prosody encoders, we add a stop gradient operation between the two modules. The stop gradient ensures that the training of the prosody encoder will not negatively affect the training of the phone encoder. We also use teacher-forcing of the prosodic features to efficiently train the network. The model architecture is illustrated in Fig. 1.

### 2.2. Prosodic Features

To learn a prosodically meaningful latent space for prosody control, we use acoustic features extracted from the original speech to condition the model. We use fundamental frequency (pitch), phone duration, speech energy, and spectral tilt to model the prosodic space. These features are easy to calculate from speech signals and are robust against background noise or other recordings conditions. These features are also disentangled to a large degree so that they can be independently varied. Overall, these features provide an intuitive control over the prosodic space.

We extract the pitch of speech using 3 pitch estimators, and then vote for the most likely pitch trajectory. We compute two features from the voted pitch contour: average log-pitch and log-pitch range of voiced speech. The latter is calculated as the difference between the 0.05 and 0.95 quantile values of the frame-wise log-pitch contour for each utterance. We use automatic speech recognition to force-align the text and audio to obtain the phone durations, and calculate the average log-phone duration per utterance. We extract the frame-wise log-energy<sup>2</sup> of each utterance (excluding silences), and calculate the frame-wise spectral tilt of voiced speech using first order all-pole modeling<sup>3</sup>, and finally average them per utterance.

The utterance-wise acoustic features, log-pitch, log-pitch range, log-phone duration, log-energy, and spectral tilt, are then

<sup>2</sup>Speech energy  $E$  is calculated as  $E = 20 \log_{10}(\hat{x})$ , where  $\hat{x}$  is the average absolute sample amplitude, excluding silence parts.

<sup>3</sup>Spectral tilt is measured using the predictor coefficient of the first order all-pole filter.

normalized to a range  $[-1, 1]$  by first calculating the median ( $M$ ) and the standard deviation ( $\sigma$ ) of each feature, and then projecting the data in the range  $[M - 3\sigma, M + 3\sigma]$  into  $[-1, 1]$ . Finally, we clip values  $|x| > 1$  so that all data is in the range  $[-1, 1]$ . This process is similar to the approach in [12].

### 3. Experiments

#### 3.1. Data

We use an internal 36-hour dataset of an American English female voice. In addition, we have a 16-hour conversational expressive dataset from the same speaker to expand the prosodic coverage. All recordings were produced in a professional studio and post-processed for high quality. The original speech data is down-sampled to 24 kHz for training the neural TTS system. 80-dimensional Mel-spectrograms are computed from pre-emphasized speech using short-time Fourier transform (STFT) with 25 ms frame length and 10 ms shift.

#### 3.2. Models

We trained the following three models:

1. **Baseline:** High-quality baseline model trained with the 36-hour dataset.
2. **Prosody 36h:** Our proposed prosody control model trained with the 36-hour dataset.
3. **Prosody 52h:** Our proposed prosody control model trained with the 36-hour dataset combined with the 16-hour conversational dataset for better prosodic coverage.

For the prosody control models, we used pitch, pitch range, phone duration, speech energy, and spectral tilt as the conditioning features. We train all the models for 3 million steps using a single GPU and batch size of 16. All systems use the same back-end WaveRNN model [19], trained with the 36-hour dataset, to generate speech from the Mel-spectrograms.

#### 3.3. Objective Measures

To measure how well the model can control each prosodic dimension, we synthesized speech at different points in the  $[-1, 1]$  scale. Each dimension was varied independently. As text material, we used 199 sentences with general text and responses typical to a voice assistant. A total of 16,517<sup>4</sup> utterances were synthesized. Then we compared how well the output speech reflects the given target prosodic bias by measuring the corresponding acoustic features from the synthetic utterances.

Fig. 2 shows the measured acoustic features at each target bias value, normalized by the procedure in Sec. 2.2, using the original normalization values calculated over the whole database. The original acoustic feature values with respect to the scale  $[-1, 1]$  are shown in Table 1. The results show that the target bias values of the prosodic features are well reflected in the output synthetic speech. Pitch shows almost ideal correlation between the target and measured values, whilst the remaining features show good correlation. The system trained with expressive speech (Prosody 52h) shows better correlation at the low-end of duration, energy, and spectral tilt.

#### 3.4. Listening Tests

We carried out a number of listening tests to evaluate the subjective performance of the proposed approach. First, we evaluated

<sup>4</sup>16,517 = 5 dimensions  $\times$  9 values  $\times$  199 samples  $\times$  2 systems + 199 baseline samples - (5-1)  $\times$  2  $\times$  199 samples repeated at 0 bias.

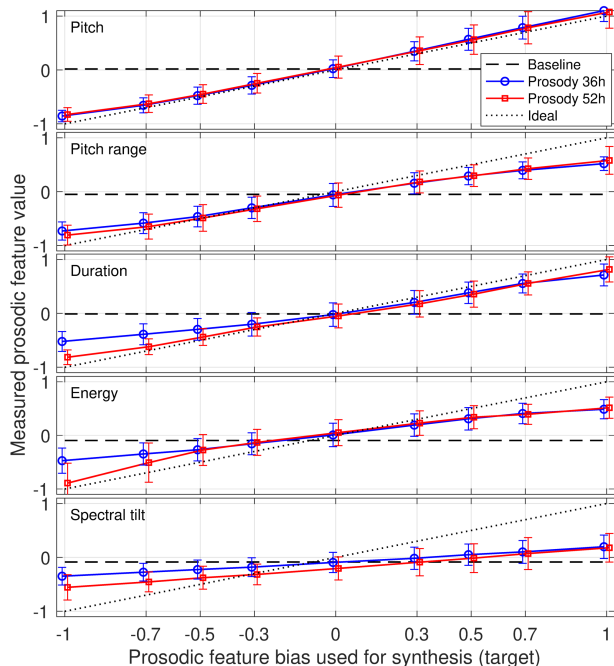


Figure 2: Means and standard deviations of the measured prosodic features with respect to target bias values.

Table 1: Prosodic feature values in normalized scale  $[-1, 1]$ .

Feature	Unit	-1.0	0.0	1.0
Pitch	Hz	144.2	234.0	323.7
Pitch range	Hz	50.9	355.8	660.8
Duration	ms	32.7	117.6	202.5
Energy	dB	-26.2	-20.7	-15.2
Spectral tilt	-	-0.997	-0.978	-0.958

the overall TTS naturalness of all the models without prosody control (zero bias). The intent was to assess the feasibility of the proposed architecture in place of the baseline system, even when manual prosody control is not specifically used.

From each system, we used the 199 synthetic utterances described in Sec. 3.3. A 5-point mean opinion score (MOS) test was performed by 66 individual American English native speakers using headphones, resulting in a total of 8,986 responses. The results in Table 2 show that all the three systems yield a high MOS, and that there is no statistically significant difference ( $p > 0.10$ ) in quality between the three systems when listening to the speech samples in isolation. Therefore, we can conclude that the proposed prosody control systems can generate high quality synthetic speech.

MOS test can be somewhat insensitive to small differences in speech samples. For a more accurate assessment between the three systems, we performed AB listening tests where listeners were presented with a pair of speech samples, and they were asked to choose the one that sounded better, or choose no preference. We evaluated both prosody control models against the baseline. Ten proficient English speakers evaluated a total of 50 sample pairs using headphones, each test subject having a randomized set of different samples. The results, shown in Table 2, indicate that the baseline system is preferred over both prosody control models ( $p < 0.005$ ). This degradation in quality could arise from the explicit prosody control that adds complexity, and thus makes it harder to make consistent prosody predictions. Also, matching the quality of our precisely tuned baseline system is a challenging task. The AB test also shows

Table 2: MOS (first row) and AB test results for the three systems without any manual prosody control. 95% confidence intervals computed from the  $t$ -distribution are shown for MOS.

Baseline	Prosody 36h	Prosody 52h	No pref.
4.26 ( $\pm$ 0.030)	4.23 ( $\pm$ 0.032)	4.18 ( $\pm$ 0.033)	
48.3%	25.9%	-	25.9%
43.5%	-	25.7%	30.9%

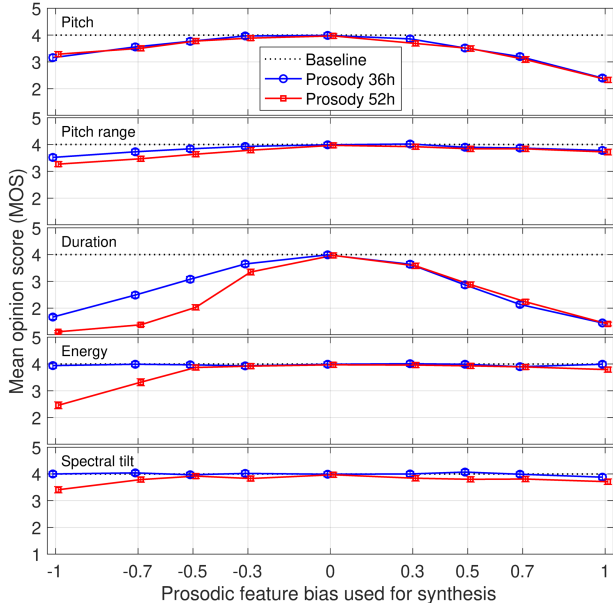


Figure 3: Means and 95% confidence intervals of MOS measured over all features and target bias values.

that the prosody control model trained with additional expressive speech is rated lower than the other models. This degradation could arise from the more challenging and varying speech material, which reduces the consistency of the output.

We also evaluated the MOS of the prosody control models at different points in the  $[-1, 1]$  feature scale. We used 40 sentences described in Sec. 3.3. A total of 3,320 utterances were synthesized, and each was given 10 ratings, totaling 33,200 individual ratings by American English native speakers using headphones. The results in Fig. 3 show that the MOS has a tendency to slightly decrease towards both ends of the prosodic scale  $[-1, 1]$ . This makes sense as the resulting speaking styles are less typical in comparison to the expected default delivery style of the voice talent. However, for some features, there is clear degradation in quality at the extreme ends of the scale as it results in speech production modes that do not occur in human speech. For example, when pitch bias is set to 1.0, the utterance exhibits a sustained high pitch on the limits of the voice talent’s voice. Similarly, very low or high phone duration in all the phones of an utterance does not correspond to normal speech production either. Due to wider prosodic range, the Prosody52h system shows more degradation at the extremes, especially with low energy where speech turns to almost whisper.

### 3.5. Additional Experiments

One of the aims of the prosody control architecture was to enable the generation of new improved, alternative, or more expressive versions of utterances that are key to the user experience of a voice assistant. We hypothesize that the intuitive control of prosodic features enables easy tuning of synthetic utterances. To test this hypothesis, we gathered 10 utterances that

were lacking appropriate prosody for a voice assistant, consisting of either synthetic speech or original recordings from the studio. In a small proof-of-concept experiment, we presented dialog writers with a tool that had levers in the range  $[-1, 1]$  for each of the prosodic features, and let them generate new synthetic utterances by adjusting the prosodic dimensions. The Prosody 52h model was used for generating the samples in order to have greater prosodic flexibility. We then created a listening test to assess if the new speech samples were rated more appropriate than the existing ones. Listeners were presented with the original synthesis (or recording) and two new modified versions of the same utterance, and were asked to choose the one they preferred. Eight American English native speakers (different from those that created the data) performed the listening test. The prosody-modified samples were preferred 63% of the time, and 44% of the time even over original studio recordings. Although this test was very small in scale and without statistical power, it demonstrates a potential procedure of creating synthetic speech with improved prosody.

The best way to demonstrate the expressive capability of our proposed system is by listening. We present a set synthetic speech samples in [20], generated by varying the bias of each of the prosodic dimension gradually in the range  $[-1, 1]$ .

## 4. Discussion

There are multiple benefits in the proposed architecture. First, the prosodic dimensions are inherently disentangled so that they can be independently varied. Second, the prosodic features are intuitive so that it is easy to understand their effect on prosody. Third, the prosodic features are independent from other acoustic factors such as background noise or other recordings conditions, thus making them robust for prosody modeling.

On the downside, the proposed model only provides sentence-level prosody control, and is lacking finer-grained control as in [14, 16]. Although the sentence-level features could be varied at the decoder input, accurate modification requires alignment information [21]. Alternatively, the prosodic features can be concatenated to the encoder output to enable phone-level prosodic modification. However, defining a good reference prosody becomes a problem, and the method would become similar to fine-grained prosody transfer [15].

Like [13], the proposed model predicts appropriate prosody directly from the text. Thus, the model can be used for synthesizing any text with appropriate style, provided with enough training data for the text and style. However, in our study, we did not achieve an overall better speech naturalness in comparison to our baseline system; sometimes the prosody was evaluated less natural in the AB test, indicating less stable prosody prediction for some text inputs and target style. For future work, we will investigate the robustness of the prosodic features, and explore different model architectures to overcome these issues.

## 5. Conclusions

We proposed a unified neural TTS system with prosody control capability using intuitive prosodic features. Subjective results show that the proposed model can synthesize speech with quality similar to our baseline, while being able to reproduce various prosodic styles. The objective results also show a clear correlation between the target prosodic feature values and the measured values from synthetic speech. Additional experiments demonstrate how the model can be used for creating new improved versions of synthetic speech with a desired prosody.

## 6. References

- [1] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to sequence learning with neural networks,” in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds., 2014, pp. 3104–3112.
- [2] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. Le, Y. Agiomyrgiannakis, R. Clark, and R. A. Saurous, “Tacotron: Towards end-to-end speech synthesis,” in *Interspeech*, 2017.
- [3] 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. Agiomyrgiannakis, and Y. Wu, “Natural TTS synthesis by conditioning WaveNet on Mel spectrogram predictions,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 4779–4783.
- [4] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” in *9th ISCA Speech Synthesis Workshop*, 2016.
- [5] N. Kalchbrenner, E. Elsen, K. Simonyan, S. Noury, N. Casagrande, E. Lockhart, F. Stimberg, A. van den Oord, S. Dieleman, and K. Kavukcuoglu, “Efficient neural audio synthesis,” *arXiv:1802.08435*, 2018.
- [6] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brebisson, Y. Bengio, and A. Courville, “MelGAN: Generative adversarial networks for conditional waveform synthesis,” in *NeurIPS*, 2019.
- [7] M. Wagner and D. Watson, “Experimental and theoretical advances in prosody: A review,” *Language, Cognition and Neuroscience*, vol. 25, no. 7, pp. 905–945, 2010.
- [8] R. Skerry-Ryan, E. Battenberg, Y. Xiao, Y. Wang, D. Stanton, J. Shor, R. J. Weiss, R. Clark, and R. A. Saurous, “Towards end-to-end prosody transfer for expressive speech synthesis with Tacotron,” in *ICML*, 2018.
- [9] Y. Wang, D. Stanton, Y. Zhang, R. J. Skerry-Ryan, E. Battenberg, J. Shor, Y. Xiao, F. Ren, Y. Jia, and R. A. Saurous, “Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis,” in *ICML*, 2018.
- [10] N. Campbell and P. Mokhtari, “Voice quality: The 4th prosodic dimension,” in *Proc 15th ICPHS*, 2003, pp. 2417–2420.
- [11] A. M. C. Sluijter and V. J. van Heuven, “Spectral balance as an acoustic correlate of linguistic stress,” *The Journal of the Acoustical Society of America*, vol. 100 4 Pt 1, pp. 2471–2485, 1996.
- [12] S. Shechtman and A. Sorin, “Sequence to sequence neural speech synthesis with prosody modification capabilities,” *10th ISCA Speech Synthesis Workshop*, Sep 2019.
- [13] D. Stanton, Y. Wang, and R. Skerry-Ryan, “Predicting expressive speaking style from text in end-to-end speech synthesis,” in *2018 IEEE Spoken Language Technology Workshop (SLT)*, 2018, pp. 595–602.
- [14] V. Wan, C. an Chan, T. Kenter, J. Vit, and R. Clark, “CHiVE: Varying prosody in speech synthesis with a linguistically driven dynamic hierarchical conditional variational network,” in *ICML*, 2019.
- [15] V. Klimkov, S. Ronanki, J. Rohnke, and T. Drugman, “Fine-grained robust prosody transfer for single-speaker neural text-to-speech,” in *Interspeech*, 2019, pp. 4440–4444.
- [16] G. Sun, Y. Zhang, R. J. Weiss, Y. Cao, H. Zen, and Y. Wu, “Fully-hierarchical fine-grained prosody modeling for interpretable speech synthesis,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [17] C. Raffel, M.-T. Luong, P. J. Liu, R. J. Weiss, and D. Eck, “Online and linear-time attention by enforcing monotonic alignments,” in *ICML*, 2017.
- [18] M. He, Y. Deng, and L. He, “Robust sequence-to-sequence acoustic modeling with stepwise monotonic attention for neural TTS,” in *Interspeech*, 2019.
- [19] S. Achanta, A. Antony, L. Golipour, J. Li, T. Raitio, R. Rasipuram, F. Rossi, J. Shi, J. Upadhyay, D. Winarsky, and H. Zhang, “On-device neural speech system,” in *Interspeech*, 2020, submitted.
- [20] “Audio samples from Controllable neural text-to-speech synthesis using intuitive prosodic features,” <https://apple.github.io/neural-tts-with-prosody-control/>, 2020.
- [21] Y. Lee and T. Kim, “Robust and fine-grained prosody control of end-to-end speech synthesis,” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.