

JDI-T: Jointly trained Duration Informed Transformer for Text-To-Speech without Explicit Alignment

Dan Lim¹, Won Jang², Gyeonghwan O², Heayoung Park², Bongwan Kim², Jaesam Yoon²

¹Kakao Corp., Seongnam, Korea

²Kakao Enterprise Corp., Seongnam, Korea

satoshi.2018@kakaocorp.com

Abstract

We propose Jointly trained Duration Informed Transformer (JDI-T), a feed-forward Transformer with a duration predictor jointly trained without explicit alignments in order to generate an acoustic feature sequence from an input text. In this work, inspired by the recent success of the duration informed networks such as FastSpeech and DurIAN, we further simplify its sequential, two-stage training pipeline to a single-stage training. Specifically, we extract the phoneme duration from the autoregressive Transformer on the fly during the joint training instead of pretraining the autoregressive model and using it as a phoneme duration extractor. To our best knowledge, it is the first implementation to jointly train the feed-forward Transformer without relying on a pre-trained phoneme duration extractor in a single training pipeline. We evaluate the effectiveness of the proposed model on the publicly available Korean Single speaker Speech (KSS) dataset compared to the baseline text-to-speech (TTS) models trained by ESPnet-TTS.

Index Terms: text-to-speech, speech synthesis, Transformer, Korean speech

1. Introduction

Deep learning approaches to text-to-speech (TTS) task have made significant progress in generating highly natural speech close to human quality. Especially attention-based encoder-decoder models such as Tacotron [1] and Tacotron2 [2] are dominant in this area. They generate an acoustic feature sequence, mel-spectrogram, for example, from an input text autoregressively using an attention mechanism where the attention mechanism plays the role of implicit aligner between the input sequence and the acoustic feature sequence. Finally, Griffin-Lim algorithm [3] or a neural vocoder such as WaveNet [4], WaveGlow [5] or Parallel WaveGAN [6] is used to convert the predicted acoustic feature sequence to corresponding audio samples.

Besides Recurrent Neural Network (RNN) based TTS models (Tacotron [1], Tacotron2 [2]), Transformer [7] has also been applied for TTS in the attention-based encoder-decoder framework successfully achieving the quality of human recording. The self-attention module, followed by a nonlinear transformation in the Transformer [8], solves the long-range dependency problem by constructing a direct path between any two inputs at different time steps and improves the training efficiency by computing the hidden states in an encoder and a decoder in parallel.

Despite its success in synthesizing high-quality speech, the attention-based encoder-decoder models are prone to the synthesis error, which prevents its commercial use. The unstable attention alignment at synthesis causes the synthesized speech

to be imperfect, e.g., phoneme repeat, skip, or mispronunciation. To solve this problem, duration informed networks such as FastSpeech [9] and DurIAN [10] reduce the errors by relying on a duration predictor instead of the attention mechanism.

The reliance on the duration predictor instead of the attention mechanism is more robust since the duration predictor guarantees stepwise and monotonic alignments between a phoneme sequences and mel-spectrogram. Although the duration informed networks can synthesize high-quality speech without the synthesis error, the training process is tricky: a pre-trained model must be prepared as a phoneme duration extractor since the duration informed networks cannot be trained without a reference phoneme duration sequence.

For example, FastSpeech [9] extracts a phoneme duration sequence from attention alignments matrix of pre-trained autoregressive Transformer for training a feed-forward Transformer and a duration predictor. On the other hand, DurIAN [10] uses the forced alignment, which is commonly used in statistical parametric speech-synthesis systems to train their duration models. Thus, the previous duration informed networks have sequential, two-stage training pipeline, which may slow down the model training time. More recently, duration informed network not requiring pre-trained model [11] has been proposed, but it still requires a multi-stage training phase.

We are motivated by a simple idea that if the attention mechanism of the autoregressive Transformer can provide reliable phoneme duration sequences from the early in joint training, the previous two-stage training pipeline of the duration informed networks could be combined. In this paper, we simplify the training pipeline of the duration informed networks by jointly training the feed-forward Transformer and the duration predictor with the autoregressive Transformer. The contributions of our work are as follow:

- We propose a novel training framework where the feed-forward Transformer and the duration predictor are trained jointly with the autoregressive Transformer. By acquiring reference phoneme duration from the autoregressive Transformer during the training on the fly, the previous two-stage training pipeline of the typical duration informed networks such as FastSpeech [9] or DurIAN [10] is simplified to a single-stage training.
- We remedy an instability of the attention mechanism of the autoregressive Transformer by adding an auxiliary loss and adopt a forward attention mechanism [12]. This makes the phoneme duration sequence extracted from the attention mechanism reliable from the early on in training.
- We prove the effectiveness of the proposed model by comparing it with popular TTS models (Tacotron2, Transformer, and FastSpeech) implemented from

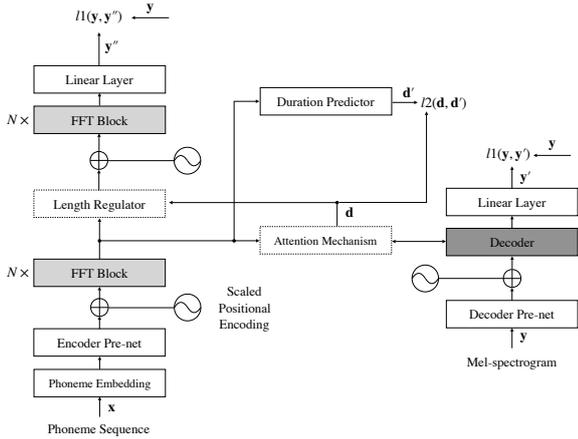


Figure 1: An illustration of our proposed joint training framework (Auxiliary loss for attention is omitted for brevity.)

ESPnet-TTS [13] on the publicly available Korean dataset.

2. Model description

The main idea of JDI-T is to train feed-forward Transformer and duration predictor with autoregressive Transformer jointly. In this section, we describe each component in detail.

2.1. Feed-Forward Transformer

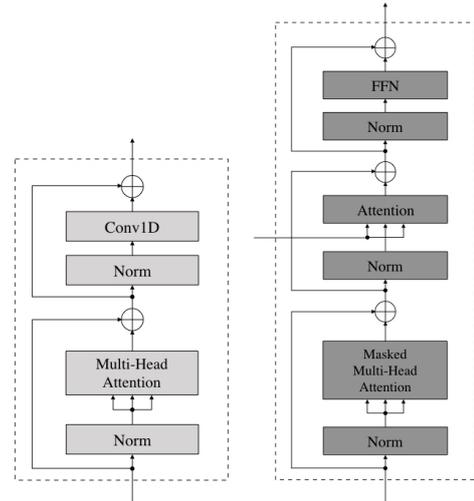
Feed-forward Transformer located on the left in Figure 1 consists of a phoneme embedding, an encoder pre-net, scaled positional encodings, multiple Feed-Forward Transformer (FFT) blocks, a length regulator and a linear layer for phoneme sequence to mel-spectrogram transformation.

As shown in Figure 2a, the structure of FFT blocks is composed of a multi-head attention and a single layer 1D convolution network where residual connections, layer normalization, and dropout are used. It is slightly different from what described in [9]. Note that the output of each FFT block on both sides is normalized by layer normalization. Encoder pre-net and scaled positional encoding have the same configuration as described in Transformer TTS [7].

The stacked modules from the phoneme embedding to the FFT blocks below the length regulator works as the encoder, which is also shared by autoregressive Transformer and duration predictor. The length regulator regulates an alignment between the phoneme sequences and the mel-spectrogram in the same way described in FastSpeech [9], expanding the output sequences of FFT blocks on phoneme side according to reference phoneme duration so that total length of it matches the total length of mel-spectrogram. The feed-forward Transformer is trained to minimize $l1$ loss between predicted and reference mel-spectrogram.

2.2. Autoregressive Transformer

Autoregressive Transformer is an attention-based encoder-decoder model as shown in Figure 1 where it shares the encoder with the feed-forward Transformer on the left bottom and the output of the encoder is attended by the stacked modules on the right consisting of a decoder pre-net, a scaled positional encoding, a decoder, and a linear layer. The decoder depicted in



(a) FFT Block

(b) Decoder

Figure 2: (a) The feed-forward Transformer block. (b) The decoder for autoregressive Transformer

Figure 2b has the similar structure with the FFT block in Figure 2a but there are differences of having masked multi-head attention instead of multi-head attention, additional sub-layer in the middle for attention mechanism over the outputs of the encoder and position-wise Feed Forward Network (FFN) instead of 1D convolution network. Note that the output of the decoder is normalized by layer normalization. The FFN of the decoder and the decoder pre-net have the same configuration as described in Transformer TTS [7].

The autoregressive Transformer provides reference phoneme duration during training for each phoneme on the fly. Specifically, while training the autoregressive Transformer to minimize $l1$ loss between predicted and reference mel-spectrogram, phoneme duration is extracted from the attention alignment matrix by counting the number of mel frames which scores the highest value for each distinct phoneme. The phoneme duration acquired in this way is used as reference phoneme duration for the length regulator of feed-forward Transformer and the duration predictor. Note that the decoder, unlike described in Transformer TTS [7], neither has been stacked nor has multi-head attention over the outputs of the encoder since it was sufficient as a phoneme duration extractor.

2.3. Duration Predictor

Duration predictor consists of a 2-layer 1D convolution network with ReLU activation followed by the layer normalization, dropout layer and an extra linear layer to output a scalar value which is interpreted as phoneme duration in the logarithmic domain. The depiction of duration predictor is omitted since it is exactly same as described in FastSpeech [9]. As shown in the middle top of Figure 1, the duration predictor also shares the encoder with feed-forward, autoregressive Transformer so it is stacked on top of the FFT blocks on the phoneme side. In order to minimize the difference between predicted and reference phoneme duration, $l2$ loss is used, where the reference phoneme duration is extracted from an attention alignment matrix of autoregressive Transformer during training. Note that we stop gradients propagation from the duration predictor to the

FFT blocks during joint training, since the training fails without it.

3. Auxiliary loss and Attention mechanism

It is very important to stabilize the attention alignments of autoregressive Transformer, since the training of the proposed model requires precise phoneme duration and the quality of phoneme duration is affected by the stability of attention alignment of autoregressive Transformer. In this section, we describe some methods for acquiring reliable phoneme duration by stabilizing attention alignments of autoregressive Transformer from the early in training steps.

3.1. CTC recognizer loss

In [14], authors insist that the synthesis error rate of attention based TTS models can be alleviated if the model is guided to learn the dependency between the input text and the predicted acoustic feature sequences by maximizing the mutual information between them. After formulating that maximizing the mutual information is equivalent to training an auxiliary recognizer, they show that training CTC recognizer as auxiliary loss for the Tacotron TTS model can reduce the synthesis error.

Although the method of maximizing the mutual information is proposed for reducing the synthesis error of the attention-based model, we found that it also helps autoregressive Transformers learn stable attention alignments during training, which is a crucial factor for the proposed model to be jointly trained successfully. In this work, we implement it just by adding an extra CTC loss layer. Specifically, the extra linear layer is stacked on top of the decoder of the autoregressive Transformer so that it is trained to predict input phoneme sequences.

3.2. Forward attention with Guided attention loss

As reported in [15], location-relative attention mechanisms are more preferred over the content-based attention mechanisms in that they not only reduce synthesis error but also help the model align quickly during training. However, it is not proper for parallel computing. If it is used in the autoregressive Transformer such as Transformer TTS [7], the computation efficiency is sacrificed. Therefore we adopt forward attention mechanism [12] on top of content-based attention mechanism for minimal computational overhead instead of location-relative attention mechanisms. The forward attention mechanism guarantees fast convergence and stability of the attention alignments by considering only monotonic attention alignment paths, which is a natural assumption between the input phoneme sequences and the output acoustic feature sequences.

Algorithm 1 states the forward attention mechanism adopted for the proposed model. The algorithm has same procedures as in original paper [12] except that it prepares the attention weights $w_{1:T}(1:N)$ for all time steps N, T in parallel using the attention mechanism *Attention*. The *Attention* is a single-head, content-based attention from the decoder of autoregressive Transformer using an output sequences of the FFT blocks on phoneme side $h_{1:N}$ and an output sequences of the decoder $s_{1:T}$ which has the length of N and T respectively.

Although the forward attention mechanism guarantees monotonic attention alignments in principle, it is observed that the training often fails without learning any valid alignments. We found that guiding attention alignments with auxiliary loss is useful in solving the problem [16, 17]. In the proposed model, the guided attention loss with the same configuration as de-

Algorithm 1 Forward Attention

Initialize:

$$\hat{\alpha}_0(1) \leftarrow 1$$

$$\hat{\alpha}_0(n) \leftarrow 0, n = 2, \dots, N$$

$$w_{1:T}(1:N) \leftarrow \text{Attention}(h_{1:N}, s_{1:T})$$

for $t = 1$ to T **do**

$$\hat{\alpha}'_t(n) \leftarrow (\hat{\alpha}_{t-1}(n) + \hat{\alpha}_{t-1}(n-1))w_t(n)$$

$$\hat{\alpha}_t(n) \leftarrow \hat{\alpha}'_t(n) / \sum_{m=1}^N \hat{\alpha}'_t(m)$$

$$c_t \leftarrow \sum_{n=1}^N \hat{\alpha}_t(n)h_n$$

end for

scribed in [16] is added as auxiliary loss term for training the model since it keeps a single joint training property without depending on external modules. It applies constraint on attention alignments matrix in the form of the diagonal mask based on the idea that the input phoneme sequences and the output acoustic feature sequences have nearly diagonal correspondence.

Consequently, the loss function of the proposed model consists of two $l1$ losses for the mel-spectrogram, a $l2$ loss for phoneme duration, and two auxiliary losses, which are CTC recognizer loss and Guided Attention (GA) loss. It can be formulated as follows:

$$L = \|\mathbf{y} - \mathbf{y}'\| + \|\mathbf{y} - \mathbf{y}''\| + \|\mathbf{d} - \mathbf{d}'\|_2 + \text{CTCLoss} + \text{GALoss} \quad (1)$$

where \mathbf{y} , \mathbf{y}' and \mathbf{y}'' are the mel-spectrograms from the reference, autoregressive Transformer and feed-forward Transformer respectively, \mathbf{d} is a phoneme duration sequences from the attention mechanism of autoregressive Transformer and \mathbf{d}' is a phoneme duration sequences from the duration predictor. It can be future works to investigate an effect of different scaling constants for each loss term.

After joint training, the autoregressive Transformer is no longer needed, so only the feed-forward Transformer and the duration predictor are used for synthesis. The feed-forward Transformer generates a mel-spectrogram \mathbf{y}'' from the input phoneme sequences in parallel using the phoneme duration sequence predicted by the duration predictor.

4. Experiments

4.1. Datasets

We conduct experiments on two different datasets. The first dataset is an internal speech recorded by a professional Korean female speaker in studio quality. The number of utterances used for training is 25k, among which 250, 50 samples were reserved for validation and testing, respectively. The second dataset is Korean Single speaker Speech (KSS) Dataset [18], which is publicly available for the Korean text-to-speech task. It consists of 12,853 utterance audio files recorded by a professional female voice actress and transcription extracted from their books with a total audio length of approximately 12 hours. We reserved the last 250 utterances; 200 samples for validation and 50 samples for testing.

We convert the input text to phoneme sequences using an internal text processing tool. The acoustic feature is 80-dimensional mel-spectrogram extracted from audio with sampling rate of 22,050 Hz using the Librosa library [19]. FFT size and hop size are 1024, 256 respectively. Finally, the mel-spectrogram is normalized so that every element of the feature vector has zero mean and unit variance over the training set.

4.2. Model configuration

The proposed model has 6 FFT blocks in the feed-forward Transformer both on the phoneme side and the mel-spectrogram side. The number of head for self-attention is set to 8, and the kernel size of 1D convolution is set to 5. The inner-layer dimension of the position-wise feed-forward network is set to 2048, and all other dimensions not stated explicitly have been set to 512. The proposed model is implemented by PyTorch [20] neural network library and trained on 4 NVIDIA V100 GPUs with a batch size of 16 per each GPU. We optimize it using the RAdam algorithm [21] with the same learning rate schedule as in [8] about 300k training steps.

For comparative experiments, we train another three models using ESPnet-TTS [13] which is open-source speech processing toolkit supporting state-of-the-art end-to-end TTS models. The models are two autoregressive, attention-based models: Tacotron2.v3 and Transformer.v1 and a non-autoregressive, duration informed model FastSpeech.v2 following the configuration in the recipe of the toolkit. Note that the FastSpeech.v2 uses pre-trained Transformer.v1 as teacher model.

We use Parallel WaveGAN [6] as the vocoder for transforming the generated mel-spectrogram to audio samples in all experiments.

4.3. Evaluation

To evaluate the effectiveness of the proposed model, we conduct the Mean Opinion Score (MOS) test ¹. The proposed model, JDI-T, is compared with three different models, including Tacotron2, Transformer, and FastSpeech. The audio samples for the MOS test are generated using scripts of the test samples reserved for each dataset, and thirteen native Korean speakers listen to all of it for measuring the audio quality. Table 1 shows the results on two different datasets; the Internal and the KSS. Note that the results of *GT mel* are the evaluation of the audio samples converted from reference mel-spectrogram by the vocoder; thus, it indicates upper bounds that our TTS models can achieve.

The results of the MOS test on both datasets show that the score of the duration informed model (FastSpeech) is lower than the attention-based model (Tacotron2, Transformer). Therefore it can be said that duration informed model is challenging to match the audio quality of its teacher model, especially by comparing the FastSpeech with its teacher model; Transformer. In this case, a more elaborate training technique, focus rate tuning, or sequence-level knowledge distillation, as described in [9], would be required to improve the audio quality by its teacher model. On the other hand, the score of our proposed model, which is also non-autoregressive and duration informed model like FastSpeech, is better than FastSpeech and even achieves the highest score among the TTS models in the Internal dataset. These results show that the joint training of the proposed model is beneficial for improving the audio quality as well as for simplifying the training pipeline.

In contrast, the score of the proposed model is lower than Transformer in the KSS. It may be derived from the fact that the quality of audio files is poorer (i.e., the pronunciation is unclear) in KSS dataset than in the Internal dataset which is recorded in studio quality with commercial use in mind. So, it seems that the proposed model has the difficulty of learning precise attention alignments with relatively poor sound quality of audio

¹Audio samples are available at the following URL: <https://imdanboy.github.io/interspeech2020>

Table 1: Mean opinion scores (5-point scale)

Model	Internal	KSS
<i>GT mel</i>	3.92	3.87
<i>Tacotron2</i>	3.52	3.33
<i>Transformer</i>	3.55	3.72
<i>FastSpeech</i>	3.48	3.23
<i>JDI-T (ours)</i>	3.77	3.52

files. It can be future works to make the model learn better alignments during joint training even in the slightly poor sound quality of audio files.

In addition to its high-quality speech synthesis, the proposed model has benefits of the robustness and fast speed at synthesis over the autoregressive, attention-based TTS models since it has the feed-forward structure and does not rely on an attention mechanism as in FastSpeech [9]. Moreover, our internal test shows that Tacotron2 and Transformer have a high rate of synthesis error, especially when they are trained with the KSS dataset and synthesize the out-of-domain scripts. Note that the synthesized audio samples from the test scripts have no synthesis error.

5. Conclusion

In this paper, we propose Jointly trained Duration Informed Transformer (JDI-T) for TTS. The proposed model, consisting of the feed-forward Transformer, the duration predictor, and the autoregressive Transformer, is trained jointly without explicit alignments. After joint training, only the feed-forward Transformer with the duration predictor is used for fast and robust conversion from phoneme sequences to mel-spectrogram. Experimental results on publicly available Korean datasets prove the effectiveness of the proposed model by showing that it can synthesize high-quality speech. Furthermore, it achieves state-of-the-art performance in the internal studio-quality dataset compared to other popular TTS models implemented from ESPnet-TTS.

6. References

- [1] 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. Agiomvrgiannakis, R. Clark, and R. A. Saurous, "Tacotron: Towards End-to-End Speech Synthesis," in *Proceedings of the Annual Conference of the International Speech Communication Association (INTERSPEECH)*, 2017, pp. 4006–4010.
- [2] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan, R. A. Saurous, Y. Agiomvrgiannakis, and Y. Wu, "Natural TTS Synthesis by Conditioning Wavenet on MEL Spectrogram Predictions," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4779–4783.
- [3] D. Griffin and Jae Lim, "Signal estimation from modified short-time Fourier transform," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, vol. 8, 1983, pp. 804–807.
- [4] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, "WaveNet: A Generative Model for Raw Audio," *ArXiv*, vol. abs/1609.03499, 2016.
- [5] R. Prenger, R. Valle, and B. Catanzaro, "WaveGlow: A Flow-based Generative Network for Speech Synthesis," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 3617–3621.

- [6] R. Yamamoto, E. Song, and J. Kim, "Parallel WaveGAN: A Fast Waveform Generation Model Based on Generative Adversarial Networks with Multi-Resolution Spectrogram," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6199–6203.
- [7] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, "Neural Speech Synthesis with Transformer Network," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 6706–6713.
- [8] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is All you Need," in *Advances in Neural Information Processing Systems 30 (NIPS)*, 2017, pp. 5998–6008.
- [9] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, "FastSpeech: Fast, Robust and Controllable Text to Speech," in *Advances in Neural Information Processing Systems 32 (NIPS)*, 2019, pp. 3171–3180.
- [10] C. Yu, H. Lu, N. Hu, M. Yu, C. Weng, K. Xu, P. Liu, D. Tuo, S. Kang, G. Lei, D. Su, and D. Yu, "DurIAN: Duration Informed Attention Network For Multimodal Synthesis," *ArXiv*, vol. abs/1909.01700, 2019.
- [11] Z. Zeng, J. Wang, N. Cheng, T. Xia, and J. Xiao, "AlignTTS: Efficient Feed-Forward Text-to-Speech System Without Explicit Alignment," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6714–6718.
- [12] J. Zhang, Z. Ling, and L. Dai, "Forward Attention in Sequence-To-Sequence Acoustic Modeling for Speech Synthesis," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4789–4793.
- [13] T. Hayashi, R. Yamamoto, K. Inoue, T. Yoshimura, S. Watanabe, T. Toda, K. Takeda, Y. Zhang, and X. Tan, "Espnet-TTS: Unified, Reproducible, and Integratable Open Source End-to-End Text-to-Speech Toolkit," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7654–7658.
- [14] P. Liu, X. Wu, S. Kang, G. Li, D. Su, and D. Yu, "Maximizing Mutual Information for Tacotron," *ArXiv*, vol. abs/1909.01145, 2019.
- [15] E. Battenberg, R. Skerry-Ryan, S. Mariooryad, D. Stanton, D. Kao, M. Shannon, and T. Bagby, "Location-Relative Attention Mechanisms for Robust Long-Form Speech Synthesis," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6194–6198.
- [16] H. Tachibana, K. Uenoyama, and S. Aihara, "Efficiently Trainable Text-to-Speech System Based on Deep Convolutional Networks with Guided Attention," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4784–4788.
- [17] X. Zhu, Y. Zhang, S. Yang, L. Xue, and L. Xie, "Pre-Alignment Guided Attention for Improving Training Efficiency and Model Stability in End-to-End Speech Synthesis," *IEEE Access*, vol. 7, pp. 65 955–65 964, 2019.
- [18] K. Park, "KSS Dataset: Korean Single speaker Speech Dataset," <https://kaggle.com/bryanpark/korean-single-speaker-speech-dataset>, 2018.
- [19] B. McFee, C. Raffel, D. Liang, D. P. Ellis, M. McVicar, E. Battenberg, and O. Nieto, "librosa: Audio and Music Signal Analysis in Python," in *Proceedings of the 14th Python in Science Conference*, 2015, pp. 18 – 24.
- [20] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An Imperative Style, High-Performance Deep Learning Library," in *Advances in Neural Information Processing Systems 32 (NIPS)*, 2019, pp. 8024–8035.
- [21] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han, "On the Variance of the Adaptive Learning Rate and Beyond," in *Proceedings of the Eighth International Conference on Learning Representations (ICLR)*, 2020.