



Multi-speaker Emotional Text-to-speech Synthesizer

Sungjae Cho^{1,†}, Soo-Young Lee²

¹Korea Institute of Science and Technology, Republic of Korea

²Korea Advanced Institute of Science and Technology, Republic of Korea

sj.cho@snu.ac.kr, sylee@kaist.ac.kr

Abstract

We present a methodology to train our multi-speaker emotional text-to-speech synthesizer that can express speech for 10 speakers' 7 different emotions. All silences from audio samples are removed prior to learning. This results in fast learning by our model. Curriculum learning is applied to train our model efficiently. Our model is first trained with a large single-speaker neutral dataset, and then trained with neutral speech from all speakers. Finally, our model is trained using datasets of emotional speech from all speakers. In each stage, training samples of each speaker-emotion pair have equal probability to appear in mini-batches. Through this procedure, our model can synthesize speech for all targeted speakers and emotions. Our synthesized audio sets are available on our web page.

Index Terms: emotional speech synthesis, text-to-speech, machine learning, neural network, deep learning

1. Introduction

Emotional speech synthesis has been achieved through deep neural networks [1, 2, 3, 4]. However, most studies have trained models on a small number of speakers or balanced class distributions because it is challenging to guarantee speech quality for each speaker and emotion, given imbalanced data distributions with respect to speakers and emotions. In this paper, we present a methodology for training our multi-speaker emotional text-to-speech (TTS) synthesizer capable of generating speech for all targeted speakers' voices and emotions. The main methods are silence removal, curriculum learning [5], and oversampling [6]. The synthesized audios are demonstrated through a web page.

2. Datasets

4 datasets were used to train the multi-speaker emotional TTS synthesizer. The first dataset, the Korean single speaker speech (KSS) dataset [7], is publicly available and contains speech samples of a single female speaker: kss-f. We labeled their emotion as neutral. The remaining 3 datasets consist of speech of the Ekman's 7 basic emotions [8]: neutral, anger, disgust, fear, happiness, sadness, and surprise.

The first Korean emotional TTS (KETTS) dataset consists of 1 female and 1 male speaker: ketts-30f and ketts-30m, which are abbreviations of a 30's female and male in KETTS. The 2 speakers were assigned to different sets of sentences; however, the same sentences were recorded across 7 emotions for a single speaker. In the female case, only happy speech samples have a different set of sentences. KETTS is balanced with respect to speakers and emotions, except for the female's happy speech subset (Table 1).

The second Korean emotional TTS (KETTS2) dataset consists of 3 female and 3 male speakers, totally 6 speakers: kett2-

[†] work done at KAIST

Table 1: Hours of preprocessed training datasets

| Speaker | all | neu | ang | dis | fea | hap | sad | sur |
|------------|--------|-------|-------|-------|-------|-------|-------|-------|
| kss-f | 12.59 | 12.59 | | | | | | |
| ketts-30f | 26.61 | 3.52 | 3.46 | 3.51 | 3.68 | 5.13 | 3.75 | 3.56 |
| ketts-30m | 24.12 | 3.37 | 3.29 | 3.31 | 3.51 | 3.50 | 3.73 | 3.40 |
| ketts2-20m | 5.09 | 0.72 | 0.72 | 0.74 | 0.76 | 0.69 | 0.75 | 0.70 |
| ketts2-30f | 4.69 | 0.66 | 0.65 | 0.67 | 0.65 | 0.70 | 0.68 | 0.68 |
| ketts2-40m | 4.98 | 0.73 | 0.69 | 0.70 | 0.75 | 0.69 | 0.74 | 0.69 |
| ketts2-50f | 4.98 | 0.73 | 0.71 | 0.71 | 0.70 | 0.72 | 0.71 | 0.69 |
| ketts2-50m | 4.73 | 0.68 | 0.68 | 0.69 | 0.67 | 0.68 | 0.68 | 0.65 |
| ketts2-60f | 4.90 | 0.77 | 0.68 | 0.67 | 0.68 | 0.72 | 0.72 | 0.67 |
| ketts3-f | 9.64 | 3.96 | 1.34 | | 1.27 | 1.44 | 1.64 | |
| ketts3-m | 9.38 | 3.90 | 1.43 | | 1.18 | 1.39 | 1.48 | |
| all | 111.70 | 31.63 | 13.65 | 11.01 | 13.85 | 15.64 | 14.87 | 11.05 |

20m, ketts2-30f, ketts2-40m, ketts-50f, ketts2-50m, and ketts2-60f. The same sentences were recorded across 7 emotions and 6 speakers. Hence, KETTS2 is balanced with respect to speakers and emotions (Table 1).

The third Korean emotional TTS (KETTS3) dataset consists of 1 female and 1 male speaker: ketts3-f and ketts3-m. It includes 5 emotions, excluding disgust and surprise. The same sentences were recorded across 2 speakers; however, different sentences were spoken for the 5 emotions. KETTS3 is balanced for speakers but not for emotions. Therefore, the whole training dataset is balanced for neither speakers nor emotions (Table 1).

3. Methodology

3.1. Preprocessing

The WebRTC voice activity detector, py-webrtcvad¹, is utilized to remove unvoiced segments in audios, with its settings of an aggressiveness level of 3, frame duration 30ms, and padding duration 150ms. These settings remove silences at the start, end, and middle of speech. However, the amount of silence removed does not distort emotional expression. All audios are resampled to sampling rate 22,050Hz. Mel spectrograms are computed through a short-time Fourier transform (STFT) using frame size 1024, hop size 256, window size 1024, and a Hann window function. The STFT magnitudes are transformed to the librosa Slaney mel scale using an 80-channel mel filterbank spanning 0Hz to 8kHz, and the results are then clipped to a minimum value of 10^{-5} , followed by log dynamic range compression.

Every Korean character in an input sentence is decomposed into 3 elements: an onset, nucleus, and coda. In total, 19 onsets, 21 nuclei, and 28 codas including the empty coda are employed as defined by Unicode. A sequence of these elements becomes a grapheme sequence taken as input by our synthesizer.

¹<https://github.com/wiseman/py-webrtcvad>

3.2. Model

Our multi-speaker emotional TTS synthesizer takes 3 inputs — the grapheme sequence of a Korean sentence, 1 of 10 speakers (5 females, 5 males), and 1 of the 7 Ekman’s emotion classes. It then generates a waveform in which the speaker utters the input sentence with the given emotion. Our synthesizer consists of 2 sub-models: *Tacotron 2* [9], mapping a grapheme sequence to a mel spectrogram, and *WaveGlow* [10], transforming the mel spectrogram to a waveform. Tacotron 2 is an auto-regressive sequence-to-sequence neural network with a location-sensitive attention mechanism. WaveGlow is a flow-based generative neural network without auto-regression. We adapted NVIDIA Tacotron 2 and WaveGlow repositories^{2,3} to synthesize speech for multiple speakers and emotions. The WaveGlow model was utilized without modification but the Tacotron 2 model was modified as outlined in the following paragraph.

Speaker identity is represented as a 5-dimensional trainable *speaker vector*. Emotion identity is represented as a 3-dimensional trainable *emotion vector*, except for the neutral emotion vector, which is a non-trainable zero vector. To synthesize speech of a given speaker and emotion, in the decoder of Tacotron 2, speaker and emotion vectors are concatenated to attention context vectors taken by the first and second LSTM layers and the linear layer estimating a mel spectrogram.

3.3. Training

Tacotron 2 was trained with a batch size of 64 equally distributed to 4 GPUs. The Adam optimizer [11] of the default settings ($\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-6}$) was used with a learning rate of 10^{-3} and L_2 regularization with weight 10^{-6} . If the norm of gradients exceeded 1, their norm was normalized to 1 to ensure stable learning.

Using a curriculum learning [5] strategy, Tacotron 2 was trained to learn single-speaker neutral speech, multi-speaker neutral speech, and multi-speaker emotional speech in this order. More specifically, the model was trained with the KSS dataset for 20,000 iterations, then additionally with all datasets of neutral speech for 30,000 iterations, and finally with all training datasets for 65,000 iterations. Transitioning to training on the next dataset was done when the model stably pronounced given whole sentences for all training speaker-emotion pairs.

In each training stage, we oversampled [6] the training set with respect to speaker-emotion pairs, which means samples of each speaker-emotion pair appear in a mini-batch with equal probability. For example, samples of (ketts-30f, neutral) and those of (ketts2-20m, happy) appear in a mini-batch with equal probability. This helped overcome difficulty in learning to synthesize speech of speaker-emotion pairs with relatively scarce samples.

WaveGlow was trained with a batch size of 24, equally distributed to 3 GPUs using 24 clips of 16,000 mel spectrogram frames randomly chosen from each training sample. Training samples shorter than 16,000 mel frames were excluded from the training set since these samples padded with zeros caused unstable learning such as exploding gradients. Similar to Tacotron 2, we oversampled the training set with respect to speaker-emotion pairs. The Adam optimizer was used with the default settings and learning rate 10^{-4} . Weight normalization was applied, as described in the original paper [10]. To ensure stable learning, if the norm of gradients exceeded 1, their norm was normalized

²<https://github.com/NVIDIA/tacotron2>

³<https://github.com/NVIDIA/waveglow>

to 1. The model was initialized with the pretrained weights⁴ offered in the WaveGlow repository. The network was trained for 400,000 iterations until its loss curve formed a plateau. The z elements were sampled from Gaussians with standard deviation 1 during training and 0.75 during inference.

4. Results and Discussion

Through this procedure, our speech synthesizer is able to synthesize speech for all available 10 speakers and 7 emotions. Unexpectedly, disgusted and surprised expressions of the KETTS3 speakers can be synthesized even without training supervision. Synthesized speech samples can be found on this web page⁵.

Although our model expresses speaker and emotion identities, there are some minor inconsistencies in the quality of synthesized samples across speakers and emotions. Thus, in production, it is reasonable to fine-tune for each speaker and respectively preserve the model parameters.

Our silence removal settings substantially accelerated the learning of Tacotron 2. This was probably because silence removal at the start, end, and middle of speech resulted in the linear relationship between text and speech, and this relationship helped the location-sensitive attention network easily learn text-to-speech alignments.

5. Acknowledgements

This work was supported by *Ministry of Culture, Sports and Tourism* and *Korea Creative Content Agency* [R2019020013, R2020040298].

6. References

- [1] Y. Lee, A. Rabiee, and S.-Y. Lee, “Emotional end-to-end neural speech synthesizer,” *ArXiv*, vol. abs/1711.05447, 2017.
- [2] H. Choi, S. Park, J. Park, and M. Hahn, “Multi-speaker emotional acoustic modeling for CNN-based speech synthesis,” in *ICASSP*, 2019.
- [3] S.-Y. Um, S. Oh, K. Byun, I. Jang, C. Ahn, and H.-G. Kang, “Emotional speech synthesis with rich and granularized control,” in *ICASSP*, 2020.
- [4] T.-H. Kim, S. Cho, S. Choi, S. Park, and S.-Y. Lee, “Emotional voice conversion using multitask learning with text-to-speech,” in *ICASSP*, 2020.
- [5] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, “Curriculum learning,” in *ICML*, 2009.
- [6] M. Buda, A. Maki, and M. A. Mazurowski, “A systematic study of the class imbalance problem in convolutional neural networks,” *Neural Networks*, vol. 106, 2018.
- [7] K. Park, “KSS dataset: Korean single speaker speech dataset,” <https://kaggle.com/bryanpark/korean-single-speaker-speech-dataset>, 2018.
- [8] P. Ekman and D. Cordaro, “What is meant by calling emotions basic,” *Emotion Review*, vol. 3, no. 4, 2011.
- [9] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R.-S. Ryan, R. A. Saurous, Y. Agiomyrgiannakis, and Y. Wu, “Natural TTS synthesis by conditioning wavenet on MEL spectrogram predictions,” in *ICASSP*, 2018.
- [10] R. Prenger, R. Valle, and B. Catanzaro, “Waveglow: A flow-based generative network for speech synthesis,” in *ICASSP*, 2019.
- [11] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *ICLR*, 2015.

⁴“waveglow_256channels_universal_v5.pt” was used.

⁵https://sungjae-cho.github.io/InterSpeech2021_STDemo/