



CSS10: A Collection of Single Speaker Speech Datasets for 10 Languages

Kyubyong Park^{*†}, Thomas Mulc^{^†}

^{*}Kakao Brain, South Korea

[^]Expedia Group, USA

kyubyong.park@kakaobrain.com, tmulc@expedia.com

Abstract

We describe our development of CSS10, a collection of single speaker speech datasets for ten languages. It is composed of short audio clips from LibriVox audiobooks and their aligned texts. To validate its quality we train two neural text-to-speech models on each dataset. Subsequently, we conduct Mean Opinion Score tests on the synthesized speech samples. We make our datasets, pre-trained models, and test resources publicly available. We hope they will be used for future speech tasks.

Index Terms: text-to-speech, TTS, multi-lingual datasets, speech datasets, deep learning

1. Introduction

Text-to-speech (TTS) models aim at generating an audio sequence given some input text. Recently there have been many TTS models based on neural networks. Among them are WaveNet [1], Tacotron 1 & 2 [2][3], Char2Wav [4], DeepVoice 1, 2, & 3 [5][6][7], DCTTS [8], VoiceLoop [9], etc. All of them adopted an end-to-end approach, taking as input a sequence of characters or phonemes and returning the raw waveform or spectrogram.

All of those models have been introduced with results manufactured on their internal datasets except DCTTS, which uses the public LJ Speech Dataset [10]. This status quo has made reproducing results in the papers difficult for researchers outside. People have resorted to using less than ideal data and/or creating their own datasets that try to match some of the properties of the internal ones. What is even worse is that it is hard to compare different TTS models, because there is no benchmark dataset.

Furthermore, the research is mostly focused on English. Except Deep Voice 2 & 3, which trained on the Chinese language, all the other models attempted to model the text-to-audio mapping for the English language. The most likely explanation is simply the dearth of freely available non-English datasets. We strongly believe that languages other than English deserve to have attention from researchers as well. This is why we construct multi-lingual speech datasets and share them with the research community.

Our contribution is two-fold:

- We construct single speaker speech datasets with aligned text for ten different languages.
- We train two famous TTS models on each dataset and evaluate them with Mean Opinion Scores (MOS).

All of the resources mentioned above are available in our GitHub repository¹.

[†] equal contribution

¹ <https://github.com/Kyubyong/CSS10>

2. Related work

Not surprisingly, for the English language, there are many public speech datasets such as Blizzard [11], VCTK [12], LibriSpeech [13], TED-LIUM [14], VoxForge [15], and Common Voice [16]. While these datasets are large—some having more than 100 hours of audio—they are all from multiple speakers, which make them unideal for the single speaker TTS task (namely, generating speech from text in the voice of a single speaker) where more data from a single speaker tends to help model performance [17]. One popular, newly-created dataset is the LJ Speech dataset [10]. It consists of audio files segmented from audiobooks in LibriVox recorded by a female volunteer. It has a large cumulative audio length (20+ hrs) and has been verified to work with neural TTS models. Another is the World English Bible (WEB) dataset [18], which was sourced from bible recordings and text in the public domain. It shares the LJ Speech dataset's properties but is sampled at a relatively low rate of 12 kHz. Both are publicly available.

For the Japanese language, there is a public dataset called the JSUT dataset [19]. It is a single speaker dataset designed for speech synthesis that includes around 10 hours of utterances with aligned text. One important distinction between the JSUT dataset and the LJ Speech and WEB datasets is that the JSUT dataset's recording was carried out in a controlled environment with specially designed scripts. One drawback of the JSUT dataset is that it lacks a phonetic transcription of text. Additionally, there were no follow-up experiments on the dataset.

For the German language, there is the Pavoque dataset [20]—a single speaker, multi-style corpus of German speech. It has 12+ hours of audio clips each of which is associated with phoneme-level annotations. However, it is more suitable for speech style tasks rather than regular TTS tasks.

A dataset closer to our work is the Spoken WP Corpus Collection [21]. It has hundreds of hours of audio with aligned text for three languages: Dutch, German, and English. For each language, much of the audio comes from a single speaker. However, the audio clips are generally large—on the order of minutes—which makes using the dataset difficult for neural TTS models [22].

Perhaps the datasets closest to our work are the Tundra [22] and M-AILABS [23] datasets, which have 14 and 9 languages, respectively, and are built from audiobooks. The Tundra dataset uses a single speaker for each language, but does so by using only one audiobook per language. The M-AILABS dataset, which focuses on European languages, has nearly one-thousand hours of total audio but uses multiple speakers in each language.

3. Datasets

We choose to use audiobooks from LibriVox [24], a website for free public domain audiobooks, as the source of audio data for three major reasons. First, audiobooks are inherently ac-

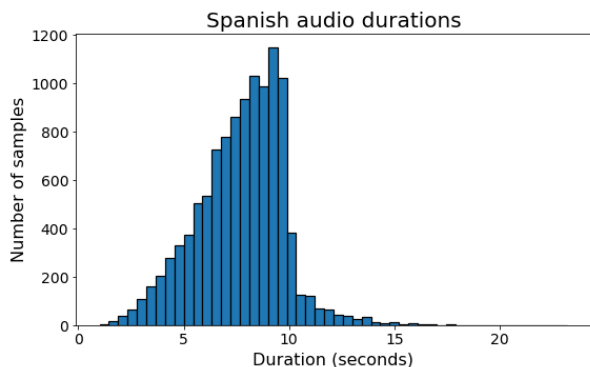


Figure 1: Histogram of audio durations for the Spanish dataset.

accompanied by text, which is essential for our purpose. Second, although many readings are performances and hence use abnormal intonation, readers tend to speak with a regular, constant speed. Third, many audiobooks are performed by a single speaker.

3.1. Selection of audiobooks

As of now, there are audiobooks for 95 languages in LibriVox. We examine how many hours of solo recordings each language has; many audiobook performers record multiple audiobooks which makes finding large quantities of single speaker data achievable. We exclude a language if it has less than 4 hours of solo audio recordings, because we are not confident that training will succeed². Then, we check for text availability and for audio quality. If text is not available or if the audio includes a noticeable amount of noise, the audiobook is excluded.

This process yields audiobooks for the following languages: Chinese (zh), Dutch (nl), French (fr), Finnish (fi), German (de), Greek (el), Hungarian (hu), Japanese (ja), Russian (ru), and Spanish (es)³. All audio files are sampled at 22 kHz.

3.2. Audio processing

The audio from LibriVox usually comes in large files with lengthy audio clips that do not suit the TTS task, so we fragment them into many small files. We use the audio editor Audacity [25] to programmatically find split-points of the audio anytime there is more than a 0.5-second duration of silence, except for Spanish audiobooks, where we use a 0.25-second duration. Next, we adjust the points such that neighboring clips are joined to have a duration around 10 seconds. We found these tricks improve computational efficiency. The distribution of audio lengths for the Spanish dataset is shown in Figure 1. We see about 85% of the samples have a duration between 5 and 11 seconds. All other languages have distributions similar to this.

3.3. Text processing

We have experts align the text with each segmented audio clip to create an <audio, text> pair. At first we considered using a forced aligner such as Gentle [26]. However, we forfeited the idea after we realized that it does not guarantee the correct

²This number was chosen based on our prior successes of training DCTTS on an audiobook with 4 hours of data.

³The parenthetical expression next to each language is the language code from ISO 639-1.

alignments and that it is language dependent.

3.3.1. Text normalization

Once we secure the <audio, text> pairs, we request that our experts normalize text; all abbreviations are expanded (e.g., Dra. → Doctora) and Arabic numerals are spelled out to match the context. Unlike Deep Voice 3 or DCTTS, which ignore case, we decide to retain case, because it can be a cue for new sentences or proper nouns or nouns. In addition, we remove infrequent symbols other than punctuations such as the period, question mark, exclamation point, colon, comma, and semi-colon.

3.3.2. Phonetic transcription

Because we take text as input, it is important to understand the writing system of our target languages. Dutch, German, Finnish, French, Hungarian, and Spanish are based on Latin alphabets, while Greek and Russian use Greek letters and Cyrillic letters, respectively. Chinese uses Chinese characters, and Japanese employs three different scripts (i.e., Hiragana, Katakana, and Chinese characters). All the writing systems except Chinese characters are phonetic. In other words, they are written as they are pronounced. On the other hand, Chinese characters are ideographic, so they are not directly associated with pronunciations. Thus, Chinese people typically use the Romanization system for pronunciation, which is called pinyin, to input Chinese characters in digital settings. Therefore, we convert original Chinese text into pinyin (e.g., 思想言采都有什大 → sī xiǎng yán lùn jǔ dòng fēng cǎi dū méi yǒu shén me dài qū bié) using the Chinese segmenter Jieba [27] and the open-source dictionary CC-CEDICT [28]. For Japanese, first we use a morphological analyzer MeCab [29] to get the pronunciations of the text and subsequently use romkan [30] to convert them into Roman letters. When MeCab fails to return the pronunciations for words, we have a native speaker create them manually.

Now, each of our datasets contains <audio, original text, processed text> triplets.

4. Experiments

In order to validate our datasets, we train two neural TTS models—Tacotron [2] and DCTTS [8]—on each dataset. We then synthesize audio samples and perform Mean Opinion Score tests.

4.1. Models

Tacotron and DCTTS are well-known neural TTS models that have an attention-based sequence-to-sequence architecture. Tacotron, which was introduced in 2017, is an impactful model that produced high quality results from an end-to-end approach and has since served as an important benchmark. DCTTS is one of the TTS models inspired by Tacotron. DCTTS is different from Tacotron in a few ways. Tacotron’s computational backbone is the CBHG, a combination of convolution layers, fully-connected layers, and recurrent layers, whereas DCTTS only uses convolution layers. Tacotron is trained end-to-end whereas DCTTS has two networks trained independently. DCTTS uses a few tricks that help training such as guided attention and forcibly incremental attention.

One reason for choosing these models is that we have working implementations of both [31][32] which have been tested on multiple datasets. We have found that reproducing original work in neural TTS is non-trivial. Model performance is depen-

dent on many factors, such as data, model architecture, hyperparameters, etc. The neural TTS field has yet to have the ossification of helpful techniques for training that has seemingly already happened in the neighboring field of neural computer vision. Thus, we find it paramount to use models that have successful implementations in order to evaluate new datasets.

4.2. Training

We use Tesla P40 GPUs for training. For both models, we mostly use the given hyperparameters in their respective papers. See our repository for the full list of values. We use the same hyperparameters for all languages. The original Tacotron paper uses more than 2 million training steps, which we find impractical given our resources. In our preliminary tests, we found 400k steps produced good results for DCTTS. For simplicity, we train both models for 400k steps. It takes around ten days and three days to train Tacotron and DCTTS, respectively.

4.3. Evaluation

4.3.1. Test sentences

To evaluate each model’s performance, 20 sentences per language are collected from Tatoeba [33], a web database of sentences for multiple languages. These sentences are carefully chosen to maximize the cover of letters in the vocabulary so we can check the utterance quality of various phonemes. For Chinese and Japanese, phonetic transcriptions, not the original text, are considered. Some letters which are very rare in the language are left out. All sentences are available in our repository.

4.3.2. Mean Opinion Scores

We leverage Amazon Mechanical Turk (MTurk) to gather workers to score the test sentences. MTurk allows requesters to post Human Intelligence Tasks (HITs) for a worker to complete. For each HIT, we ask the worker to listen to an audio clip and score it.

We adopt the standard absolute category rating (ACR) test, where workers are required to give integer scores between 1 and 5. We split evaluation into two categories: speech naturalness and pronunciation accuracy. For speech naturalness, we refer to the rubric for MOS scores in [34], which we include in Table 1 for completeness. We design a new simple rubric for pronunciation accuracy as shown in Table 2.

Table 1: Rubric of naturalness.

Score	Quality	Definition
5	Excellent	Imperceptible distortions (dist.)
4	Good	Just perceptible but not annoying dist.
3	Fair	Perceptible and slightly annoying dist.
2	Poor	Annoying but not objectionable dist.
1	Bad	Very annoying and objectionable dist.

With each HIT, we give the worker a reference sample from the audiobook and tell the worker that the sample should score highly on both naturalness and pronunciation; we then give the worker an audioclip with the corresponding text and ask them score it.

For some languages, MTurk allows us to list qualifications for basic proficiency. While we could opt for this requirement for some languages (e.g., French), we want consistency in the

Table 2: Rubric of pronunciation accuracy.

Score	Quality	Definition
5	Excellent	No mispronunciations
4	Good	Few minor mispronunciations
3	Fair	Many minor mispronunciations
2	Poor	Few major mispronunciations
1	Bad	Many major mispronunciations

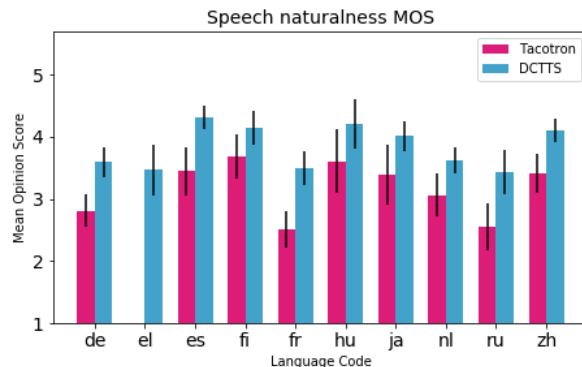


Figure 2: MOS scores for speech naturalness with 95% C.I.

MOS procedure across all languages to make the results more comparable. Thus, instead of relying on it, we require each worker to listen to a reference sample chosen from the language’s dataset and transcribe the text. We consider participants who have done this correctly to be truthful and include their scores in our results.

We use the same method as crowdMOS [34] for computing confidence intervals (C.I.). Because the number of available workers to complete HITs varies with language, we allow languages to have varying numbers of total samples.

4.4. Results

We successfully trained models for all languages except Greek for Tacotron. We believe it is due to the small data size of Greek (~4 hours of audio). Although we successfully trained DCTTS on Greek, the samples were much worse than those from other languages.

The MOS results for speech naturalness and pronunciation accuracy are shown in Figures 2 and 3, respectively. For naturalness, DCTTS is statistically more natural than Tacotron for German, French, and Spanish. For pronunciation accuracy, we find that DCTTS and Tacotron are somewhat similar. The detailed MOS scores are found in Table 3.

Although it’s common to use MOS as a performance metric for TTS models, we recognize that it may not be appropriate to take the mean of Likert scores, because each score belongs to a category and the semantic meaning of these categories need not be evenly distributed along a number line as we (and others) have implied in our rubrics. Thus, we also present the distributions of scores for each model-language pair; the distributions for naturalness and pronunciation accuracy are shown in Tables 4 and 5, respectively. We see that for both speech naturalness and pronunciation accuracy, DCTTS has distributions more skewed towards higher scores than Tacotron, but the effect is more marked for naturalness.

Table 3: MOS scores with 95% C.I. of Tacotron and DCTTS for all languages.

Lang.	Dur. (hh:mm:ss)	# Workers	Speech Naturalness		Pronunciation Accuracy	
			Tacotron	DCTTS	Tacotron	DCTTS
de	16:08:01	75	2.82 ± 0.26	3.59 ± 0.24	4.31 ± 0.22	4.25 ± 0.24
el	04:08:14	43	N/A	3.47 ± 0.41	N/A	3.70 ± 0.39
es	23:49:49	78	3.45 ± 0.38	4.31 ± 0.19	4.31 ± 0.40	4.72 ± 0.13
fi	10:32:03	62	3.69 ± 0.36	4.15 ± 0.27	4.16 ± 0.30	4.43 ± 0.27
fr	19:09:03	75	2.51 ± 0.30	3.50 ± 0.27	3.48 ± 0.48	4.24 ± 0.25
hu	10:00:25	39	3.60 ± 0.51	4.21 ± 0.39	4.27 ± 0.40	4.51 ± 0.35
ja	14:55:36	79	3.39 ± 0.48	4.01 ± 0.24	3.61 ± 0.48	4.20 ± 0.21
nl	14:06:40	97	3.06 ± 0.35	3.63 ± 0.21	3.82 ± 0.31	3.93 ± 0.20
ru	21:22:10	73	2.55 ± 0.37	3.43 ± 0.35	3.64 ± 0.39	4.02 ± 0.32
zh	06:27:04	122	3.41 ± 0.31	4.10 ± 0.19	4.19 ± 0.20	4.46 ± 0.17

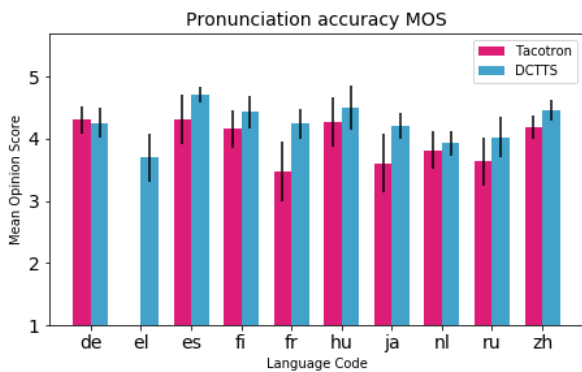


Figure 3: MOS scores for pronunciation accuracy with 95% C.I.

In general, when listening to samples, both models capture the original voice, but the samples of DCTTS sound better than those of Tacotron. We found DCTTS produced fairly clean speech, while Tacotron consistently had noise in the outputs.

Both models exhibited mumbling at the end of a few generated samples. Although this was not expected—the models should have learned trailing silences—we were able to remove mumbling via trimming.

In Japanese, we found utterances that were not relevant to the input text were generated at the end in a few samples. We suspect it is derived from the imperfection of the automatic phonetic transcription in the training data.

5. Conclusions

We discussed how we built CSS10, a collection of single-speaker speech datasets for 10 languages, and how we used it for a TTS task. Despite the fact that there were differences in model performance depending on the language, we were able to train the models successfully on our datasets. We release all the resources for this project including source code, datasets, pre-trained models, and evaluation data to the public.

6. Future work

We hope CSS10 and our experiments serve as benchmarks for future non-English TTS research. We found that some automatic phonetic transcriptions for Japanese and Chinese contain errors. If these errors are corrected, perhaps the model perfor-

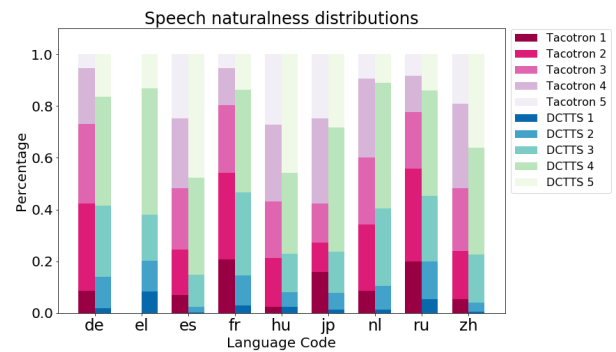


Figure 4: Percentages of naturalness scores shown by model and language.

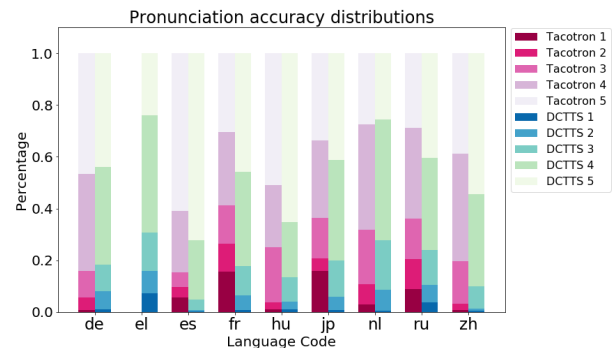


Figure 5: Percentages of pronunciation accuracy scores shown by model and language.

mance will improve. Additionally, we plan to add a Korean dataset. Because there are not enough audiobooks available for Korean in LibriVox, we are willing to produce recordings ourselves. When they are ready, we will release them with our other languages. Although we validated our datasets on TTS models in this work, CSS10 can be used for other speech tasks such as multi-lingual speech recognition.

7. Acknowledgments

We thank the LibriVox creator for his platform which allows public access to audiobooks, the performers who recorded their readings of audiobooks, and Kakao Brain for funding this work.

8. References

- [1] 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," *arXiv preprint arXiv:1609.03499*, 2016.
- [2] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio *et al.*, "Tacotron: A fully end-to-end text-to-speech synthesis model," *arXiv preprint arXiv:1703.10135*, 2017.
- [3] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan *et al.*, "Natural tts synthesis by conditioning wavenet on mel spectrogram predictions," *arXiv preprint arXiv:1712.05884*, 2017.
- [4] J. Sotelo, S. Mehri, K. Kumar, J. F. Santos, K. Kastner, A. Courville, and Y. Bengio, "Char2wav: End-to-end speech synthesis," 2017.
- [5] S. O. Arik, M. Chrzanowski, A. Coates, G. Diamos, A. Gibiansky, Y. Kang, X. Li, J. Miller, J. Raiman, S. Sengupta *et al.*, "Deep voice: Real-time neural text-to-speech," *arXiv preprint arXiv:1702.07825*, 2017.
- [6] S. O. Arik, G. Diamos, A. Gibiansky, J. Miller, K. Peng, W. Ping, J. Raiman, and Y. Zhou, "Deep voice 2: Multi-speaker neural text-to-speech," *arXiv preprint arXiv:1705.08947*, 2017.
- [7] W. Ping, K. Peng, A. Gibiansky, S. O. Arik, A. Kannan, S. Narang, J. Raiman, and J. Miller, "Deep voice 3: 2000-speaker neural text-to-speech," *arXiv preprint arXiv:1710.07654*, 2017.
- [8] H. Tachibana, K. Uenoyama, and S. Aihara, "Efficiently trainable text-to-speech system based on deep convolutional networks with guided attention," *arXiv preprint arXiv:1710.08969*, 2017.
- [9] Y. Taigman, L. Wolf, A. Polyak, and E. Nachmani, "Voice synthesis for in-the-wild speakers via a phonological loop," *arXiv preprint arXiv:1707.06588*, 2017.
- [10] K. Ito, "The lj speech dataset," <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [11] "Blizzard challenge 2018," https://www.synsig.org/index.php/Blizzard_Challenge_2018, 2018.
- [12] J. Yamagishi, T. Nose, H. Zen, Z. H. Ling, T. Toda, K. Tokuda, S. King, and S. Renals, "Robust speaker-adaptive hmm-based text-to-speech synthesis," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 17, no. 6, pp. 1208–1230, Aug 2009.
- [13] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," <http://www.openslr.org/12/>, 2015.
- [14] A. Rousseau, P. Deléglise, and Y. Estève, "Ted-lium: an automatic speech recognition dedicated corpus," <http://www-lium.univ-lemans.fr/en/content/ted-lium-corpus>, 2014.
- [15] "Voxforge," <http://www.voxforge.org/>, 2006.
- [16] Mozilla, "Common voice," <https://voice.mozilla.org/>, 2017.
- [17] P. Baljekar, "Speech synthesis from found data," 2018.
- [18] K. Park, "The world english bible," <https://www.kaggle.com/bryanpark/the-world-english-bible-speech-dataset>, 2017.
- [19] R. Sonobe, S. Takamichi, and H. Saruwatari, "Jsut corpus: Free large-scale japanese speech corpus for end-to-end speech synthesis," <http://arne.chark.eu/static/spoken-wp-corpus-collection.pdf>, <https://nats.gitlab.io/swc/>, 2017.
- [20] "Pavoque corpus of expressive speech," <https://github.com/marytts/pavoque-data>, 2009.
- [21] T. Baumann, A. Kohn, and F. Hennig, "The spoken wikipedia corpus collection," <http://arne.chark.eu/static/spoken-wp-corpus-collection.pdf>, <https://nats.gitlab.io/swc/>, 2016.
- [22] A. Stan, O. Watts, Y. Mamiya, M. Giurgiu, R. A. J. Clark, J. Yamagishi, and S. King, "Tundra: A multilingual corpus of found data for tts research created with light supervision," 08 2013.
- [23] I. Solak, "The m-ailabs speech dataset," <https://www.caito.de/2019/01/the-m-ailabs-speech-dataset/>, January 2019.
- [24] "Librivox," <https://librivox.org/>, 2018.
- [25] "Audacity," <http://audacity.sourceforge.net/>, 2018.
- [26] R. Ochshorn and M. Hawkins, "Gentle," <https://github.com/lowerquality/gentle>, 2017.
- [27] J. Sun, "Jieba," <https://github.com/fxsjy/jieba>, 2017.
- [28] P. Denisowski, "Cc-cedict," <https://cc-cedict.org/editor>, 2018.
- [29] "Mecab: Yet another part-of-speech and morphological analyzer," <http://taku910.github.io/mecab>, 2006.
- [30] M. Yao, "python-romkan," <https://www.soimort.org/python-romkan/>, 2013.
- [31] K. Park and T. Mulc, "A (heavily documented) tensorflow implementation of tacotron: A fully end-to-end text-to-speech synthesis model," <https://github.com/Kyubyong/tacotron>, 2018.
- [32] K. Park, "A tensorflow implementation of dc-tts: yet another text-to-speech model," https://github.com/Kyubyong/dc_tts, 2018.
- [33] Sysko, "Tatoeba," <https://tatoeba.org>, 03 2013.
- [34] F. P. Ribeiro, D. Florencio, C. Zhang, and M. Seltzer, "Crowdmos: An approach for crowdsourcing mean opinion score studies," in *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, 05 2011, pp. 2416–2419.