Research Repository UCD

Title	Speech Quality Factors for Traditional and Neural-Based Low Bit Rate Vocoders
Authors(s)	Jassim, Wissam A., Skoglund, Jan, Chinen, Michael, Hines, Andrew
Publication date	2020-05-28
Publication information	Jassim, Wissam A., Jan Skoglund, Michael Chinen, and Andrew Hines. "Speech Quality Factors for Traditional and Neural-Based Low Bit Rate Vocoders." IEEE, 2020.
Conference details	International Conference on Quality of Multimedia Experience (QoMEX), Dublin, Ireland, 26-28 May 2020
Publisher	IEEE
Item record/more information	http://hdl.handle.net/10197/12212
Publisher's statement	Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes, or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works, must be obtained from the IEEE.
Publisher's version (DOI)	10.1109/QoMEX48832.2020.9123109

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/340223989>

Speech Quality Factors for Traditional and Neural-Based Low Bit Rate Vocoders

Preprint · March 2020

CITATIONS

0

READS

32

4 authors, including:

[Jan Skoglund](#)

Google LLC, San Francisco, United States

78 PUBLICATIONS 821 CITATIONS

[SEE PROFILE](#)

[Michael Chinen](#)

16 PUBLICATIONS 57 CITATIONS

[SEE PROFILE](#)

[Andrew Hines](#)

Technological University Dublin - City Campus

49 PUBLICATIONS 275 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:

[Ambiquat View project](#)

Speech Quality Factors for Traditional and Neural-Based Low Bit Rate Vocoders

Wissam A. Jassim
School of Computer Science
University College Dublin
Dublin, Ireland
wissam.a.jassim@gmail.com

Jan Skoglund
Chrome Media
Google
San Francisco, CA, USA
jks@google.com

Michael Chinen
Chrome Media
Google
San Francisco, CA, USA
mchinen@google.com

Andrew Hines
School of Computer Science
University College Dublin
Dublin, Ireland
andrew.hines@ucd.ie

Abstract—This study compares the performances of different algorithms for coding speech at low bit rates. In addition to widely deployed traditional vocoders, a selection of recently developed generative-model-based coders at different bit rates are contrasted. Performance analysis of the coded speech is evaluated for different quality aspects: accuracy of pitch periods estimation, the word error rates for automatic speech recognition, and the influence of speaker gender and coding delays. A number of performance metrics of speech samples taken from a publicly available database were compared with subjective scores. Results from subjective quality assessment do not correlate well with existing full reference speech quality metrics. The results provide valuable insights into aspects of the speech signal that will be used to develop a novel metric to accurately predict speech quality from generative-model-based coders.

Index Terms—speech quality assessment, neural speech synthesis, WaveNet, LPCNet, Opus, vocoder.

I. INTRODUCTION

A vocoder system refers to the process that analyzes and synthesizes human speech for a wide range of applications such as speech compression, voice transformation, and multiplexing. It has been studied for decades and many well-developed algorithms have been introduced [1] [2]. These algorithms are designed to compress speech signals with low bit rates and high speech quality. The performance of any vocoder system is determined by other desirable properties such as robustness to different speakers/languages, robustness to channel errors, computational complexity and memory aspects, and minimizing the coding delay [3].

Recently, several deep neural network (DNN) speech synthesis algorithms that contain a generative model have been proposed for text-to-speech synthesis and low bit-rate coding (compression) [4] [5] [6] [7]. WaveNet [8], a deep neural network for generating raw audio waveforms, was the first method of these algorithms. It was originally designed for text-to-speech synthesis and has since been used for parametric coding [6], in which the speech signal is represented as a sequence of parameters extracted at the encoder. To improve memory efficiency, SampleRNN [9] uses a sparse recurrent neural network (RNN) instead of the convolutional networks. It has been shown to be useful for speech coding [10] as it is able to overcome the problem of modeling extremely

high-resolution temporal data. WaveRNN [11] also focused on finding more efficient models in order to reduce the complexity of speech synthesis compared to WaveNet. Furthermore, the LPCNet model [12] has shown to be able to provide lower complexity and real-time operation by combining linear prediction with WaveRNN. It achieved significantly higher quality than WaveRNN for the same network size. All of the models generate samples by sampling from a distribution, and most are autoregressive. This means the number of waveforms that are likely for a given sequence of conditioning input increases with the model's uncertainty. The typical errors generated by model failures are incorrect mappings on the speech manifold, such as phoneme mismatches or slurred speech, as opposed to non-speech artifacts introduced by traditional coders.

The neural-based generative networks have recently been adopted to improve the performance of other types of coders (waveform and parametric coders). For example, a backward-compatible way of improving the quality of the low bit rate Opus coder [13] by re-synthesizing speech from the decoded parameters using two different neural generative models, WaveNet (high complexity, and high-latency architecture) and LPCNet (low-complexity, low-latency RNN-based generative model) was proposed in [14]. The two systems were used to extract conditioning features from the Opus bit stream coded at 6 kb/s. These conditioning features represent the spectral shape and the pitch of the signal for performing signal decoding and reconstruction. According to the subjective tests, synthesized speech using LPCNet outperformed the output of the standard Opus decoder for the same 6 kb/s bit stream.

The speech quality assessments summarised in Section II rated the neural-based coders highly. However, in Section III we observe that traditional full reference objective speech quality metrics such as PESQ (Perceptual Evaluation of Speech Quality) measure [15] fail to provide accurate quality scores for the speech signals processed by this kind of speech coders. Therefore, a new quality measure is needed that is robust enough to work with different types of coders. As a first step, making meaningful comparisons between the performance of the most recent neural-based coders is necessary to determine the related factors that need to be taken into account when designing a new quality measures.

Consequently, in Section IV of this study, the quality of

TABLE I
CODERS ADOPTED IN THIS STUDY.

Abbreviation	Coder Name	Bit rate	Type	Database	Notes	Reference
LPCNetUnquant	Unquantized LPCNet	—	Generative-model-based coder	Set 1 [16]	LPCNet operating on Opus unquantized features	[16]
Opus9.0	Opus	9 kb/s	Hybrid waveform-matching coder	Set 1 [16], Set 2 [14]	Wideband vocoder (SILK mode)	[13] [14]
WaveNet6.0	WaveNet	6 kb/s	Generative-model-based coder	Set 2 [14]	WaveNet operating on Opus quantized features	[14]
LPCNet1.6	LPCNet	1.6 kb/s	Generative-model-based coder	Set 1 [16]	WaveRNN + linear prediction	[12] [16]
LPCNet6.0	LPCNet	6 kb/s	Generative-model-based coder	Set 2 [14]	LPCNet operating on Opus quantized features	[14]
MELP2.4	MELP	2.4 kb/s	Source-filter coder	Set 1 [16]	Narrowband vocoder	[17]
Opus6.0	Opus	6 kb/s	Hybrid waveform-matching coder	Set 2 [14]	Narrowband vocoder (SILK mode)	[13] [14]
Speex4.0	Speex	4 kb/s	Hybrid waveform coder?	Set 1 [16], Set 2 [14]	Wideband vocoder (wideband quality 0)	[18]

speech coders is evaluated in terms of multi-dimensional aspects. The coded speech signals obtained from most recent coders are evaluated in terms of the following quality aspects: subjective listening scores with a MUSHRA-inspired methodology [19], accuracy of pitch period estimation, word error rate for automatic speech recognition, robustness to speaker gender, and robustness to coding delays.

II. SPEECH QUALITY DATASETS FOR CODED SPEECH

In this study, we employed the same speech materials with their subjective listening scores that were used to test the quality of the LPCNet coder at 1.6 kb/s in [16] and the modified Opus coder at 6 kb/s in [14]. As stated in [16], 16 samples from 3 male and 3 female speakers (Set 1) were used in the testing stage. On the other hand, the speech database employed to test the quality of the modified Opus coder in [14] consists of 16 samples from 3 male and 3 female speakers (Set 2). For the two studies, all samples are part of the NTT Multi-Lingual Speech Database for Telephony database [20]. In addition, all samples from the selected speakers for the test were excluded from the training set. The subjective listening tests with a MUSHRA-inspired methodology [19] were conducted on 100 listeners for each test.

Table I lists the speech coders that were employed for performance analysis in this study. The coders are sorted in descending order according to their MUSHRA scores. Fig. 1 shows these scores as a function of coder type. It can be seen that the best results were obtained by the LPCNet coder operating on unquantized features (LPCNetUnquan). The 9 kb/s Opus (Opus9.0) and 6 kb/s WaveNet (WaveNet6.0) coders achieved comparable results. The RNN-based coder at 1.6 kb/s (LPCNet1.6) outperformed the same coder operating on Opus quantized features at 6 kb/s (LPCNet6.0). Both WaveNet6.0 and LPCNet6.0 at provided higher quality results than that of the standard Opus decoder for the same 6 kb/s bit stream. The quality of MELP coder (source-filter coder) at 2.4 kb/s obtained better results than that of the Speex coder at 4 kb/s.

III. OBJECTIVE SPEECH QUALITY EVALUATION

In order to evaluate the quality of the speech signals coded by different coders, PESQ [15] scores were computed for the reference speech samples and their corresponding coded (synthesized) samples. Results were averaged by coder type and compared to the average MUSHRA scores per coder type. Fig. 2 shows the PESQ scores as a function of MUSHRA

Fig. 1. Subjective listening results of different types of coders taken from the LPCNet [16] and OpusNet [14] studies. Note that the 9 kb/s Opus and 4 kb/s Speex coders were used in both studies.

Fig. 2. Objective PESQ measure against subjective MUSHRA scores for different coders. Results were averaged by coder type and compared to the corresponding average MUSHRA scores.

scores. In the plot, the relationship points are highlighted in two different colors. The points highlighted in green represent the results for the parametric coders (Opus9.0, MELP2.4, Opus6.0 and Speex4.0) whereas the points highlighted in red are the results for the generative-model-based coders (LPCNetUnquan, WaveNet6.0, LPCNet1.6 and LPCNet6.0). It can be seen that the PESQ scores of the parametric coders obtained good correlation with the subjective listening results. On the other hand, the generative-model-based coders provided poor correlation in terms of PESQ against MUSHRA scores

Fig. 3. Example of pitch (fundamental frequency) plots for a reference speech sample with its corresponding synthesised sample. In this example the speech signal was processed by the LPCNet6.0 coder. The spectrogram frequencies are shown up to 2 kHz only.

relationship. This suggests that the traditional speech quality metrics may not be suitable for evaluating the performance of generative-model-based coders. Standard objective speech quality measures such as PESQ are unable to reflect the distortions in the speech signals of these coders. Similar behaviour was observed by the authors with other full reference quality metrics (e.g. POLQA [21] and ViSQOL [22]) motivating this study and more comprehensive experiments to design a new quality measure for these kind of coders.

IV. COMPARATIVE ANALYSIS OF CODER QUALITY

It is important to investigate how speech coders, especially the generative-model-based coders, behave in terms of other quality aspects as illustrated in the following subsections. Having an idea about the effect of these aspects on coded speech, it is useful to improve the performance of the current coders and to develop more robust coders in the future.

A. Accuracy of pitch estimation

As an example of how the pitch (fundamental frequency) looks like, Fig. 3 shows the plots of the pitch sequence as a function of time for a reference speech sample taken from the NTT database with its corresponding speech obtained from a generative-model-based coder. The pitch frequencies were computed using Praat software [23]. The spectrogram representation of the two signals are also shown in Fig. 3.

The `mir_eval` [24] python package was used to compute pitch accuracy given two pitch sequences. The results are compared in terms of two measures [25] [26]: raw pitch accuracy (RPA) and raw chroma accuracy (RCA) with a threshold of 50 cents. Fig. 4 shows the RPA and RCA results against coder type. It can be seen that the best results were obtained by the Opus coder at 9 kbps. The other parametric coders (MELP2.4,

Fig. 4. Pitch accuracy scores as a function of coder types.

Opus6.0 and Speex4.0) also outperformed by the generative-model-based coders. Among these coders, the LPCNetUnquan and LPCNet6.0 achieved comparable accuracies in terms of pitch estimation, whereas the WaveNet6.0 obtained the lowest RPA and RCA scores. Furthermore, the RPA results averaged by coder were plotted as a function of the average MUSHRA scores as shown in Fig. 5. Again, the neural-based coders obtained poor correlation compared to that of the other coders.

B. Word error rates for automatic speech recognition

In this analysis, the speech signal obtained from each speech coder was fed as an input to the speech recognition library [27] with the Google speech-to-text API [28]. The word

Fig. 5. Pitch accuracy scores against subjective MUSHRA scores for different coders. Results were averaged by coder type and compared to the corresponding average MUSHRA scores.

error rate (WER) between the ground-truth transcript and the hypothesis transcript of the API was then computed. Fig. 6 shows the plot of WER results for all coders. As expected, the Opus9.0 and LPCNetUnquan coders achieved the best results. The neural-based coders at low bit rates (WaveNet6.0, LPCNet1.6 and LPCNet6.0) outperformed the traditional parametric coders (MELP2.4, Opus6.0 and Speex4.0). The MELP2.4 coder obtained the worst error rates whereas the Opus6.0 and Speex4.0 coders achieved comparable results. The WER versus MUSHRA scores relationship is shown in Fig. 7. The overall trend is an inverse correlation of WER to quality with the parametric coders (except Opus9.0) displaying higher WERs than the generative-model-based coders.

The pitch tracking results highlighted that while there are differences in pitch tracking observed from the neural-based coders, these do not influence perceived quality and a new objective quality metric needs to take this into account. From the WER results, a similar trend is observed for intelligibility. Fig. 7 highlights that there is a strong inverse correlation between WER and the MUSHRA scores, i.e. WER is a better predictor of quality than the objective quality metric results seen in Fig. 2. This points to the potential to use automatic speech recognition (ASR) as part of an objective quality metric solution although as the ASR model is language dependant, any such solution would be less flexible than the current objective speech quality models that are not language specific in design.

C. Robustness to speaker's gender

In this analysis, we investigate the influence of speakers gender on the quality of coded speech. Before describing the outcomes of this experiment, it is necessary to mention that each speech sample in the employed database contains two utterances, and each utterance was spoken by a different speaker. This approach generated speech samples with four categories of speaker's gender: Male-Male, Male-Female, Female-Male and Female-Female. The Male-Male category means that the two utterances in the sample were spoken by

Fig. 6. Results of WER as a distance measure between the ground-truth transcript and the hypothesis transcript.

Fig. 7. WER against subjective MUSHRA scores for different coders. Results were averaged by coder type and compared to the corresponding average MUSHRA scores.

two male speakers, whereas female-female category contains speech samples that their utterances were spoken by two female speakers, and so on for other categories.

The subjective listening results were reported as a function of speaker's gender for each coder, and the results are shown in Fig. 8. The speech samples with Female-Male category achieved highest subjective scores for Opus coders. For the generative-model-based coders, the highest subjective scores were obtained when the first utterance of the speech sample is spoken by a female speaker, i.e. the speech samples with Female-Male and Female-Female categories. This behaviour is consistent with that of the Speex4.0 coder. The Male-Male speech samples achieved the best results for the MELP2.4 coder. It can be seen that the speaker's gender has an impact on the overall subjective listening results as the Mushra scores seemed to be affected by speaker's gender. This factor needs to be further investigated when designing new speech coders especially for the neural-based models.

Fig. 8. Subjective MUSHRA scores for different coders and speaker's gender category

Fig. 9. Results of average coding delays between reference and coded speech samples.

D. Robustness to coding delays

The effect of coding delay on the quality of different speech coders was also tested in this study. The cross-correlation was computed between the reference and coded speech signals, and the index of the maximum values (argmax) was taken as time leads or lags. Fig. 9 shows the average coding delays as a function of coder type. As expected, the generative-model-based coders incur longer delays in time compared to that of the other coders. For delays plotted against the subjective scores as shown in Fig. 10 there is some correlation apparent (except Speex4.0). Identifying delay characteristics from the coded speech could be useful for classifying coder behaviour for data driven objective quality models.

Another experiment was run to show the micro-alignment time difference between reference speech and its corresponding coded speech. Fig. 11 shows waveform of a reference speech sample taken from the same database. The same sample was processed by two coders: the WaveNet6.0 and Speex4.0. A Mel spectrogram with 64 band to 8 kHz was then computed for each signal. The 2 dimensional plot of each Mel spectrogram is shown in the same figure. The dynamic time wrapping (DTW) algorithm was performed on the two Mel spectrograms to measure similarity (in terms of temporal information) between the reference and coded speech waveforms [29]. The optimal path

Fig. 10. Average coding delays against subjective MUSHRA scores for different coders. Results were averaged by coder type and compared to the corresponding average MUSHRA scores.

that has the minimal cost (summation of absolute differences) was then tracked and plotted. It can be seen that, although the Speex coder obtained poor subjective scores (see Fig. 1), it provided very low micro delays as its optimal path is a straight diagonal. On the other hand, the micro-alignment time shifts are very apparent for the WavNet6.0 coder.

V. CONCLUSION

In this study, different types of speech coders were compared in terms of several objective measures. Experiments were conducted using speech samples taken from the NTT database with its subjective MUSHRA scores. The existing full reference speech quality metrics failed to provide scores consistent with subjective quality assessment. Taking into account the influence of the quality aspects considered in this study is necessary to design a new quality metric that is robust enough to work with different coders at different bit rates. One possible approach is to combine results from the objective measures employed in this study based on multi modelling algorithm to design the required speech quality metric.

Fig. 11. Example of 1.5 s of Reference, WaveNet6.0 and Speex4.0 coded speech. The waveforms, Mel spectrogram in dB, and DTW cost matrix plots are shown to highlight the micro-alignment changes in neural-based vocoded speech.

ACKNOWLEDGMENTS

This publication has emanated from research supported in part by the Google Chrome University Program and research grants from Science Foundation Ireland (SFI) co-funded under the European Regional Development Fund under Grant Number 13/RC/2289_P2 and 13/RC/2077.

REFERENCES

- [1] J. Markel and A. Gray, "A linear prediction vocoder simulation based upon the autocorrelation method," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 22, no. 2, pp. 124–134, April 1974.
- [2] B. S. Atal and S. L. Hanauer, "Speech analysis and synthesis by linear prediction of the speech wave," *The Journal of the Acoustical Society of America*, vol. 50, no. 2B, pp. 637–655, 1971.
- [3] W. B. Kleijn and K. K. Paliwal, *Speech Coding and Synthesis*. USA: Elsevier Science Inc., 1995.
- [4] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan, R. A. Saurous, Y. Agiomvrgianakis, 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)*, April 2018, pp. 4779–4783.
- [5] S. Ö. Arik, G. F. Diamos, A. Gibiansky, J. Miller, K. Peng, W. Ping, J. Raiman, and Y. Zhou, "Deep voice 2: Multi-speaker neural text-to-speech," *CoRR*, vol. abs/1705.08947, 2017. [Online]. Available: <http://arxiv.org/abs/1705.08947>
- [6] W. B. Kleijn, F. S. C. Lim, A. Luebs, J. Skoglund, F. Stimberg, Q. Wang, and T. C. Walters, "Wavenet based low rate speech coding," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 676–680, 2018.
- [7] R. Prenger, R. Valle, and B. Catanzaro, "Waveglow: A flow-based generative network for speech synthesis," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 3617–3621.
- [8] 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 *Arxiv*, 2016. [Online]. Available: <https://arxiv.org/abs/1609.03499>
- [9] S. Mehri, K. Kumar, I. Gulrajani, R. Kumar, S. Jain, J. Sotelo, A. C. Courville, and Y. Bengio, "SampleRNN: An unconditional end-to-end neural audio generation model," *Arxiv*, vol. abs/1612.07837, 2016.
- [10] J. Klejsa, P. Hedelin, C. Zhou, R. Fejgin, and L. Villemoes, "High-quality speech coding with sample RNN," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 7155–7159.
- [11] 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," *CoRR*, vol. abs/1802.08435, 2018. [Online]. Available: <http://arxiv.org/abs/1802.08435>
- [12] J. Valin and J. Skoglund, "LPCNET: Improving Neural Speech Synthesis through Linear Prediction," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 5891–5895.
- [13] J.-M. Valin, K. Vos, and T. Terriberry, "Definition of the Opus Audio Codec," RFC 6716, Sep. 2012. [Online]. Available: <https://rfc-editor.org/rfc/rfc6716.txt>
- [14] J. Skoglund and J.-M. Valin, "Improving Opus low bit rate quality with neural speech synthesis," *ArXiv*, vol. abs/1905.04628, 2019.
- [15] ITU, "Perceptual Evaluation of Speech Quality (PESQ): an objective method for end-to-end speech quality assessment of narrow-band telephone networks and speech codecs," in *ITU-T Rec. P.862*, 2001.
- [16] J.-M. Valin and J. Skoglund, "A Real-Time Wideband Neural Vocoder at 1.6kb/s Using LPCNet," in *Proc. Interspeech 2019*, 2019, pp. 3406–3410. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2019-1255>
- [17] A. McCree, Kwan Truong, E. B. George, T. P. Barnwell, and V. Viswanathan, "A 2.4 kbit/s MELP coder candidate for the new U.S. Federal Standard," in *1996 IEEE International Conference on Acoustics, Speech, and Signal Processing Conference Proceedings*, vol. 1, May 1996, pp. 200–203 vol. 1.
- [18] J.-M. Valin, *The Speex codec manual*, Xiph.Org Foundation, 2007. [Online]. Available: <https://speex.org/docs/manual/speex-manual.pdf>
- [19] ITU-R, "Recommendation BS.1534-1: Method for the subjective assessment of intermediate quality level of coding systems," in *International Telecommunications Union*, 2001.
- [20] NTT Advanced Technology, *Multilingual Speech Database for Telephonometry, CD-ROM*, 1994.
- [21] ITU, "Perceptual objective listening quality assessment," Int. Telecomm. Union, Geneva, Switzerland, ITU-T Rec. P.863, 2018.
- [22] A. Hines, J. Skoglund, A. C. Kokaram, and N. Harte, "ViSQOL: An objective speech quality model," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2015, no. 1, p. 13, 2015.
- [23] P. Boersma and D. Weenink, "PRAAT, a system for doing phonetics by computer," *Glott international*, vol. 5, pp. 341–345, 01 2001.
- [24] C. Raffel, B. McFee, E. J. Humphrey, J. Salamon, O. Nieto, D. Liang, and D. P. W. Ellis, "MIR-EVAL: A Transparent Implementation of Common MIR Metrics," in *Proc. of the 15th International Society for Music Information Retrieval Conference ISMIR*, 2014.
- [25] G. E. Poliner, D. P. W. Ellis, A. F. Ehmann, E. Gomez, S. Streich, and B. Ong, "Melody Transcription From Music Audio: Approaches and Evaluation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 4, pp. 1247–1256, May 2007.
- [26] E. Gómez, S. Streich, B. Ong, R. P. Paiva, S. Tappert, J. M. Batke, G. Poliner, D. Ellis, and J. P. Bello, "A Quantitative Comparison of Different Approaches for Melody Extraction from Polyphonic Audio Recordings," 2006. [Online]. Available: <http://hdl.handle.net/10230/41912>
- [27] A. Zhang, *Speech Recognition (Version 3.8)*, 2017. [Online]. Available: <https://github.com/Uberi/speech-recognition>
- [28] Google Speech API. [Online]. Available: <https://cloud.google.com/speech>
- [29] M. Müller, *Fundamentals of music processing: Audio, analysis, algorithms, applications*. Springer, 2015.