



| | |
|-------------------------------------|---|
| Title | ViSQOL v3: An Open Source Production Ready Objective Speech and Audio Metric |
| Authors(s) | Chinen, Michael, Lim, Felicia S. C., Skoglund, Jan, Gureev, Nikita, O’Gorman, Feargus, Hines, Andrew |
| Publication date | 2020-05-28 |
| Publication information | Chinen, Michael, Felicia S. C. Lim, Jan Skoglund, Nikita Gureev, Feargus O’Gorman, and Andrew Hines. “ViSQOL v3: An Open Source Production Ready Objective Speech and Audio Metric.” IEEE, May 28, 2020. https://doi.org/10.1109/QoMEX48832.2020.9123150 . |
| Conference details | The 12th International Conference on Quality of Multimedia Experience (QoMEX), Athlone, Ireland (held online due to coronavirus outbreak), 26-28 May 2020 |
| Publisher | IEEE |
| Item record/more information | http://hdl.handle.net/10197/12214 |
| 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/QoMEX48832.2020.9123150 |

Downloaded 2026-05-01 23:51:03

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

ViSQOL v3: An Open Source Production Ready Objective Speech and Audio Metric

Michael Chinen, Felicia S. C. Lim, and Jan Skoglund
Chrome Media Audio
Google LLC
 San Francisco, USA
 mchinen,flim,jks@google.com

Nikita Gureev
Hangouts Meet
Google LLC
 Stockholm, Sweden
 gureev@google.com

Feargus O’Gorman and Andrew Hines
School of Computer Science
University College of Dublin
 Dublin, Ireland
 feargusog@gmail.com, andrew.hines@ucd.ie

Abstract—Estimation of perceptual quality in audio and speech is possible using a variety of methods. The combined v3 release of ViSQOL and ViSQOLAudio (for speech and audio, respectively,) provides improvements upon previous versions, in terms of both design and usage. As an open source C++ library or binary with permissive licensing, ViSQOL can now be deployed beyond the research context into production usage. The feedback from internal production teams at Google has helped to improve this new release, and serves to show cases where it is most applicable, as well as to highlight limitations. The new model is benchmarked against real-world data for evaluation purposes. The trends and direction of future work is discussed.

Index Terms—Perceptual audio quality assessment, mean opinion score estimation, ViSQOLAudio, ViSQOL, PESQ, POLQA, PEAQ, PEMO-Q

I. INTRODUCTION

There are numerous objective metrics available, i.e, metrics obtained by measurements on the audio signal, to assess the quality of recorded audio clips. Examples of physical measures include signal-to-noise ratio (SNR), total harmonic distortion (THD), and spectral (magnitude) distortion. When estimating perceived quality, PESQ [1], [2] and POLQA [3], [4] have become standards for speech, and in practice also for general audio, despite being originally designed to target only speech quality. There are other notable examples, e.g., PEAQ [5] and PEMO-Q [6]. Most of these metrics require commercial licenses. ViSQOL [7] and ViSQOLAudio [8] (referred to collectively as ViSQOL below), are freely available alternatives for speech and audio. These metrics are continually being expanded to cover additional domains. For example the work on AMBIQUAL [9] extends the same principles used in ViSQOLAudio into the ambisonics domain.

Advancements in speech and audio processing, such as denoising and compression, propel the need for improvements in quality estimation. For example, speech and audio codecs reach lower and lower useful bitrates. As such, it may be worthwhile to analyze the performance of ViSQOL for this extended domain. Furthermore, there have been a number of deep neural network (DNN) generative models that recreate

the waveform by sampling from a distribution of learned parameters. One example is the WaveNet-based low bitrate coder [10], which is generative in nature. There are other DNN-based generative models, including SampleRNN [11] and WaveGlow [12], with promising results that suggest that this trend will continue. These generative models typically do not lend themselves to being analyzed well by existing full reference speech quality metrics. While the work described in this paper does not propose a solution to the generative problem, the limitations of the current model should be acknowledged to encourage development of solutions.

ViSQOL was originally designed with a polynomial mapping of the neurogram similarity index measure (NSIM) [13] to MOS, and ViSQOLAudio was extended to use a model trained for support vector regression. Since then, deep neural network models have emerged and been applied to speech quality models [14], [15]. Such approaches are promising and potentially can resolve some of the issues that the current architectures cannot. While such new directions are clearly interesting and warrant further investigation, they are rapidly evolving.

we present a new version of ViSQOL, v3, which contains incremental improvements to the existing framework based on real-world feedback, rather than fundamental changes such as end-to-end DNN modeling. Since ViSQOL has been presented and benchmarked in a large number of experiments that have validated its application to a number of use cases [7], [8], [16]–[19] we consider it relatively well analyzed for the known datasets, which tend to be smaller and relatively homogeneous. We instead turn our attention to the data and types of problems encountered “in the wild” at Google teams that were independent of ViSQOL development, and the iterative improvements that have come from this analysis. Adapting it to these cases has yielded various improvements to usability and performance, along with feedback and insights about the design of future systems for estimating perceptual quality. Since the nature these improvements fill the ‘blind spots’ in the datasets, they are not expected to improve its results on these datasets. Until there is the creation of more diverse subjective score datasets, real-world validation seems to be a reasonable compromise.

Additionally and alongside improving the quality of MOS

estimation from real-world data, we are concerned with how to make ViSQOL more useful to the community from a practical tooling perspective. Even though ViSQOL was available through a MATLAB implementation, there were still unnecessary hurdles to use it in certain cases, e.g. production and continuous integration testing, (which may need to run on a server), or may not have MATLAB licenses available. As a result, we chose to re-implement it in C++ because it is a widely available and extensible language that can be wrapped in other languages. We decided to put the code on GitHub for ease of access and contribution.

This paper is structured as follows: in section II, a case study of the findings and challenges encountered when integrating ViSQOL into various Google projects. In section III, we present the general design and algorithmic improvements that are in the new version. Then in section IV, the improvements with respect to the case studies are discussed. Finally, we summarize in a concluding section V.

II. CASE STUDIES AND USER FEEDBACK

This version of ViSQOL is the result of the integration process of ViSQOL, using real production and integration testing cases at Google. The case studies described in this section were initiated by individual teams that were independent of prior ViSQOL development. They typically consulted with a ViSQOL developer to verify appropriate usage, or read the documentation and integrated ViSQOL on their own.

A. Hangouts Meet

The Meet team has been successfully using ViSQOL for assessing audio quality in Hangouts Meet. Hangouts Meet is a video communication service that uses WebRTC [20] for transmitting audio. Meet uses a testbed that is able to reliably replicate adverse network conditions to assess the quality of audio during the call. For this use case they have 48 kHz-sampled reference and degraded audio samples and use ViSQOLAudio for calculating the results.

In order to ensure that ViSQOL works reliably for this use case, it was compared to an internal no-reference audio quality metric that is based on technical metrics of a WebRTC-based receiver. The metric is on a scale from 0 to 1, with lower scores being better. ViSQOL’s MOS is able to correlate to this metric, as seen in Figure 1.

In this use case, Meet developers were mostly interested in the sensitivity of ViSQOL to audio degradations from network impairments. In Figure 2 there is a comparison between mean ViSQOL scores during a call that shows that the metric is sensitive to how audio quality changed from a good network conditions scenario with scores ranging from 4.21 to 4.28, to a medium impaired scenario with scores ranging from 4.04 to 4.16, to finally an extremely challenging network scenario with scores from 3.72 to 3.94. Although the exact network conditions can not be shared, here good network conditions indicated that the connection should allow for both video and audio to be near perfect in the call, medium conditions indicate that the call might have issues, but the audio should continue to




Fig. 1. Hangouts Meet’s internal no-reference metric has components to detect audio degradations. ViSQOL successfully detected these degradations in audio that contained them (blocks 1 and 4), while in the audio blocks that were not affected the scores from ViSQOL were higher (blocks 2 and 3).

be good, while in extremely challenging conditions we expect to see both video and audio perceptually degraded, but the call would still go through.

Mean ViSQOL Per-Block Scores




Fig. 2. Comparison between mean ViSQOL MOS and network degradations. Some of the calls were run with good network conditions (green), some were simulating average network conditions, where the product should still perform well (yellow), while others were simulating extremely challenging network conditions, where it is expected to for issues to appear (red).

In order to ensure that ViSQOL performs reliably, several hundreds of calls were collected from the testbed. The mean values obtained from ViSQOL and the internal metric from these calls were plotted in Figure 3. The results were reliably reproduced. Following the positive results from this investigation, ViSQOL is currently one of the main objective audio quality metrics deployed by the Hangouts Meet product team at Google.

B. Opus Codec

Google contributes to the development of the Opus codec. ViSQOL and POLQA were used to benchmark the quality of the Opus coder for both speech and music at various bitrates and computational complexities. In previous studies ViSQOLAudio has been shown to perform reasonably on low bitrate audio [8]. However, ViSQOL’s speech mode did not specifically target the low bitrate case. Additionally, recent




Fig. 3. Scatter plot of ViSQOL MOS versus a no-reference internal metric. Each point represents a call.

advancements in Opus have pushed the lower bound of the range of bitrates further downwards for a given bandwidth since the time ViSQOL was introduced. For example, Opus 1.3 can produce a wideband signal at 9 kbps, whereas the TCDAudio14 [21], CoreSV14 [22], and AACvOpus15 [19] datasets that ViSQOL’s support vector regression was trained on have bitrates that only go as low as 24 kbps.

POLQA and the original version of ViSQOL in speech mode display similar trends that are consistent with expectations with respect to the bitrate and complexity settings. The differences in the *lower* bitrates are more pronounced according to POLQA. The differences in *higher* bitrates are more pronounced according to ViSQOL. Although subjective scores were not available, the developers expected that MOS should be less sensitive to changes in higher bitrates, giving POLQA a better match. After the improvements described in section 3, ViSQOL v3 MOS was a closer match to the expectation as can be seen in Figure 4.

For musical examples, the developers found that both metrics display similar trends with respect to bitrates. However, POLQA shows higher discrimination between 6-8 kbps, 10-12 kbps and 16-24 kbps. ViSQOL is able to discriminate between the different bitrates with monotonic behavior, but one point of concern is that this results in ViSQOLAudio being relatively insensitive to differences in complexity settings. In light of this, we would not recommend using ViSQOL for automated regression tests without retraining the model. The improvements made in section 3 slightly ameliorate these issues, as can be seen in Figure 5. On the other hand, ViSQOL identified a spurious bandwidth ‘bump’ at 12 kbps for the 5 and 6 complexity settings (which was perceived as higher quality in informal listening), where POLQA did not.

Lastly, ViSQOL was used to analyze the results for both clean and noisy references. This is not a case ViSQOL was designed for, as it presumes a clean reference, similar to PESQ [2] and POLQA [4]. However, it was found to perform in a similar fashion to the clean cases for both speech and audio in the noisy cases.

It was concluded that ViSQOL could be used for regression




Fig. 4. Estimated MOS for varying bitrates (6-64 kbps) and complexity settings on Opus-encoded speech.

testing for speech. However, formal listening tests would be desirable for two reasons: to better interpret the differences between POLQA and ViSQOLAudio, and to allow training a model that represented the low bitrate ranges.

C. Other Findings

A number of other teams have also adapted ViSQOL for their products. In the majority of cases, their use case vaguely resembles the training data (e.g. wideband speech network degradations or music coding), but often has marked differences. For example, one team chose to analyze the network loop with a digital and analog interface, requiring a rig to be built for continuous automated testing. Typically these teams also had access to PESQ, POLQA or subjective scores for their cases and wanted to evaluate the accuracy of ViSQOL measurements as well as identify limitations. A frequent issue was related to the duration and segmentation of the audio that would be used with ViSQOL when used in an automated framework. While ViSQOL in speech mode has a




Fig. 5. Estimated MOS for varying bitrates and complexity settings on Opus-encoded music (legend as per Fig. 4). The bitrates follow the same key as Figure 4. The bump at complexity 5, 6, and 10 for 12 kbps is related to Opus deciding to use a 12 kHz bandwidth for some fraction of the files instead of the 8 kHz bandwidth it used for complexities 2-4 and 7-9.

voice activity detector, it was found that ViSQOLAudio would perform poorly for segments where the reference was silent, because of either the averaging effects, or because of the lack of log-scale thresholding which was overly sensitive to small absolute differences in ambient noise levels. To resolve the averaging effects, it was recommended to extract segments of audio of 3 to 10 seconds where there was known activity. A solution to the thresholding issues is discussed in the next section.

III. DESIGN AND IMPROVEMENTS

This section summarizes the previous version and describes the changes made to the new version. Figure 6 shows the overall program flow and highlights the new components of the system that are referred to in the subsections.




Fig. 6. System Diagram. The inputs and outputs have white fill, and the processing components have blue fill. New components have thick edges. The dashed and dotted fill represent speech-only and audio-only components, respectively.

A. General Design

The ViSQOL algorithms described in [7] and [8] share many components by design, such as the gammatone spectrogram and NSIM calculation. It then seems reasonable that the common components be shared and developed together. The differences between the two algorithms are related to differences in the characteristics of speech and music. For example, the use of voice activity detection (VAD) for speech, and analysis of the higher bands (up to 24 kHz) for general audio/music. The common components of both speech and audio systems include creating a gammatone spectrogram using equivalent rectangular bandwidth (ERB) filters, creation of patches on the order of a half-second, aligning them, computing the NSIM from the aligned patches, and then mapping the NSIM values to MOS.

There were minor changes to some of these components because of practical reasons, such as modifying dependencies, or fixing issues found in case studies or test failures. For example, the VAD implementation uses a simple energy-based VAD, which should be sufficient given the requirement of clean references. As another example, window sizes were updated to be 80 ms with a hop of 20 ms after discovering an issue with the windowing of previous versions.

B. C++ Library and Binary

To make ViSQOL more available, we uncoupled the dependency on MATLAB by implementing a C++ version with only open source dependencies. The new version, v3, is available as a binary or as a library. The codebase was made available on GitHub because we wish for it to be easy to use by the public, and to invite external contributions.

The majority of users were binary users, but some had requirements for finer control. For this purpose we designed a library with protobuf support and error checking, which the binary depends on. This library would also be useful for a user that wishes to wrap the functions in a different language, such as with python bindings.

There were several changes to the input and output. Verbose output has also changed to include the average NSIM values per frequency band and mean NSIM per frame. Because ViSQOL is continuously changing to adapt to new problems, a conformance version number is included in the output. Whenever the MOS changes for known files, the conformance number will be incremented. Lastly, batch processing via comma-separated value (csv) files are also supported.

A number of Google-related projects were used to build this version. The application binary was implemented using the Abseil C++ application framework [23]. The Google Test C++ testing framework [24] was integrated and various tests were implemented to ensure correctness, detect regressions, and increase stability for edge cases. 23 test classes with multiple tests were implemented. These include not only unit tests, but also a test to check the conformance of the current version to known scores. The Bazel framework [25] was used to handle building and dependency fetching, as well as test development.

C. Fine-scaled Time Alignment

Although the previous versions of ViSQOL did two levels of alignment (global and patch), there were still issues with the patch alignment due to the spectrogram frames being misaligned at a fine scale. To address this, we implemented an additional alignment step that offsets by the lag found in a cross correlation step on the time-domain regions that corresponds to the aligned patches as described in [26]. Next, the gammatone spectrogram is recomputed for sample-aligned patch audio and the NSIM score is taken.

D. Silence Thresholds

To deal with problem of log-scale amplitudes discussed in II-C, we introduce silence thresholds on the gammatone

spectrogram. Because NSIM is calculated on log-amplitudes, we found that it was too sensitive to different levels of ambient noise. For example, a near-digital silence reference compared against a very low level of ambient noise would still have a very low NSIM score, despite being perceptually transparent. The silence threshold introduces an absolute floor as well as a relative floor that may be higher for high amplitude frames.

The thresholded amplitude $y_{t,f}(x)$ for a time t and frequency band f given an input spectrogram x is subject to:

$$y_{t,f}(x) = \max(Y_{min}, Y_{fmin}(t), x(t, f)) \quad (1)$$

where:

$$Y_{fmin}(t) = \max(r_f(t), d_f(t)) - Y_{fmin} \quad (2)$$

given reference and degraded log amplitudes $r_{t,f}$ and $d_{t,f}$, and global absolute threshold Y_{min} , and relative per-frame threshold Y_{fmin} .

E. NSIM to MOS Model

The changes above ultimately affect the NSIM scores. This requires that a new SVR model is trained to map the frequency band NSIM to MOS using libsvm [27]. We conducted a grid search to minimize the 4-way cross validation loss on the same training set (TCDAudio14, CoreSV14, AACvOpus15). However, we observed in Section II that this model was too specific to the training data and would behave poorly on very low bitrate (6-18 kbps) audio. This appears to be related to the fact that there is no monotonicity constraint in the SVR model used by ViSQOL (a strictly higher NSIM for out of distribution data produced lower MOS). To address this issue for the default model, we relaxed the SVR parameters by lowering the cost and gamma parameters to have a slightly higher cross validation error while providing behavior that was closer to monotonic behavior.

Additionally, this version includes some tooling and documentation that allows for users to train their own SVR model by the use of CSV input files if the user can provide subjective scores for degraded/reference pairs. By following the grid search methods described by libsvm authors, users should be able to tailor a model that is able to represent their data.

IV. DISCUSSION

Here we present a discussion of the use cases and feedback in light of the improvements. This is followed by reflection on trends and the areas that are promising as future work.

The case studies mentioned in Section II highlight the challenges with real world applications of ViSQOL. The findings are generally that ViSQOL can be used for various applications, but careful investigation is required for any use case. The users of these tools are the very developers of new audio processing and coding techniques, and are often analyzing new types of audio that is “out of distribution”. In some cases, we can allow the user to retrain a model to match the new data.

We find that developers are reasonably skeptical about how well ViSQOL will apply to their problem, given that it almost

always has unique characteristics. Although ViSQOL is not guaranteed to give a meaningful absolute MOS for cases that are significantly different from what it was originally designed with, the developers in our case studies found some correlation that was useful for their use case. However, this conclusion is often facilitated by the use of additional metrics that can be used to validate ViSQOL's application.

In other cases, for example, in the generative case, it is possible that it requires a redesign of the algorithm at a fundamental level, which could include different spectrogram representations or DNNs. Projects like LibriTTS [28] have curated large amounts of freely available speech data, which has been a boon to speech-related DNNs, there is yet no standard and widely available subjective score dataset that is of similar scale. A larger dataset would enable new development, but also require rethinking of existing tools, such as support vector regression, used by ViSQOLAudio, which is intended for use on smaller datasets on the order of hundreds of points.

V. CONCLUSION

We have presented a new version of ViSQOL which is available for use on GitHub. The integration to real world problems by different teams at Google yielded a number of insights and improvements to the previous version. There are a number of promising avenues for future work, including DNN based approaches, a more general model, and taking the new generative audio approaches into account.

ACKNOWLEDGMENTS

We acknowledge Colm Sloan for his work on the initial C++ version of ViSQOL. 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] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (PESQ)-a new method for speech quality assessment of telephone networks and codecs," in *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No. 01CH37221)*, vol. 2. IEEE, 2001, pp. 749–752.
- [2] 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," Int. Telecomm. Union, Geneva, Switzerland, ITU-T Rec. P.862, 2001.
- [3] J. G. Beerends, C. Schmidmer, J. Berger, M. Obermann, R. Ullmann, J. Pomy, and M. Keyhl, "Perceptual objective listening quality assessment (POLQA), the third generation ITU-T standard for end-to-end speech quality measurement part I," *Journal of the Audio Engineering Society*, vol. 61, no. 6, pp. 366–384, 2013.
- [4] ITU, "Perceptual objective listening quality assessment," Int. Telecomm. Union, Geneva, Switzerland, ITU-T Rec. P.863, 2018.
- [5] T. Thiede, W. C. Treurniet, R. Bitto, C. Schmidmer, T. Sporer, J. G. Beerends, and C. Colomes, "PEAQ - The ITU standard for objective measurement of perceived audio quality," *Journal of the Audio Engineering Society*, vol. 48, no. 1/2, pp. 3–29, 2000.
- [6] R. Huber and B. Kollmeier, "PEMO-Q-A new method for objective audio quality assessment using a model of auditory perception," *IEEE Transactions on audio, speech, and language processing*, vol. 14, no. 6, pp. 1902–1911, 2006.
- [7] A. Hines, J. Skoglund, A. Kokaram, and N. Harte, "ViSQOL: The virtual speech quality objective listener," in *IWAENC 2012; International Workshop on Acoustic Signal Enhancement*. VDE, 2012, pp. 1–4.
- [8] A. Hines, E. Gillen, D. Kelly, J. Skoglund, A. Kokaram, and N. Harte, "ViSQOLAudio: An objective audio quality metric for low bitrate codecs," *The Journal of the Acoustical Society of America*, vol. 137, no. 6, pp. EL449–EL455, 2015.
- [9] M. Narbutt, A. Allen, J. Skoglund, M. Chinen, and A. Hines, "AMBI-QUAL - A full reference objective quality metric for ambisonic spatial audio," in *2018 Tenth International Conference on Quality of Multimedia Experience (QoMEX)*. IEEE, 2018, pp. 1–6.
- [10] W. B. Kleijn, F. S. Lim, A. Luebs, J. Skoglund, F. Stimberg, Q. Wang, and T. C. Walters, "WaveNet based low rate speech coding," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 676–680.
- [11] 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)*. IEEE, 2019, pp. 7155–7159.
- [12] 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.
- [13] A. Hines and N. Harte, "Speech intelligibility prediction using a neurogram similarity index measure," *Speech Communication*, vol. 54, no. 2, pp. 306 – 320, 2012.
- [14] A. R. Avila, H. Gamper, C. Reddy, R. Cutler, I. Tashev, and J. Gehrke, "Non-intrusive speech quality assessment using neural networks," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 631–635.
- [15] H. Gamper, C. K. Reddy, R. Cutler, I. J. Tashev, and J. Gehrke, "Intrusive and non-intrusive perceptual speech quality assessment using a convolutional neural network," in *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2019, pp. 85–89.
- [16] A. Hines, J. Skoglund, A. Kokaram, and N. Harte, "Robustness of speech quality metrics to background noise and network degradations: Comparing ViSQOL, PESQ and POLQA," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2013, pp. 3697–3701.
- [17] A. Hines, E. Gillen, and N. Harte, "Measuring and monitoring speech quality for Voice over IP with POLQA, ViSQOL and P. 563," in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [18] C. Sloan, N. Harte, D. Kelly, A. C. Kokaram, and A. Hines, "Bitrate classification of twice-encoded audio using objective quality features," in *2016 Eighth International Conference on Quality of Multimedia Experience (QoMEX)*. IEEE, 2016, pp. 1–6.
- [19] —, "Objective assessment of perceptual audio quality using ViSQOLAudio," *IEEE Transactions on Broadcasting*, vol. 63, no. 4, pp. 693–705, 2017.
- [20] (2011) WebRTC. [Online]. Available: webrtc.org
- [21] A. Hines, E. Gillen, D. Kelly, J. Skoglund, A. Kokaram, and N. Harte, "Perceived audio quality for streaming stereo music," in *Proceedings of the 22nd ACM international conference on Multimedia*. ACM, 2014, pp. 1173–1176.
- [22] CoreSV Team. (2014) CoreSV listening test. [Online]. Available: <http://listening-test.coresv.net/results.htm>
- [23] Google. (2017) Abseil. [Online]. Available: <https://abseil.io/>
- [24] —. (2008) Google test framework. [Online]. Available: <https://github.com/google/googletest>
- [25] —. (2015) Bazel build system. [Online]. Available: <https://bazel.build/>
- [26] T. Mo and A. Hines, "Jitter buffer compensation in Voice over IP quality estimation," in *2019 30th Irish Signals and Systems Conference (ISSC)*. IEEE, 2019, pp. 1–6.
- [27] C.-C. Chang and C.-J. Lin, "Libsvm: A library for support vector machines," *ACM transactions on intelligent systems and technology (TIST)*, vol. 2, no. 3, pp. 1–27, 2011.
- [28] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, "LibriTTS: A corpus derived from librispeech for text-to-speech," *arXiv preprint arXiv:1904.02882*, 2019.