



Benchmarking Neural Speech Codec Intelligibility with SITool

Anna Leschanowsky¹, Kishor Kayyar Lakshminarayana¹, Anjana Rajasekhar¹, Lyonel Behringer¹,
Ibrahim Kilinc², Guillaume Fuchs¹, Emanuël A. P. Habets³

¹Fraunhofer Institute for Integrated Circuits (IIS), Erlangen, Germany

²Department of Electrical and Computer Engineering, University of California, San Diego, USA

³International Audio Laboratories Erlangen, Germany

anna.leschanowsky@iis.fraunhofer.de

Abstract

Speech intelligibility assessment is essential for evaluating neural speech codecs, yet most evaluation efforts focus on overall quality rather than intelligibility. Only a few publicly available tools exist for conducting standardized intelligibility tests, like the Diagnostic Rhyme Test (DRT) and Modified Rhyme Test (MRT). We introduce the Speech Intelligibility Toolkit for Subjective Evaluation (SITool), a Flask-based web application for conducting DRT and MRT in laboratory and crowdsourcing settings. We use SITool to benchmark 13 neural and traditional speech codecs, analyzing phoneme-level degradations and comparing subjective DRT results with objective intelligibility metrics. Our findings show that, while neural speech codecs can outperform traditional ones in subjective intelligibility, only STOI and ESTOI – not WER – significantly correlate with subjective results, although they struggle to capture gender and wordlist-specific variations observed in subjective evaluations.

Index Terms: speech coding, speech intelligibility, subjective evaluation, Rhyme Test

1. Introduction

Neural speech codecs have gained increasing interest for their ability to outperform traditional coding schemes, especially at very low bitrates [1]. Unlike traditional approaches, they use data-driven methods that avoid costly tuning via listening tests [1] but may hallucinate content due to their generative nature. Thus, performance evaluation remains essential. Evaluation can be subjective (e.g., P.800 [2], MUSHRA [3]) or objective (e.g., VISQOL [4], ESTOI [5, 6, 7]). A recent subjective study utilized Degradation Category Ratings (DCR) to benchmark neural codecs [8]. While P.800 and its crowdsourcing variant P.808 are widely used, they assess overall speech quality rather than speech intelligibility. Yet, as neural speech codecs are increasingly deployed in digital communication, broadcasting, and media streaming, intelligibility becomes a crucial evaluation criterion. Speech intelligibility, a distinct dimension of speech quality, is not always strongly correlated with overall quality [9]. Consequently, dedicated intelligibility assessments offer deeper insights into codec performance.

Speech intelligibility can be assessed at phoneme, word, or sentence level. At the sentence level, Benoît et al. [10] proposed a Semantically Unstable Sentence (SUS) test, which measures intelligibility through Word Error Rates (WER) but can be unreliable due to participants' spelling errors. Word-level tests such as the Phonetically Balanced Word Test (PBWT), Diagnostic Rhyme Test (DRT), and Modified Rhyme Test (MRT) have been standardized by the American National Standards Institute (ANSI) [11]. PBWT is an open-set test requiring trans-

cription of a word, while DRT and MRT are closed-set tests where listeners select the heard word from a set of two (DRT) or six (MRT) options. Closed-set tests are easier to administer, require fewer listeners, and show high reliability, whereas open-set tests are less prone to ceiling effects [12]. The DRT consists of 96 rhyming word pairs that differ only in their initial or final consonant sounds [13]. A key advantage of DRT is that it allows for detailed analysis of phoneme-specific degradations by providing diagnostic feature scores across six categories [13].

Crowdsourced listening tests are widely adopted for speech quality [14] and, more recently, for intelligibility evaluation [15]. While open-source tools for speech quality assessment exist (e.g., WebMUSHRA [16], HultiGen-v2¹), open-source DRT and MRT software toolkits remain limited or outdated [17]. For instance, the scripts for deploying DRT, provided in [15], are restricted to Qualtrics and do not support MRT, whereas the Python-based MRT tool provided by NIST is not optimized for crowdsourced test environments².

To fill this gap, we introduce SITool – a Flask-based web application for DRT and MRT – and benchmark 13 speech codecs with English male and female talkers. An overview of benchmarked codecs is given in Section 2, and the tool is described in Section 3. While subjective tests remain the gold standard, the correlation between objective intelligibility metrics (e.g., STOI, ESTOI, and WER [5, 18, 19]) and subjective DRT scores for neural speech codecs is underexplored. Therefore, we present subjective and objective results and analyze their relationship in Section 4.

2. Benchmarked Speech Codecs

We evaluate the intelligibility of various publicly available state-of-the-art neural codecs and, for a full performance evaluation, compare them with prominent traditional speech codecs.

In this work, we consider **LPCNet**³ [20] at 1.6 kbps, an early solution that combines classical signal processing and deep neural networks to decode a bitstream generated by a conventional encoder. The hybrid solution allows an autoregressive neural vocoder to run in real-time on a CPU (3 GFLOPs). Further, we evaluate different GAN-based end-to-end autoencoder approaches, a paradigm introduced by SoundStream [21] built on the discrete learned representation of VQ-VAE [22]. This principle is adopted by most state-of-the-art neural speech codecs. We test **Lyra V2**⁴ [23], an open-source codec derived

¹<https://github.com/APL-Huddersfield/HULTI-GENv2/>

²<https://github.com/usnistgov/MRT>

³<https://github.com/xiph/LPCNet/commit/7dc9942>

⁴<https://github.com/google/lyra/tree/v1.3.2>

directly from the original SoundStream model and optimized to run in real-time on a smartphone CPU. Various improvements to the paradigm have been proposed by adding recurrent layers [24], increasing the model size [7], or carefully selecting the training data [25]. It has also been proposed to modify them by exploiting the strong generative capacity of diffusion models [26, 27], or by semantically disentangling the latent space to achieve even lower bitrates [27, 28]. Transformers and attention layers have also been considered to get an even more compact representation [28, 29]. Among them, we evaluate two extensions of DAC [7], namely a speech fine-tuned version, **DACS** [25], at 1.5 kbps, and the multi-scale quantization **SNAC** [30], at 980 bps. **Mimi** [28], evaluated at 0.55 and 1.1 kbps, is characterized by transformer bottlenecks before and after quantization, as well as the inclusion of semantic tokens via distillation. We also include the diffusion-based **SemantiCodec** [27], which combines acoustic and semantic encoders. With 507M parameters, it is the largest of the codecs evaluated, allowing for very low bitrates (340 and 680 bps).

For comparison purposes, three generations of traditional 3GPP communication codecs, **AMR** [31], **AMR-WB** [32], and **EVS** [33] are included in the study at 4.75, 6.6 and 8 kbps, respectively. They are all based on the popular CELP paradigm, a hybrid approach aiming at preserving the waveform through an analysis-by-synthesis optimization of its coded parameters. We employ them at, or near, their minimum bitrate. In addition, the purely parametric **Codec2** [34] speech coder is considered at its minimum bitrate of 0.7 and at 2.4 kbps.

3. SITool - Speech Intelligibility Toolkit for Subjective Evaluation

We developed a standalone Flask-based web application that supports both DRT and MRT testing and can be deployed in laboratory settings or on crowdsourcing platforms like Amazon Mechanical Turk (AMT)⁵. Our tool allows the integration of consent forms, language proficiency assessments, predefined trap questions, and enforced breaks through an easily customizable configuration file. The application outputs detailed results files, which automatically analyze for accuracy across test conditions and, in the case of DRT, across acoustic features. Additionally, the toolkit includes scripts for further analysis.

3.1. Word Lists, Speech Samples and Test Blocks

For benchmarking, we used English word lists from [13] and audio stimuli provided by [15]. Testing was done at 16 kHz sampling rate. Audio for codecs operating at a different sampling rate was resampled before coding and converted back to 16 kHz afterward⁶. To reduce test time, we randomly selected male and female talker samples from the original dataset. Importantly, talkers varied only within and across word lists but remained constant across codecs. Word pairs were divided into four sessions [13], resulting in a total of eight sessions.

3.2. Trap Questions, Gold Standard and Lower Anchor

Following best practices for crowdsourced listening tests [14], the DRT included trap questions, the original signal as a gold standard, and a lower anchor for post-screening purposes. Trap

⁵We make the framework publicly available at <https://github.com/audio4labs/SITool>

⁶Some examples are provided at <https://s.fhg.de/sitoolweb>




Figure 1: Graphical user interface for the DRT test

questions consisted of two phonetically distinct words from the DRT word list to ensure participants’ attentiveness. The clean speech signal served as a reference for codec performance and ensured that participants could reliably identify clean recordings. The lower anchor was designed as a control condition for unintelligibility and reliability of the interface to ensure no unintended cues helped with word recognition. Therefore, we synthesized words without their distinguishing consonants using a synthesis system consisting of an acoustic model based on ForwardTacotron⁷, with extensions similar to [35] and StyleMelGAN as neural vocoder [36]. To entirely obscure the incomplete word and ensure no residual information could help to distinguish between words, we added cough noise from [37] to the location where the consonant was removed.

3.3. Test Procedure

Prior to the test, informed consent was obtained, and demographic information on age, gender, and self-assessed language proficiency was collected. In addition, to ensure a sufficient level of English language proficiency, participants completed a B1/B2 Level test based on official Cambridge English Listening Test examples⁸. To qualify for the main test, they had to answer at least four out of five questions correctly and were instructed to use headphones. After passing, participants completed a trial run to familiarize themselves with the test interface (see Figure 1). The main test included 15 conditions from either male or female talkers and was split into two parts of roughly 40 minutes with a mandatory five-minute break in between.

3.4. Participants and Post-Screening

Participants were recruited from AMT and an internal listener pool with an average hourly compensation of \$10.50 USD, aligning with prior crowdsourcing studies in this context [15]. For the final sample, we excluded participants who failed any trap questions but not based on their lower anchor performance, as no outliers were identified. The final sample consisted of 82 participants, unevenly distributed across the eight tests, but with at least eight participants per test as recommended by [13]. Of these, 54.88% identified as male, 43.90% as female, and 1.22% as other. The sample was diverse in age ($M=43.74$, $SD=10.68$), ranging from 23 to 70 years. 66 participants were native English speakers (US: 56, UK: 3, Other: 7), 16 non-native speakers.

3.5. Score Calculation

Intelligibility scores were calculated according to [13] and adjusted for guessing with R being the number of correct answers, W being the number of incorrect answers: $P(c) = \frac{R-W}{R+W} * 100$.

⁷<https://github.com/as-ideas/ForwardTacotron>

⁸<https://www.cambridgeenglish.org/learning-english/activities-for-learners/>

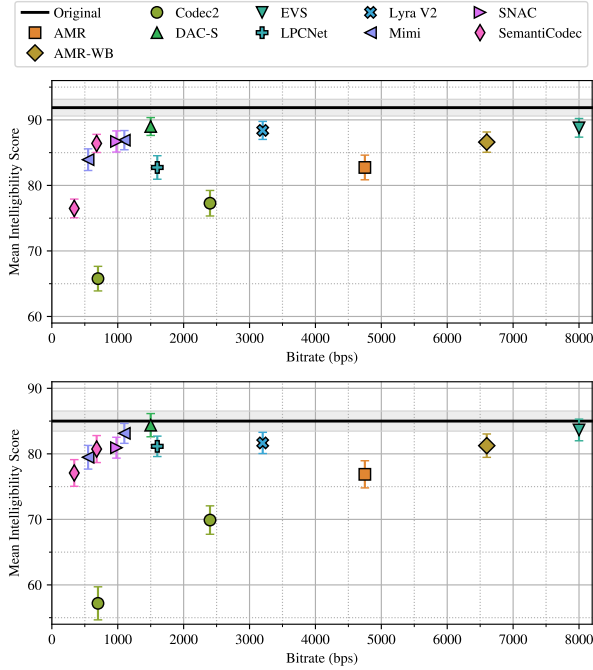


Figure 2: Mean subjective intelligibility scores per talker gender, with confidence intervals by bitrate (top: male, bottom: female). The low anchor (mean of 2.98 for male and -1.46 for female) is omitted for clarity.

4. Results

4.1. Subjective Speech Intelligibility Evaluation

Subjective intelligibility results for male and female talkers are shown in Figure 2. The results show that the female reference signal ($M=85$, $SD=9.66$) has a noticeably lower intelligibility score than the male reference signal ($M=91.87$, $SD=8.23$). This pattern can be observed across most codecs, except for SemantiCodec at 340 bps, where female talkers slightly outperformed males (77.8% vs. 76.49%). Moreover, despite operating at a lower bitrate, DAC-S slightly outperformed EVS for both genders. The lower anchor results ($M=0.81$ $SD=10.52$) confirmed its effectiveness as an unintelligibility benchmark and are excluded from further analysis.

To assess performance differences across codec conditions and talker gender, we employed a Linear Mixed-Effects Model (LMM) with fixed effects for codec condition (14 levels), talker gender (2 levels), and word list (4 levels) and random effects for individual participants to account for within-subject variability. A likelihood-ratio test showed that adding interaction effects for all factors significantly improved the model fit compared to a purely additive model ($\chi^2(94)=217.48$, $p<.0001$). Visual inspection of residual plots showed no clear deviations from homoscedasticity or normality. The model’s fixed effect intercept was 87.5 [82.32, 92.68]. A repeated measures ANOVA (type III, Satterthwaite approximation of degrees-of-freedom) on a restricted maximum likelihood estimation (REML) re-fitted model revealed significant main effects for codec condition ($F(13, 1065.68)=98.10$, $p<.0001$), talker gender ($F(1, 79.71)=12.59$, $p<.0001$), and wordlist ($F(3, 89.65)=17.31$, $p<.0001$) as well as significant interaction effects of condition-gender ($F(13, 1065.68)=3.035$, $p<.0001$), condition-wordlist

($F(39, 1065.68)=2.96$, $p<.0001$) and gender-wordlist ($F(3, 89.65)=16.30$, $p<.0001$). Yet, wordlist differences might not be reliably detectable with the current sample and should be confirmed with a larger one. Post-hoc tests using “emmeans”⁹ showed that, when averaged over gender and word lists, all codecs except EVS and DAC-S performed significantly worse than the reference ($p<.01$). However, interaction effects revealed that most codecs differed significantly from the reference for male talkers, whereas only Codec 2, AMR, Mimi at 550 bps and SemantiCodec at 340 bps scored significantly lower for females, likely due to consistently lower intelligibility scores for the female references. Although the crowdsourced audio samples were carefully curated [15], these gender-specific differences for the reference signal warrant further investigation. Future studies should explicitly control for talker-specific effects and assess the potential impact of consumer recording devices on crowdsourced recordings for intelligibility assessment [38].

Further, we found that intelligibility scores for male talkers outperformed those of female talkers on Wordlist 1, 2 and 4 for all codecs but not on Wordlist 3. Wordlist 4 yielded highest intelligibility scores for both genders and all codecs, whereas Wordlist 2 performed worst for females and second worst for male talkers. These gender-specific differences already exist at the reference level and are not due to codec performances. Even when using curated datasets as in this test, it is essential to include reference signals and possibly conduct pre-tests on the intelligibility of reference signals. A heatmap (Figure 3) shows that the distinctive feature “graveness” in Wordlist 2 had notably low scores across codecs, including the reference. Closer inspection revealed that low performance was linked to high error rates for word pairs ‘TROUGH-TROTH’ and ‘FIN-THIN’, which repeatedly test the consonants /f/ and /θ/. Similarly, Wordlist 4 showed more errors for the word pair ‘REEF-WREATH’, with little impact on average performance. We did not observe similar patterns for word pairs of “graveness” including consonants /f/ and /θ/ in Wordlist 3, suggesting that both phoneme characteristics and individual pronunciation influenced the results.

4.2. Objective Speech Intelligibility Evaluation

We show objective results¹⁰ in Table 1. First, we observe that WER is rather high even for the reference signals, which might indicate that single-word utterances pose a challenge to Whisper. Across all codecs, female ESTOI scores are at least 0.04 higher than male ones, while no such trend is visible for STOI and WER. Looking at traditional and neural codecs separately, EVS and DAC-S show the best results. However, while DAC-S outperforms EVS and AMR-WB in the subjective tests, it surpasses EVS in all objective metrics and AMR-WB in WER. Further, SemantiCodec at 680 bps ranks below Mimi at 550 bps w.r.t. STOI and ESTOI, differing from the subjective results.

Performance differences were evaluated across codec conditions, gender, and word lists using two linear regression models. For STOI and ESTOI, additive models were chosen over interaction models because they explained equivalent variance ($R^2_{STOI} = 0.79, R^2_{ESTOI} = 0.71$) while showing lower AIC and BIC values. A Type III ANOVA on the STOI model revealed significant effects of codec condition ($F(13) = 1349.2$, $p<.0001$) and word list ($F(3) = 28.82$, $p<.0001$), but not gender. In contrast to subjective results, pairwise com-

⁹<https://rvlenth.github.io/emmeans/>

¹⁰WER calculated using Whisper “large” from <https://github.com/openai/whisper/>.

Table 1: List of tested codecs, along with intelligibility objective metrics. 45 items were too short for STOI/ESTOI computation and were excluded. Confidence intervals are omitted since they are almost negligible, all under 0.015.

Codec	Params (M)	kHz	kbps	STOI f/m (↑)	ESTOI f/m (↑)	WER f/m (↓)
Original	-	-	-	1.00 / 1.00	1.00 / 1.00	0.25 / 0.19
Codec2 [34]	-	8	0.7	0.57 / 0.60	0.44 / 0.39	0.86 / 0.81
Codec2 [34]	-	8	2.4	0.75 / 0.74	0.63 / 0.52	0.54 / 0.54
AMR [31]	-	8	4.75	0.87 / 0.88	0.80 / 0.75	0.43 / 0.39
AMR-WB [32]	-	16	6.6	0.90 / 0.90	0.83 / 0.78	0.30 / 0.28
EVS [33]	-	16	8	0.95 / 0.94	0.91 / 0.87	0.29 / 0.30
LPCNet [20]	0.073	16	1.6	0.78 / 0.80	0.67 / 0.60	0.46 / 0.55
Lyra V2 [23]	n/a	16	3.2	0.90 / 0.89	0.82 / 0.75	0.38 / 0.35
DAC-S [25]	76	24	1.5	0.93 / 0.92	0.86 / 0.79	0.34 / 0.31
SNAC [30]	20	24	0.98	0.89 / 0.89	0.80 / 0.72	0.44 / 0.47
Mimi [28]	80	24	0.55	0.85 / 0.86	0.73 / 0.65	0.51 / 0.47
Mimi [28]	80	24	1.1	0.91 / 0.91	0.83 / 0.76	0.35 / 0.35
SemantiCodec [27]	507	16	0.34	0.79 / 0.80	0.64 / 0.55	0.68 / 0.69
SemantiCodec [27]	507	16	0.68	0.84 / 0.85	0.72 / 0.65	0.49 / 0.44

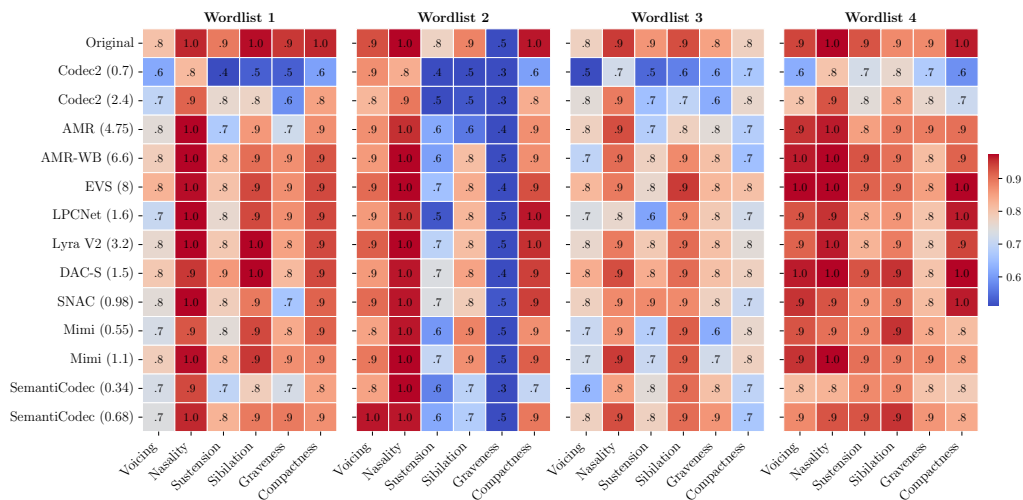


Figure 3: Accuracy heatmap for codecs across distinctive features and wordlists as outlined in [13]. The x-axis shows the distinctive features and the y-axis shows the codecs with their respective bitrates (kbps) in brackets.

parisons indicated that all codecs were significantly different from the reference, with Wordlist 4 showing the lowest and Wordlist 2 and 3 the highest performances. For ESTOI, condition ($F(13)=861.15, p<.0001$), gender ($F(1)=564.89, p<.0001$), and word list ($F(3)=11.49, p<.0001$) all significantly affected scores. However, unlike subjective results, female talkers obtained higher scores across all word lists and codec conditions than male talkers. Finally, fitting a model to WER yielded results similar to those for STOI, indicating that objective measures are misaligned with subjective performance when considering subtle differences between gender and word lists.

4.3. Correlation Analysis

We conducted a correlation analysis on the mean subjective and objective results disaggregated by codecs, wordlists, and talker gender. The analysis showed that the correlation between subjective and objective metrics – STOI and ESTOI – was considerably higher when averaging over gender and wordlist. For STOI, Pearson’s r increased from 0.595 to 0.958, and for ESTOI from 0.499 to 0.890. This, in combination with the previous analysis, suggests that the objective metrics do not fully capture the variability of talker gender and wordlist obtained in the DRT. In contrast, subjective results and WER were uncor-

related (non-averaged: -0.11, averaged: -0.15), questioning the usefulness of WER as an objective metric for speech intelligibility in the context of DRT.

5. Conclusions

We benchmarked 13 speech codecs, from traditional to state-of-the-art GAN-based codecs, on intelligibility using the newly developed *SITool*. Unlike speech quality evaluation, DRT offered a more detailed analysis of phoneme-specific degradations and distinctive acoustic features. Our results show that below 1100 bps, the intelligibility of neural codecs degrades compared to the uncoded signal, leaving room for potential improvement. Notably, the female reference signal yielded overall lower intelligibility scores, which may have impacted the evaluation of codecs and calls for further investigation of audio material. While objective measures, i.e., STOI and ESTOI, correlate well with subjective intelligibility when averaged over talker gender and word lists, they fail to capture finer variations related to gender and wordlist differences identified in the subjective evaluation, and WER showed no correlation with subjective results. As neural codecs improve, subjective tests remain crucial for detecting subtle impairments. Given the observed ceiling effects, future work should explore open-set tests like SUS, less prone to ceiling effects and not reliant on wordlists [12].

6. Acknowledgements

This research was partially supported by the Free State of Bavaria in the DSGenAI project and by the Fraunhofer-Zukunftsstiftung.

7. References

- [1] M. Kim and J. Skoglund, "Neural speech and audio coding: Modern AI technology meets traditional codecs [special issue on model-based and data-driven audio signal processing]," *IEEE Signal Processing Magazine*, vol. 41, no. 6, pp. 85–93, 2025.
- [2] ITU-T Rec. P.800, "Methods for subjective determination of transmission quality," 1996.
- [3] International Telecommunication Union, "ITU-R Recommendation BS.1534-3: Method for the subjective assessment of intermediate quality level of audio systems," October 2015.
- [4] M. Chinen, F. S. Lim, J. Skoglund, N. Gureev, F. O’Gorman, and A. Hines, "ViSQOL v3: An open source production ready objective speech and audio metric," in *2020 twelfth international conference on quality of multimedia experience (QoMEX)*. IEEE, 2020, pp. 1–6.
- [5] J. Jensen and C. H. Taal, "An algorithm for predicting the intelligibility of speech masked by modulated noise maskers," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 11, pp. 2009–2022, 2016.
- [6] A. Brendel, N. Pia, K. Gupta, L. Behringer, G. Fuchs, and M. Multrus, "Neural speech coding for real-time communications using constant bitrate scalar quantization," *IEEE Journal of Selected Topics in Signal Processing*, 2024.
- [7] R. Kumar, P. Seetharaman, A. Luebs, I. Kumar, and K. Kumar, "High-fidelity audio compression with improved RVQGAN," in *Advances in Neural Information Processing Systems*, vol. 36. Curran Associates, Inc., 2023, pp. 27980–27993.
- [8] T. Muller, S. Ragot, L. Gros, P. Philippe, and P. Scalart, "Speech quality evaluation of neural audio codecs," in *Proc. Interspeech*, 2024.
- [9] J. E. Preminger and D. J. V. Tasell, "Quantifying the relation between speech quality and speech intelligibility," *Journal of Speech, Language, and Hearing Research*, vol. 38, no. 3, pp. 714–725, 1995.
- [10] C. Benoît, M. Grice, and V. Hazan, "The SUS test: A method for the assessment of text-to-speech synthesis intelligibility using semantically unpredictable sentences," *Speech communication*, vol. 18, no. 4, pp. 381–392, 1996.
- [11] Acoustical Society of America, "ANSI/ASA S3.2-2020 method for measuring the intelligibility of speech over communication systems," 2020.
- [12] A. Schmidt-Nielsen, "Intelligibility and acceptability testing for speech technology," *Applied speech technology*, pp. 194–231, 1995.
- [13] ITU-T Rec. P.807, "Subjective test methodology for assessing speech intelligibility," 2016.
- [14] ITU-T Rec. P.808, "Subjective evaluation of speech quality with a crowdsourcing approach," 2021.
- [15] L. Lechler and K. Wojcicki, "Crowdsourced multilingual speech intelligibility testing," in *ICASSP*, 2024, pp. 1441–1445.
- [16] M. Schoeffler, F.-R. Stöter, B. Edler, and J. Herre, "Towards the next generation of web-based experiments: A case study assessing basic audio quality following the ITU-R recommendation BS.1534 (MUSHRA)," in *Proc. Web Audio Conference*, Paris, France, January 2015.
- [17] R. J. Calistri and H. J. Kallman, "Drt/***: Programs to administer and score the diagnostic rhyme test," *Behavior Research Methods, Instruments, & Computers*, vol. 18, no. 1, pp. 57–58, 1986.
- [18] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, "A short-time objective intelligibility measure for time-frequency weighted noisy speech," in *ICASSP*, 2010, pp. 4214–4217.
- [19] M. Karbasi and D. Kolossa, "Asr-based speech intelligibility prediction: A review," *Hearing Research*, vol. 426, p. 108606, 2022.
- [20] J.-M. Valin and J. Skoglund, "LPCNet: Improving neural speech synthesis through linear prediction," in *ICASSP*. IEEE, 2019, pp. 5891–5895.
- [21] N. Zeghidour, A. Luebs, A. Omran, J. Skoglund, and M. Tagliasacchi, "Soundstream: An end-to-end neural audio codec," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 495–507, 2022.
- [22] A. van den Oord, O. Vinyals, and k. kavukcuoglu, "Neural discrete representation learning," in *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc., 2017.
- [23] W. B. Kleijn, A. Storus, M. Chinen, T. Denton, F. S. C. Lim, A. Luebs, J. Skoglund, and H. Yeh, "Generative speech coding with predictive variance regularization," in *ICASSP*, 2021, pp. 6478–6482.
- [24] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High fidelity neural audio compression," *Transactions on Machine Learning Research*, 2023.
- [25] S. Shechtman and A. Dekel, "Low bitrate high-quality rvqgan-based discrete speech tokenizer," in *Interspeech 2024*, 2024, pp. 4174–4178.
- [26] H. Yang, I. Jang, and M. Kim, "Generative de-quantization for neural speech codec via latent diffusion," in *ICASSP*, 2024, pp. 1251–1255.
- [27] H. Liu, X. Xu, Y. Yuan, M. Wu, W. Wang, and M. D. Plumbley, "Semanticodec: An ultra low bitrate semantic audio codec for general sound," *IEEE Journal of Selected Topics in Signal Processing*, pp. 1–14, 2024.
- [28] A. Défossez, L. Mazaré, M. Orsini, A. Royer, P. Pérez, H. Jégou, E. Grave, and N. Zeghidour, "Moshi: A speech-text foundation model for real-time dialogue," 2024.
- [29] A. Siahkoobi, M. Chinen, T. Denton, W. B. Kleijn, and J. Skoglund, "Ultra-low-bitrate speech coding with pretrained transformers," in *Interspeech 2022*, 2022, pp. 4421–4425.
- [30] H. Siuzdak, F. Grötschla, and L. A. Lanzendörfer, "SNAC: Multiscale neural audio codec," in *Audio Imagination: NeurIPS Workshop AI-Driven Speech, Music, and Sound Generation*, 2024.
- [31] E. Ekudden, R. Hagen, I. Johansson, and J. Svedberg, "The adaptive multi-rate speech coder," in *IEEE Workshop on Speech Coding Proceedings. Model, Coders, and Error Criteria (Cat. No. 99EX351)*, 1999, pp. 117–119.
- [32] B. Bessette, R. Salami, R. Lefebvre, M. Jelinek, J. Rotola-Pukkila, J. Vainio, H. Mikkola, and K. Jarvinen, "The adaptive multi-rate wideband speech codec (AMR-WB)," *IEEE Transactions on Speech and Audio Processing*, vol. 10, no. 8, pp. 620–636, 2002.
- [33] S. Bruhn, H. Pobloth, M. Schnell, B. Grill, J. Gibbs, L. Miao, K. Jarvinen, L. Laaksonen, N. Harada, N. Naka, S. Ragot, S. Proust, T. Sanda, I. Varga, C. Greer, M. Jelínek, M. Xie, and P. Usai, "Standardization of the new 3GPP EVS codec," in *ICASSP*, 2015, pp. 5703–5707.
- [34] D. Rowe, "Codec 2 - open source speech coding at 2400 bits/s and below," *TAPR and ARRL 30th Digital Communications Conference*, pp. 80–84, 2011.
- [35] F. Zalkow, P. Sani, M. Fast, J. Bauer, M. Joshaghani, K. Kayyar, E. A. P. Habets, and C. Dittmar, "The AudioLabs system for the Blizzard Challenge 2023," in *Proceedings of the Blizzard Challenge Workshop*, Grenoble, France, 2023, pp. 63–68.
- [36] A. Mustafa, N. Pia, and G. Fuchs, "StyleMelGAN: An efficient high-fidelity adversarial vocoder with temporal adaptive normalization," in *ICASSP*, Toronto, Canada, 2021, pp. 6034–6038.
- [37] S. Patel, A. Rivas, and D. Psaltos, "Dataset of sounds of symptoms associated with respiratory sickness," Nov 2020. [Online]. Available: osf.io/tmkud
- [38] F. Busquet, F. Efthymiou, and C. Hildebrand, "Voice analytics in the wild: Validity and predictive accuracy of common audio-recording devices," *Behavior Research Methods*, vol. 56, no. 3, pp. 2114–2134, 2024.