



CQNV: A combination of coarsely quantized bitstream and neural vocoder for low rate speech coding

Youqiang Zheng¹, Li Xiao^{1,†}, Weiping Tu^{1,2,3,✉}, Yuhong Yang^{1,3}, Xinmeng Xu¹

¹NERCMS, School of Computer Science, Wuhan University, China

²Hubei LuoJia Laboratory, China

³Hubei Key Laboratory of Multimedia and Network Communication Engineering, Wuhan University, China

{youqiangzheng, xiaoli1996, tuweiping, yangyuhong, xuxinmeng}@whu.edu.cn

Abstract

Recently, speech codecs based on neural networks have proven to perform better than traditional methods. However, redundancy in traditional parameter quantization is visible within the codec architecture of combining the traditional codec with the neural vocoder. In this paper, we propose a novel framework named CQNV, which combines the coarsely quantized parameters of a traditional parametric codec to reduce the bitrate with a neural vocoder to improve the quality of the decoded speech. Furthermore, we introduce a parameters processing module into the neural vocoder to enhance the application of the bitstream of traditional speech coding parameters to the neural vocoder, further improving the reconstructed speech's quality. In the experiments, both subjective and objective evaluations demonstrate the effectiveness of the proposed CQNV framework. Specifically, our proposed method can achieve higher quality reconstructed speech at 1.1 kbps than Lyra and Encodec at 3 kbps.

Index Terms: speech coding, low bitrate, coarse quantization, neural vocoder

1. Introduction

Speech coding has emerged as a prevalent technique in various communication fields, such as mobile and satellite communication [1]. Its ability to significantly diminish the bandwidth and storage costs needed for information transmission and storage is well-known [2]. Existing traditional parametric codecs based on the linear prediction model have achieved significant success at bitrates above 2.4 kbps, such as MELP [3] and Codec2 [4]. These codecs rely on codebooks to quantize speech parameters and generate speech perceptually similar to the original [5]. However, as the bitrate is reduced, traditional parametric codecs face a performance challenge due to the inevitable increase in quantization error.

With advanced deep learning, the deep neural network has been universally applied to speech coding tasks, leading to promising results. In particular, the end-to-end architecture employs an encoder-decoder architecture, where both the encoder and decoder utilize neural networks and are jointly trained. This method has achieved a high quality while passing speech through a compact latent representation corresponding to very low bitrates [6–10]. Among them, Soundstream [7] is a whole end-to-end approach that compresses 24 kHz audio at the bitrates from 3 kbps to 18 kbps. Encodec [10], on the other hand, employs a streaming encoder-decoder architecture with quantized latent space and introduces a Transformer [11] language

model to further compress the obtained representation by up to 40%. It reduces the bandwidth of the 1.5 kbps model to 0.9 kbps and demonstrates superior performance compared to Opus [12] at 6 kbps. While end-to-end techniques have proven superior to conventional methods in low-bitrate speech coding, their black-box nature poses a significant challenge for model interpretation and improvement. Therefore, analyzing and enhancing end-to-end models is still a daunting task.

The neural vocoder-based approach, on the other hand, combines traditional and deep learning-based methods and provides more interpretability compared to the end-to-end architecture. Typically, the bitstream is generated by a traditional encoder and decoded by a neural vocoder. In [13], WaveNet [14] is the first explored as a generic generative model combined with Codec2 for speech coding, which achieves high performance at 2.4 kbps. Similarly, vocoders, including SampleRNN [15], and LPCNet [16] have been used in speech coding with a traditional encoder and present good performance [17, 18]. In order to boost speed in synthesizing speech, StyleMelgan [19] has been used to decode the bitstream of LPCNet at 1.6 kbps and generated state-of-the-art quality with low delay [20] by a high degree of parallelization.

The neural vocoder-based approaches mentioned above have demonstrated that in combination with parametric codecs, neural vocoders can achieve good quality in low-bitrate speech coding. However, we summarise two main drawbacks based on these methods: 1) **Redundancy in parameter quantization.** The bitrates associated with existing research [13] at 2.4 kbps and [16] at 1.6 kbps, exceed 1.2 kbps compared to the end-to-end approach, which implies the bitstream of traditional speech coding parameters exists redundancy that is not normally used by the neural vocoder. 2) **Possible performance limitations.** The neural vocoder employs a frame-based intermediate acoustic representation as its input, such as mel-spectrograms. However, the neural vocoder used in speech coding primarily relies on traditional speech coding parameters as input. This divergence in application contexts may impose limitations on the performance of the vocoders. Thus, a direct combination of bitstream from parametric codecs with a neural network vocoder may not be optimal.

The analysis above motivates us to design a new framework that decreases the bitrate while significantly enhancing the quality of the decoded speech. In this paper, we propose a framework that involves coarsely quantizing the speech parameters obtained from Codec2 to reduce the bitrate. We explore three ways of coarse quantization with different parameters. The coarsely quantized parameters are then fed into the HiFi-GAN [21] vocoder at a minimum of 1 kbps, which can solve the first drawback. This framework simplifies the codebook design of quantizers in traditional parametric codecs and achieves

✉Corresponding author: Weiping Tu(tuweiping@whu.edu.cn)

† Equal contribution. The numerical calculations in this paper have been done on the supercomputing system in the Supercomputing Center of Wuhan University.

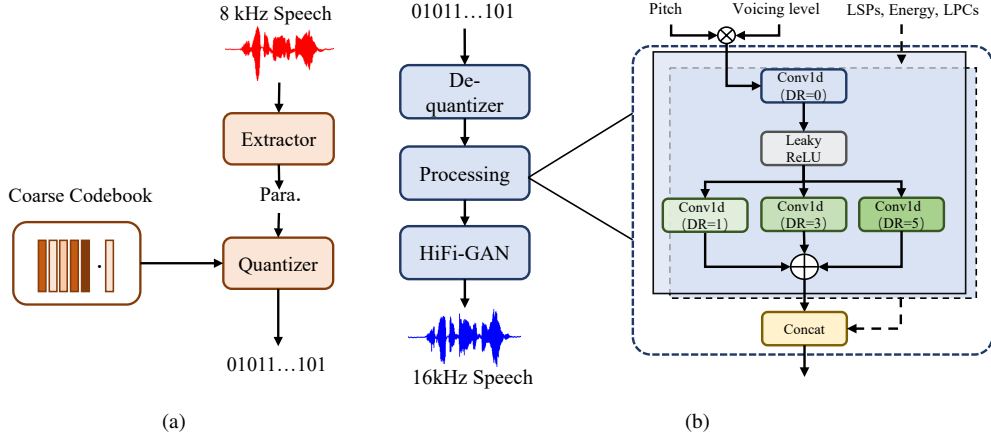


Figure 1: *The overall architecture of the proposed models. Parameters are coarsely quantized to a bitstream at (a) the encoder. At (b) the decoder bitstream is de-quantized to parameters, and it passes through the processing module as the conditional input of the vocoder.*

high-efficiency coding. Meanwhile, a parameters processing module is presented to tackle the second drawback by adopting three branches that carry different dilation rates of convolutional layers, which can achieve high-quality decoded speech reconstruction. This approach can significantly overcome the drawbacks mentioned above. The contributions of this study are summarized as follows:

- We explore the ability of the neural network-based vocoder to produce speech with coarsely quantized parameters of the traditional speech codecs, which provides a simplified approach to reduce the bitrate of speech while maintaining high quality.
- We present a parameters processing module that enhances neural vocoders' quality based on the bitstream of traditional speech coding parameters. Meanwhile, we introduce a dynamic hyper-parameter to balance the losses during training.
- Experiments show that our method is preferred to previous approaches even if they use more than 3x the bitrate.

2. Method

As shown in Figure 1, the proposed framework consists of Codec2 components (encoder and de-quantizer) and HiFi-GAN vocoder. The parameters quantizer is redesigned and trained on the encoder, which uses fewer bits to quantize the speech parameters extracted from the source speech sampled at 8 kHz. On the decoder, HiFi-GAN synthesizes 16 kHz speech with the parameters from the de-quantizer.

2.1. Encoder

Codec2 extracts parameters from each frame of the speech with a length of 10 ms. When working at 1.2 kbps, the parameters of every four consecutive frames are jointly quantized. Furthermore, 27, 16, and 4 bits are used to quantize the following parameters: line spectrum pairs (LSPs), pitch and energy, and voicing level. In order to reduce bitrate and explore the generative power of the neural vocoder with less input information, we propose a coarse quantization of LSPs, pitch and energy. For the voicing level, each frame is characterized by one-hot encoding, so it is not possible to cut the bitrate further.

The quantization of LSPs : LSPs are sampled every four

frames, and the three frames in between are interpolated at the de-quantizer. We use two-stage vector quantization [22] and split vector quantization for LSP, which keeps the same with Codec2 1.2 kbps. The 10-order LSPs are quantized in the first stage by using a codebook with the size of 512, which indicates that 9 bits are used to index the code vectors. And the quantization residual is computed. In the second stage, the quantization residual is split into two 5-dimension vectors, composed respectively of the odd-order and even-order components of the residual vector, and then quantized independently using two codebooks with the same size of 128 rather than 512. This operation indicates that index bits decrease from 9 to 7. All three codebooks are trained by using the LBG algorithm [23]. The quantization of LSPs achieves a reduction in the number of bits from 27 to 23 for four consecutive frames.

Joint quantization of Pitch and Energy : Pitch and energy are quantized by joint vector quantization. They are sampled every two frames, and the frame in between is interpolated in the decoder. The dimension of the codebook is 2. The first item is the \log_2 of the pitch compared to the lowest frequency, and the second item is the energy in dB, defined in Equation 1 and Equation 2, where \mathbf{W}_o is the pitch, and \mathbf{e} is the energy.

$$x_p = \log_2 \left(\frac{W_o}{\pi} \cdot \frac{4000}{50} \right) \quad (1)$$

$$x_e = 10 \times \log_{10} (e + 10^{-4}) \quad (2)$$

A codebook with a size of 256 is used in Codec2 at 1.2 kbps for quantizing pitch and energy. We designed and trained a new codebook with a size of 64, reducing the bits from 16 to 12 for a packet composed of 4 continuous frames. The codebook is designed using a method that combines prediction and vector quantization [24], which uses the prediction residuals and inter-frame correlation to improve the performance of the codebook, i.e., the prediction vector of the current frame is based on the prediction coefficients and the previous frame. The prediction coefficients are 0.8 for pitch and 0.9 for energy, following Codec2. Our codebook is trained on the residual vectors calculated by prediction and initial vectors. Moreover, weighted errors are computed for different features, e.g., pitch and energy errors are given much higher weights on stationary speech than on non-stationary speech or silence.




Figure 2: Comparing spectrograms of a test sample. Starting from (a) 16 kHz ground truth, (b) Lyra at 3 kbps (c) Encodec at 3 kbps, (d) CQNV-v3 at 1 kbps.

2.2. Decoder

In the decoder, a de-quantizer obtains parameters from the bit-stream. Then the parameters are introduced to condition the HiFi-GAN model at 100 Hz to produce speech signals. To condition HiFi-GAN, this work uses 23-dimension features, LSPs, energy, pitch, voicing level, and linear prediction coefficients (LPCs) calculated from LSPs. Instead of providing the features straight to the network, we employ a module as shown in Figure 1(b) to process the parameters. In this module, we adopt three branches to carry different dilation rates of convolutional layers to generate feature maps with different receptive field sizes. More precisely, the pitch is multiplied elementwise by voicing level to obtain a representation through this module. It is processed separately from the rest of the conditioning parameters because the pitch is critical for high-quality synthesis. Then LSPs, energy, and normalized LPCs parameters are passed through the same block to obtain another representation. Two representations are passed through concatenation to condition the HiFi-GAN. With the module, the neural vocoder based on the bitstream of traditional speech coding parameters can effectively reconstruct speech.

HiFi-GAN, a generative adversarial network consisting of one generator and two discriminators, achieves more efficient and high-fidelity speech than auto-regressive or flow-based models. It is trained to produce an audio waveform when conditioned on mel-spectrograms, along with GAN losses and two additional feature losses for improving training stability and model performance. Following the similar research [25], we introduce a dynamic ratio of the feature matching loss λ_{FM} to replace the fixed hyper-parameter used in [21]. It is the ratio of the mel-spectrograms loss to the feature matching loss during training. Our final objectives for the generator and discriminator are as follows:

$$\mathcal{L}_G = \mathcal{L}_{Adv}(G; D) + \lambda_{fm}\mathcal{L}_{FM}(G; D) + \lambda_{mel}\mathcal{L}_{Mel}(G) \quad (3)$$

$$\mathcal{L}_D = \mathcal{L}_{Adv}(D; G) \quad (4)$$

, where we set $\lambda_{mel} = 45$ and $\lambda_{fm} = \mathcal{L}_{Mel}(G)/\mathcal{L}_{FM}(G; D)$. This implementation effectively balances the losses and improves convergence speed during training.

3. Experimental Setup

3.1. Datasets

We use VCTK multi-speaker datasets [26] to train and evaluate the proposed coding method. The total length of the audio clips is approximately 40 hours, and the sample rate of the audio is

44.1 kHz. We downsample the speech data to 8 kHz for training codebook and 16 kHz for training HiFi-GAN. The codebook and HiFi-GAN are trained with the same corpus from 100 speakers, including 57 female and 43 male speakers, and tested with 8 unseen speakers, including 4 females and 4 males.

3.2. Implementation

Four codebooks are trained, three for quantizing LSPs and one for pitch and energy parameters. Based on the coding architecture of Codec2 at 1.2 kbps and the newly trained codebooks, three versions of the coding scheme with coarsely quantized parameters are developed. Table.1 shows the bit allocation of parameters in different versions for each packet with a length of 40 ms. **i) CQNV-v1**: bits of LSPs are reduced from 27 to 23, and the bitrate is reduced to 1.1 kbps; **ii) CQNV-v2**: bits of pitch and energy is reduced from 16 to 12, and the bitrate is reduced to 1.1 kbps; **iii) CQNV-v3**: bits of LSPs are reduced to 23, and bits of pitch and energy is reduced to 12, so 1 kbps.

Table 1: Bit allocation for a 40 ms packet. Comparison of our three different versions and Codec2 at 1.2 kbps.

	Codec2	CQNV-v1	CQNV-v2	CQNV-v3
LSPs	27 bits	23 bits	27 bits	23 bits
Pitch and Energy	16 bits	16 bits	12 bits	12 bits
Voicing level	4 bits	4 bits	4 bits	4 bits
Spare	1 bit	1 bit	1 bit	1 bit
Total	48 bits	44 bits	44 bits	40 bits

We train HiFi-GAN using one NVIDIA RTX 3090 GPU. Our batchsize is set to 32. For the generator, we set kernel sizes of transposed convolutions, the dilation, and kernel sizes of the residual stack to [5,4,4,2],[1,3,5],[3,7,11]. The FFT, window, and hop sizes are respectively set to 1024, 640, and 160. The optimizer, weight decay, and learning rate decay of the network are consistent with those in [21]. The adversarial training lasts for about 800k steps.

3.3. Evaluation

Evaluation Metrics. The evaluation of the proposed method involves both subjective and objective metrics. To measure the subjective quality of the reconstructed speech, we utilized the MUSHRA methodology [27]. A group of ten listeners, including four females and six males aged between 19 to 26, participated in the subjective listening tests. Twenty utterances, randomly selected from the test set, were evaluated. As for objective metrics, we employ the ViSQOL [28] to measure the objec-

tive quality of the proposed method. Objective evaluation and ablation studies are measured by “speech” ViSQOL, which operates on speech resampled at 16 kHz. All the utterance signals are resampled to 16 kHz in our experiments.

Baselines. We compare our proposed method with both traditional codecs and neural codecs, including Codec2 [24] at 3.2 kbps, Opus [12] at 6 kbps, Encodect [10] at 3 kbps, and Lyra [29] at 3 kbps which is trained on thousands of hours of speech data with speakers in over 70 languages¹. For Encodect, we select the 24 kHz pre-trained model to synthesize speech at 3 kbps without using Transformer language model². In addition, we use Speex [30] at 4 kbps as a low anchor.

4. Results

4.1. Subjective test

Figure 3 shows the MUSHRA score as a function of codec bitrate, where higher values and better perceptual quality. We observe that the three versions of the proposed method outperform all the baselines, including both traditional parametric and neural codecs, especially showing absolute advantages compared to traditional speech codecs. The proposed method obtains obviously higher MUSHRA scores than the traditional coders, with a sharp decrease in the bitrates. Notice CQNV at 1 kbps reaches better performance than Lyra at 3 kbps and Encodect at 3 kbps, proving our proposed method’s effectiveness.




Figure 3: *MUSHRA subjective test. The indicated interval in black represents the 95% confidence interval for each score.*

Figure 2 illustrates a test sample’s spectrograms. It can be seen from the area indicated by the dotted box that the spectrogram in the medium and low frequency band of speech synthesized by CQNV well matches the original signal. This indicates that CQNV reconstructs the harmonic structure of medium and low frequency bands better than that of Lyra and Encodect.

4.2. Objective test

For objective evaluation, we use all samples from the test set. We estimate the objective quality of the speech by computing ViSQOL. As shown in Table 2, we compare our methods with Lyra at 3 kbps and Encodect at 3 kbps. Table 2 shows that our proposed method *CQNV-v1* outperforms all other models even if they use more bandwidth. We further observe that *CQNV-v1* achieves higher ViSQOL than *CQNV-v2*, indicating that there is much redundancy in parameters LSPs than in parameters pitch and energy in Codec2. We also notice that the ViSQOL of

CQNV-v3 is comparable to that of *CQNV-v2* with saving 4 bits, demonstrating that our method effectively removes the redundancy in speech parameters.

Table 2: *Objective evaluation of CQNV based on ViSQOL.*

Method	Bitrate	ViSQOL
CQNV-v1	1.1 kbps	3.123
CQNV-v2	1.1 kbps	3.030
CQNV-v3	1 kbps	2.957
Lyra	3 kbps	2.865
Encodect	3 kbps	3.083

4.3. Ablation studies

We conduct several additional experiments to evaluate the effectiveness of the parameters processing module and the dynamic hyper-parameter. We compare our method with the one without the processing module and another without a dynamic hyper-parameter in the same setup. As shown in Table 3, we can see that the processing module and the hyper-parameter significantly improve the performance in speech reconstruction. We also notice a slight degradation between the non-dynamic hyper-parameter and dynamic hyper-parameter, which is more evidence of the effectiveness of this processing module.

Table 3: *Objective evaluation of CQNV-v3 at 1kbps. Ablation studies validate the effectiveness of the parameters processing module and the dynamic hyper-parameters.*

Method	ViSQOL
CQNV-v3	2.957
-v3(w/o. λ_{FM})	2.918
-v3(w/o.processing)	2.866
-v3(w/o. λ_{FM} and processing)	2.836

5. Conclusions

In this work, we present a novel framework that combines coarsely quantized bitstream with HiFi-GAN to lower the bitrate while preserving high-quality performance. Our proposed method involves coarse LSPs, pitch and energy quantization while maintaining compatibility with existing codecs. This approach removes extra information that is not usually used by neural vocoders, and simplifies the codebook design of quantizers in traditional parametric codecs. In addition, incorporating a processing module based on traditional speech coding parameters’ bitstream and introducing a dynamic hyper-parameter to balance the losses during training has been shown to enhance the quality of decoded speech. Our experimental results demonstrate the effectiveness of the proposed method in reducing the redundancy in speech parameters and improving the performance of the vocoder in speech reconstruction.

6. Acknowledgement

This work was supported in part by the Special Fund of Hubei LuoJia Laboratory (No. 220100019), the Hubei Province Technological Innovation Major Project(No. 2021BAA034) and the Fundamental Research Funds for the Central Universities (No. 2042023kf1033).

¹<https://github.com/google/lyra>

²<https://github.com/facebookresearch/encodect>

7. References

- [1] A. V. McCree, *Low-Bit-Rate Speech Coding*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 331–350. [Online]. Available: https://doi.org/10.1007/978-3-540-49127-9_16
- [2] P. Noll, “Mpeg digital audio coding,” *IEEE signal processing magazine*, vol. 14, no. 5, pp. 59–81, 1997.
- [3] L. M. Supplee, R. P. Cohn, J. S. Collura, and A. V. McCree, “Melp: the new federal standard at 2400 bps,” in *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 2. IEEE, 1997, pp. 1591–1594.
- [4] D. Rowe, “Codec 2-open source speech coding at 2400 bits/s and below,” in *TAPR and ARRL 30th Digital Communications Conference*, 2011, pp. 80–84.
- [5] D. O’Shaughnessy, “Review of methods for coding of speech signals,” *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2023, no. 1, pp. 1–25, 2023.
- [6] C. Gărbacea, A. van den Oord, Y. Li, F. S. Lim, A. Luebs, O. Vinyals, and T. C. Walters, “Low bit-rate speech coding with vq-vae and a wavenet decoder,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 735–739.
- [7] 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, 2021.
- [8] N. Pia, K. Gupta, S. Korse, M. Multrus, and G. Fuchs, “Nesc: Robust neural end-2-end speech coding with gans,” *arXiv preprint arXiv:2207.03282*, 2022.
- [9] X. Jiang, X. Peng, C. Zheng, H. Xue, Y. Zhang, and Y. Lu, “End-to-end neural speech coding for real-time communications,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 866–870.
- [10] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, “High fidelity neural audio compression,” *arXiv preprint arXiv:2210.13438*, 2022.
- [11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [12] J.-M. Valin, K. Vos, and T. Terriberry, “Definition of the opus audio codec,” Tech. Rep., 2012.
- [13] W. B. Kleijn, F. S. Lim, *et al.*, “Wavenet based low rate speech coding,” in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 676–680.
- [14] A. v. d. Oord, S. Dieleman *et al.*, “Wavenet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*, 2016.
- [15] S. Mehri, K. Kumar, *et al.*, “SAMPLERNN: An unconditional end-to-end neural audio generation model,” *arXiv preprint arXiv:1612.07837*, 2016.
- [16] J.-M. 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)*. IEEE, 2019, pp. 5891–5895.
- [17] J. Klejsa, P. Hedelin *et al.*, “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.
- [18] J.-M. Valin and J. Skoglund, “A real-time wideband neural vocoder at 1.6 kb/s using lpcnet,” *arXiv preprint arXiv:1903.12087*, 2019.
- [19] A. Mustafa, N. Pia, and G. Fuchs, “Stylemelgan: An efficient high-fidelity adversarial vocoder with temporal adaptive normalization,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6034–6038.
- [20] A. Mustafa, J. Bütthe *et al.*, “A streamwise gan vocoder for wide-band speech coding at very low bit rate,” in *2021 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2021, pp. 66–70.
- [21] J. Kong, J. Kim, and J. Bae, “Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 17 022–17 033, 2020.
- [22] R. Gray, “Vector quantization,” *IEEE Assp Magazine*, vol. 1, no. 2, pp. 4–29, 1984.
- [23] Y. Linde, A. Buzo, and R. Gray, “An algorithm for vector quantizer design,” *IEEE Transactions on communications*, vol. 28, no. 1, pp. 84–95, 1980.
- [24] J.-M. Valin, “A pitch-energy quantizer for codec2,” <https://jmvain.dreamwidth.org/tag/codec2>, 2012.
- [25] J. Yang, J.-S. Bae, T. Bak, Y. Kim, and H.-Y. Cho, “Ganspeech: Adversarial training for high-fidelity multi-speaker speech synthesis,” *arXiv preprint arXiv:2106.15153*, 2021.
- [26] J. Yamagishi, C. Veaux, K. MacDonald *et al.*, “Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit (version 0.92),” 2019.
- [27] R. ITU-R, “1534-1, method for the subjective assessment of intermediate quality levels of coding systems (mushra),” *International Telecommunication Union*, 2003.
- [28] 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.
- [29] W. B. Kleijn, A. Storus, M. Chinen, T. Denton, F. S. Lim, A. Luebs, J. Skoglund, and H. Yeh, “Generative speech coding with predictive variance regularization,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6478–6482.
- [30] J.-M. Valin, “The speex codec manual version 1.2 beta 3,” *Xiph.org Foundation*, 2007.