

LOW BIT-RATE SPEECH CODING WITH VQ-VAE AND A WAVENET DECODER

Cristina Gârbasea*¹, Aäron van den Oord², Yazhe Li²,
Felicia S C Lim³, Alejandro Luebs³, Oriol Vinyals², Thomas C Walters²

¹University of Michigan, Ann Arbor, USA ²DeepMind, London, UK ³Google, San Francisco, USA

ABSTRACT

In order to efficiently transmit and store speech signals, speech codecs create a minimally redundant representation of the input signal which is then decoded at the receiver with the best possible perceptual quality. In this work we demonstrate that a neural network architecture based on VQ-VAE with a WaveNet decoder can be used to perform very low bit-rate speech coding with high reconstruction quality. A prosody-transparent and speaker-independent model trained on the LibriSpeech corpus coding audio at 1.6 kbps exhibits perceptual quality which is around halfway between the MELP codec at 2.4 kbps and AMR-WB codec at 23.05 kbps. In addition, when training on high-quality recorded speech with the test speaker included in the training set, a model coding speech at 1.6 kbps produces output of similar perceptual quality to that generated by AMR-WB at 23.05 kbps.

Index Terms— Speech coding, low bit-rate, generative models, WaveNet, VQ-VAE

1. INTRODUCTION AND RELATED WORK

Speech codecs typically employ a carefully hand-engineered pipeline made up of an encoder and a decoder which take into account the physics of speech production to remove redundancies in the data and yield a compact bitstream. High quality speech codecs typically operate at bit-rates over 16 kbps. Recent advances in the field of deep learning, specifically autoencoder networks which directly learn the mapping of inputs to outputs by means of an encoder-decoder framework with an information bottleneck in between, open up the possibility of learning both the encoder and the decoder directly from speech data. These models are able to learn the redundancies in signals directly by being exposed to many examples during training, and have been successfully applied in the domain of image compression [1–4]. In the speech domain, end-to-end coding operating at rates as low as 369 bps by use of extracted acoustic features [5] and an end-to-end optimized wideband codec performing on par with the adaptive multirate wideband (AMR-WB) codec at 9-24 kbps [6] have both recently been proposed. Creating a minimally re-

dundant representation of speech is inherently related to determining the true information rate of the input signal. Kuyk *et al.* [7] compute the true information rate of speech to be less than 100 bps, yet current systems typically require a rate roughly two orders of magnitude higher than this to produce good quality speech, suggesting that there is significant room for improvement in speech coding.

The WaveNet [8] text-to-speech model shows the power of learning from raw data to generate speech. Kleijn *et al.* [9] use a learned WaveNet decoder to produce audio comparable in quality to that produced by the AMR-WB [10] codec at 23.05 kbps from the bit stream generated by the encoder in Codec2 [11] (a parametric codec designed for low bit-rates) operating at 2.4 kbps; this demonstrates the effectiveness of a learned decoder over a hand-engineered one. Furthermore, van den Oord *et al.* [12] demonstrate a learned autoencoder – the vector-quantized variational autoencoder (VQ-VAE) – which is able to encode speech into a compact discrete latent representation and then reconstruct the original audio with high quality by sampling from a WaveNet-like decoder.

In this paper, we evaluate the VQ-VAE architecture introduced in [12] as an end-to-end learned speech codec, showing that it can yield high reconstruction quality while passing speech through a compact latent representation corresponding to very low bit-rates for coding. We propose and evaluate various modifications to the VQ-VAE / WaveNet architecture as described in [12] in order to make it more suited to the task of speech coding. We create a speaker-independent model which can accurately reproduce both the content and the prosody of the input utterance while passing through a compact latent representation that leads to lower bit rates and higher reconstruction quality than the current state-of-the-art.

The VQ-VAE combines a variational autoencoder (VAE) [13] with a vector quantization (VQ) layer to produce a discrete latent representation which has been shown to capture important high-level features in image, audio and video data, yielding an extremely compact and semantically meaningful representation of the input. The prior and posterior distributions are categorical, and samples drawn from these distributions index an embedding which is passed as input to the decoder network. Van den Oord *et al.* [12] demonstrate the use of VQ-VAE in the audio domain, using a convolutional

*Work performed while at DeepMind

encoder and a WaveNet decoder on speech. The authors use the VQ-VAE model to generate very high-quality speech even when using a latent representation that is 64 times smaller than the original input waveform, reducing 16-bit pulse-code modulation (PCM) encoded speech at 16 kHz sample rate (256 kbps) to a stream of discrete latent codes at 4 kbps. In this case, the network learns to represent the high-level semantic content of the speech, therefore when a speech waveform is passed through the autoencoder network, the reconstruction contains the same syllabic content, but the prosody of the utterance may be significantly altered. Further analysis of the representation learned by the network shows that the discrete latent codes are well-correlated with the stream of syllables in the input speech, suggesting that the encoder network builds a high-level representation of the content of the utterance, and then the decoder generates plausible-sounding specific details. Given the very high quality of the reconstructed audio, and the compact latent representation that the network can produce, the VQ-VAE architecture seems to lend itself naturally to the task of low bit-rate audio coding if it can be generalized to preserve the identity of arbitrary speakers.

2. VQ-VAE BASED SPEECH CODEC

Given the ability of the VQ-VAE model to learn a high-level abstract space invariant to low-level perturbations in the input and which encodes only the contents of speech [12], our main goal is to determine if the model can be used as an end-to-end speech codec (i.e. map the input audio to the discrete latent space learnt by VQ-VAE and then reconstruct the signal by sampling from the decoder of the network conditioned on the latent representation). In order to be useful as a generic speech codec, the model must be constrained to maintain both the speaker identity and the prosody of the utterance at decoding time. We present these architectural changes below.

2.1. Maintaining speaker identity

To encourage the model to be speaker-agnostic and generalize across speakers it has never seen in the training set, we modify the original VQ-VAE architecture by removing explicit conditioning on speaker identity, which was done through a one-hot code passed to both the encoder and the decoder both at training and synthesis time. In its place, we add a latent representation (with an associated codebook) that takes its input from the whole utterance and does not vary over time. The time invariant code is generated by mean pooling over the time dimension of the encoder output and fed to a separate codebook. The expectation is that the network will learn to use the time-varying set of codes to encode the message content which varies over time, while summarising and passing speaker-related information through the separate non time-varying set of codes. In our experiments, the non time-varying code is computed over the entire utterance at en-

coding time. In the experiments below, the codebooks were of similar size to those used by the time-varying latents (eg. 256 elements), which led to a small additional coding overhead of 1-2 bytes over the course of the utterance. Such a long window would be unrealistic for online speech coding; in this case the static code could be replaced with a code that varies slowly over time and which is based only on past information.

2.2. Constraining prosody

In order to constrain the model to pass prosodic as well as semantic information through the bottleneck representation, it is necessary to design a training regime in which the network is encouraged, via a loss term, to pass pitch (f_0) and timing information through to the decoder. To this end, a second decoder was added to the network in parallel with the WaveNet audio decoder; the task of this network is to predict the f_0 track of the training utterance, using a common latent representation with the audio decoder. An additional f_0 prediction term with a tuneable weight is added to the loss function, which causes the latent representation to pass pitch track information through the bottleneck. This information is then available to the waveform decoder network, which uses it to produce an utterance that has the same pitch track as the original. The f_0 predictor model uses the same architecture as the audio decoder, but only has to upsample the latent representation to the rate of the f_0 feature. When the trained model is used as a codec, the f_0 prediction network can be removed from the model, and the audio decoder alone used to reconstruct the waveform.

In our datasets f_0 information had already been extracted from the audio waveform using a standard pitch tracker, and was sampled at 200 Hz. Given the 16 kHz sample rate of the audio in the training set, we introduce an extra convolutional layer with stride 5 on the encoder side for rate compatibility between the audio samples and the f_0 samples.

3. EXPERIMENTS

In the set of experiments below our goal is to determine under which hyper-parameter settings VQ-VAE can be used as a very low rate speech codec. Therefore, we focus on establishing how the quality of the reconstructed signal varies as a function of the latent representation size of the model. Changing the dimensionality of the latent representation is equivalent to varying the bit-rate of the encoded signal. In the original VQ-VAE setup the latents consist of one feature map and the latent space is 512-dimensional. In our experiments we vary in turn the number of encoder layers, the number of latent maps, as well as the number of discrete codes per map.

Datasets We carry our experiments on two different datasets. Most experiments rely on the LibriSpeech [14] corpus (LS), which contains around 1,000 hours of read En-

glish speech from a total of 2,302 speakers at a sample rate of 16 kHz. For the majority of results presented in this paper, the `train-clean-100`, `train-clean-360` and `train-other-500` subsets of the corpus were combined and used at training time. Evaluation samples were created using the `test-clean` set. The set of speakers in the test data is disjoint with those in the training data.

For cases when we evaluate models including test speakers at training time, we held out 1/100 of the `test-clean` dataset for evaluation, and add the remaining data (totalling ~ 5 hours of extra audio) to the training set. We refer to this dataset as the augmented LibriSpeech (LSplus) corpus. The datasets were annotated with automatically extracted f_0 information (presented as $\log(f_0)$), sampled at 200 Hz.

In addition to the datasets presented above, we also had access to a proprietary corpus consisting of 10 American English speakers totalling approximately 300 hours of speech, of which at least 85% was recorded in studio conditions. This dataset, referred to as the Studio corpus, is used to investigate performance in ideal training conditions.

Evaluation To assess the reconstructed speech samples in a way that reflects informal listening impressions, we follow Kleijn *et al.* [9] and use subjective MUSHRA-type listening tests [15] to measure the perceived quality of the output from lossy audio compression algorithms. In MUSHRA evaluations the listener is presented with a labelled uncompressed signal for reference, a set of numbered samples including the test samples, a copy of the uncompressed reference, and a low-quality anchor. We compress the anchor samples with Speex [16] at 2.4kbps (below its comfortable range of use).

For some evaluations we also include other speech codecs for comparison: Codec2 [11] at 2.4kbps, MELP [17] at 2.4kbps, AMR-WB [10] at 23.05kbps and the original signal after companding and quantization using 8-bit μ -law [18] (giving a bit-rate of 128kbps). μ -law is the companding function used by the WaveNet decoder to map the 16-bit sample resolution to an 8-bit softmax distribution, and so it gives an upper bound on the quality of the reconstruction from the VQ-VAE model. Each evaluation includes 8 utterances, and each utterance is evaluated by 100 human raters. The MUSHRA scale from 0-100 allows rating very small differences between the samples.

3.1. Quality Evaluation

Initial experiments with the Studio dataset led to a choice of hyperparameters that gave very good reconstruction quality for voices in that dataset. The number of encoder layers (each layer downsampling by a factor equal to the stride), the number of latent maps (i.e. the number of separate codebooks) and the size of each codebook used were varied. We ascertain that high-quality reconstruction of audio is possible with five strided convolutional encoder layers of stride 2 and one of stride 5 (leading to a total downsampling factor of

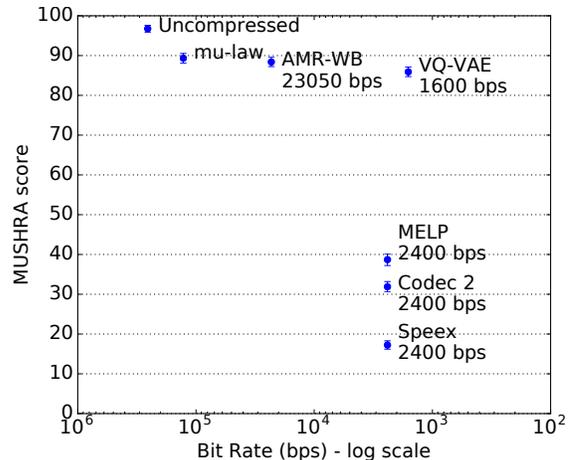


Fig. 1. MUSHRA score vs bit-rate for the VQ-VAE speech codec at 1.6 kbps, trained on Studio data and evaluated on a single studio-recorded voice present in the train set, against a variety of other codecs. Bit rate reduces from left to right, with the optimum performance for a codec suggested in [7] being in the top right corner of the graph.

$2^5 * 5 = 160$ - a feature rate of 100 Hz), two latent maps each with a 256-element codebook (each map is representable by an 8-bit value), coding a 64-dimensional latent vector. This leads to a bit rate of 1600 bps. This model is evaluated in a MUSHRA test against a range of other codecs, and the results are presented in Figure 1.

In the plot we notice the emergence of two clusters of high and low reconstruction quality: *i) the low quality cluster* includes Speex, Codec2 and MELP operating at 2.4kbps, and *ii) the high quality cluster* consists of AMR-WB at 23.05kbps, 8 bit μ -law and VQ-VAE at 1.6kbps. In these conditions, using speech from a speaker contained in the training set and trained on studio-quality utterances, the VQ-VAE model is very effective at coding the audio, achieving a factor of 14 improvement in compression rate over a standard wideband audio codec for a minimal reduction in perceptual quality.

While this is an interesting result, it does not make for a fully robust codec system, since the test speaker was present in the training set, and the train and test audio were both of very high quality. In order to evaluate the performance of the model in more natural conditions, we then switched to using the LibriSpeech corpus instead of the Studio corpus, which also has the benefit of containing many more speakers and a variety of different noise conditions.

In Figure 2 we present results for a set of models trained on the full LibriSpeech dataset. The 1600 bps model uses the same architecture parameters as the model trained on Studio data above, while the 800 bps and 400 bps models each add one further encoder layer with a stride of 2, which halves the bit rate of the model. Here we see that the model at 1600 bps is lower in perceptual quality than previously, but still

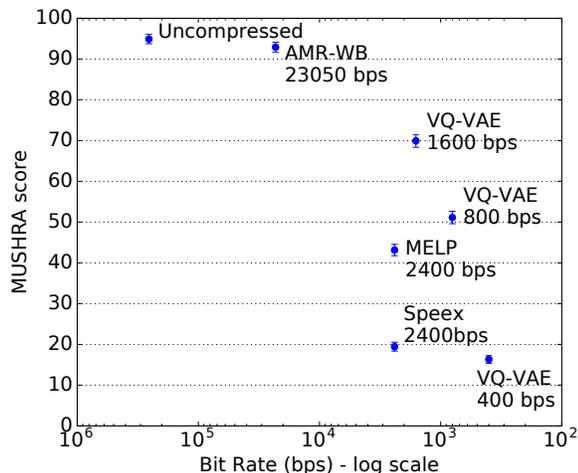


Fig. 2. MUSHRA score vs bit-rate for the VQ-VAE speech codec, trained on LibriSpeech data and evaluated on voices from the LibriSpeech test set, against various other codecs.

roughly halfway between MELP at 2400 bps and AMR-WB at 23.05 kbps. Even at 800 bps, the VQ-VAE model outperforms MELP at 3 times the bitrate, and it is only at 400 bps that performance degrades below that of Speex at 2400 bps.

Finally, we investigate the effect of including utterances from the test speaker in the train set, to determine how much benefit the model can gain from training on a specific speaker. These results are presented in Figure 3, in which we observe a small advantage in the perceptual quality for the VQ-VAE model when training on utterances from test speakers.

3.2. Speaker Transparency

We follow Jia *et al.* [19] and Chen *et al.* [20] who assess speaker similarity using mean opinion score (MOS) evaluations based on subjective listening tests. MOS evaluations [21] map ratings from bad to excellent into the range 1-5 with 0.5 point increments. To measure how similar synthesized speech is to real speech of the target speaker, each synthesized utterance is paired with a randomly selected ground truth utterance from the same speaker. Raters are explicitly asked throughout the study to not judge the content, grammar or the audio quality of the sentences, but instead to focus on the similarity of speakers to one another. Eight pairs of utterances from the LibriSpeech test set were rated by 20 listeners each.

MOS scores for speaker similarity are presented in Table 1 for versions of the VQ-VAE codec at 1.6 kbps trained on the LibriSpeech dataset and the augmented LibriSpeech dataset (LSplus, meaning that the test speakers were in the train set), and for comparison MELP at 2.4 kbps and Speex at 2.4 kbps.

We see that the speaker similarity MOS for the VQ-VAE based codec is higher than for the other two low-bitrate codecs, and is slightly higher for the case where the test speakers are included in the train set. The wider error bounds

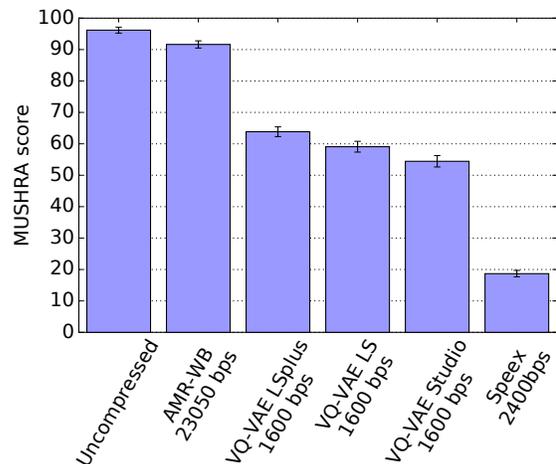


Fig. 3. MUSHRA results for the VQ-VAE speech codec at 1.6 kbps, trained on Studio data, on LibriSpeech and on LibriSpeech plus test set speakers and evaluated on voices from the LibriSpeech test set.

| Codec | Speaker Similarity MOS |
|------------------------|------------------------|
| VQ-VAE LSplus 1600 bps | 3.794 ± 0.451 |
| VQ-VAE LS 1600 bps | 3.703 ± 0.716 |
| MELP 2400 bps | 3.138 ± 0.324 |
| Speex 2400 bps | 2.534 ± 0.233 |

Table 1. MOS results for the VQ-VAE speech codec at 1600 bps trained on LibriSpeech with and without the test voices in the train set, Speex at 2400 bps and MELP at 2400 bps.

on the result for the VQ-VAE codec trained on LibriSpeech alone are at least in part due to a very low MOS, of 1.85, being assigned to one of the speakers in the evaluation (the next lowest scoring utterance has a speaker similarity MOS of 3.38); the MOS for this speaker is significantly higher, at 2.73, when the test speaker is included in the train set.

4. CONCLUSION

In this work we demonstrate initial results showing the suitability of the VQ-VAE model as a robust very low bit-rate speech codec. After adjusting the architecture to make it speaker and prosody transparent, we show that VQ-VAE at 1600 bps outperforms some popular low bit-rate speech codecs operating at 2400 bps. Listening test results demonstrate that the VQ-VAE based codec is somewhat better at preserving speaker identity than other low-rate codecs, but more work is required to understand the variability in preservation of speaker identity. Additionally we expect that model performance can be improved beyond the current 1600 bps.

Acknowledgements Thanks to Jan Chorowski, Norman Casagrande, Yutian Chen, Ray Smith, Florian Stimberg, Ron Weiss, Heiga Zen and the ICASSP reviewers.

5. REFERENCES

- [1] Lucas Theis, Wenzhe Shi, Andrew Cunningham, and Ferenc Huszár, “Lossy image compression with compressive autoencoders,” *arXiv preprint arXiv:1703.00395*, 2017.
- [2] Johannes Ballé, Valero Laparra, and Eero P. Simoncelli, “End-to-end optimized image compression,” *CoRR*, vol. abs/1611.01704, 2016.
- [3] Eirikur Agustsson, Fabian Mentzer, Michael Tschanen, Lukas Cavigelli, Radu Timofte, Luca Benini, and Luc Van Gool, “Soft-to-hard vector quantization for end-to-end learned compression of images and neural networks,” *CoRR*, vol. abs/1704.00648, 2017.
- [4] Oren Rippel and Lubomir Bourdev, “Real-time adaptive image compression,” *arXiv preprint arXiv:1705.05823*, 2017.
- [5] Milos Cernak, Alexandros Lazaridis, Afsaneh Asaei, Philip N Garner, Milos Cernak, Alexandros Lazaridis, Afsaneh Asaei, and Philip N Garner, “Composition of deep and spiking neural networks for very low bit rate speech coding,” *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, vol. 24, no. 12, pp. 2301–2312, 2016.
- [6] Srihari Kankanahalli, “End-to-end optimized speech coding with deep neural networks,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 2521–2525.
- [7] Steven Van Kuyk, W Bastiaan Kleijn, and Richard C Hendriks, “On the information rate of speech communication,” in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 5625–5629.
- [8] Aäron Van Den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W Senior, and Koray Kavukcuoglu, “Wavenet: A generative model for raw audio,” in *SSW*, 2016, p. 125.
- [9] W Bastiaan Kleijn, Felicia SC Lim, Alejandro Luebs, Jan Skoglund, Florian Stimberg, Quan Wang, and Thomas 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.
- [10] Bruno Bessette, Redwan Salami, Roch Lefebvre, Milan Jelinek, Jani Rotola-Pukkila, Janne Vainio, Hannu Mikkola, and Kari Jarvinen, “The adaptive multirate wideband speech codec (AMR-WB),” *IEEE transactions on speech and audio processing*, vol. 10, no. 8, pp. 620–636, 2002.
- [11] D Rowe, “Codec 2-open source speech coding at 2400 bit/s and below,” Tech. Rep., 2011.
- [12] Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu, “Neural discrete representation learning,” in *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., pp. 6306–6315. Curran Associates, Inc., 2017.
- [13] Diederik P Kingma and Max Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [14] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 5206–5210.
- [15] ITU-R Recommendation BS.1534-1, “Method for the subjective assessment of intermediate sound quality (MUSHRA),” 2001.
- [16] Jean-Marc Valin, “Speex: A free codec for free speech,” *arXiv preprint arXiv:1602.08668*, 2016.
- [17] Alan McCree, Kwan Truong, E Bryan George, Thomas P Barnwell, and Vishu Viswanathan, “A 2.4 kbit/s MELP coder candidate for the new US Federal standard,” in *Acoustics, Speech, and Signal Processing, 1996. ICASSP-96. Conference Proceedings., 1996 IEEE International Conference on*. IEEE, 1996, vol. 1, pp. 200–203.
- [18] ITU-T Recommendation G.711, “Pulse code modulation (PCM) of voice frequencies,” 1988.
- [19] Ye Jia, Yu Zhang, Ron J Weiss, Quan Wang, Jonathan Shen, Fei Ren, Zhifeng Chen, Patrick Nguyen, Ruoming Pang, Ignacio Lopez Moreno, et al., “Transfer learning from speaker verification to multispeaker text-to-speech synthesis,” *arXiv preprint arXiv:1806.04558*, 2018.
- [20] Yutian Chen, Yannis Assael, Brendan Shillingford, David Budden, Scott Reed, Heiga Zen, Quan Wang, Luis C Cobo, Andrew Trask, Ben Laurie, et al., “Sample efficient adaptive text-to-speech,” *arXiv preprint arXiv:1809.10460*, 2018.
- [21] ITU-T Recommendation P.800, “Methods for subjective determination of transmission quality,” pp. 800–899, 1996.