

VQPC-GAN: VARIABLE-LENGTH ADVERSARIAL AUDIO SYNTHESIS USING VECTOR-QUANTIZED CONTRASTIVE PREDICTIVE CODING

Javier Nistal^{*,1,2} Cyran Aouameur,¹ Stefan Lattner,¹ Gaël Richard²

¹ Sony Computer Science Laboratory, 6 Rue Amyot, 75005 Paris, France

² LTCI, Télécom Paris, Institut Polytechnique de Paris, France

ABSTRACT

Influenced by the field of Computer Vision, Generative Adversarial Networks (GANs) are often adopted for the audio domain using fixed-size two-dimensional spectrogram representations as the “image data”. However, in the (musical) audio domain, it is often desired to generate output of variable duration. This paper presents VQPC-GAN, an adversarial framework for synthesizing variable-length audio by exploiting Vector-Quantized Contrastive Predictive Coding (VQPC). A sequence of VQPC tokens extracted from real audio data serves as conditional input to a GAN architecture, providing step-wise time-dependent features of the generated content. The input noise z (characteristic in adversarial architectures) remains fixed over time, ensuring temporal consistency of global features. We evaluate the proposed model by comparing a diverse set of metrics against various strong baselines. Results show that, even though the baselines score best, VQPC-GAN achieves comparable performance even when generating variable-length audio. Numerous sound examples are provided in the accompanying website,¹ and we release the code for reproducibility.²

Index Terms— Generative Adversarial Networks, Audio Synthesis, Vector-Quantized Contrastive Predictive Coding

1. INTRODUCTION

In recent years, Generative Adversarial Networks (GANs) [1] have shown outstanding results in image and audio synthesis tasks [2, 3, 4, 5, 6]. As most (initial) studies on GANs focused on generating images, the resulting architectures are now often adopted for the musical audio domain, using fixed-size two-dimensional spectrogram representations as the “image data”. However, while it is a natural choice to use data of fixed dimensionality in the visual domain, fixing the length of musical audio content in generation tasks poses a significant limitation. As a result, GANs are currently mainly used to generate short audio content in the musical audio domain, like single notes of a tonal instrument or single percussion samples [4, 5].

When dealing with variable-length sequence generation, commonly utilized models are Transformer architectures [7], causal convolutional architectures (causal CNNs) [8], and recurrent neural networks (RNNs) [9]. However, those models suffer various problems like high computational cost (autoregressive), missing look-back capabilities (recurrent), and, typically, they cannot be parallelized at test time. In contrast, GANs are relatively efficient in generating high-dimensional data, as the conditioning on a single

noise vector determines the values of all output dimensions at once. Therefore, it seems reasonable to also adopt the GAN paradigm for generating variable-length musical audio content. It has been shown in text-to-speech translation [6] that GANs can be successful in generating coherent variable-length audio when conditioned on meaningful sequences of symbols (i.e., linguistic + pitch features), while the input noise z accounts for the remaining variability.

We adopt a similar strategy by first learning sequences of symbolic audio descriptors, serving as conditional inputs to a GAN architecture. These descriptors are discrete tokens learned through self-supervised training, using Vector-Quantized Contrastive Predictive Coding (VQPC) [7]. In VQPC, discrete representations are learned through contrastive learning, by confronting positive and negative examples. In contrast to reconstruction-based VQVAEs [10], VQPC allows to control to some extent which aspects of the (sequential) data are captured in the tokens, by carefully designing a negative sampling strategy, thus defining the so-called “pretext” task. In this work, the tokens are trained to represent the *temporal evolution* (i.e., the envelope) of single, pitched audios of different instruments. The proposed model is conditioned on such envelope feature sequences, on the noise vector z (static, representing the “instrument”), and on pitch information (static). This approach of sequence generation with GANs using discrete tokens is promising for future, more elaborate applications. While in this work, we are simply up-sampling token sequences to generate longer sounds, one could also *generate* plausible token sequences. Such a system could then be used to hold sounds for an arbitrary time in real-time performance with a MIDI input device. Also, token sequences could be generated conditioned on MIDI information, to represent the dynamics of a target instrument. The resulting system could then be used for naturalistic rendering of MIDI files. Furthermore, training tokens to also represent pitch information would result in a more general variable-length audio generation framework. To the best of our knowledge, this is the first work implementing a variable-length GAN for *musical* audio.

The paper is organized as follows. First, in Sec. 2, we summarize previous works on adversarial audio synthesis, time-series GANs, and contrastive predictive coding. In Sec. 3 we describe in detail the proposed framework. Sec. 4 describes the experiment setup. Next, in Sec. 5, we evaluate the proposed method and compare results with previous work and other baselines. Finally, in Sec. 6, we draw some conclusions and discuss future directions.

2. PREVIOUS WORK

In the following, we review some of the most important works on adversarial audio synthesis, variable-length time-series generation using GANs, and contrastive learning of sequences. We pay special attention to those works focused on audio data.

*Nistal received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 765068.

¹sonycslparis.github.io/vqpc-gan.io

²github.com/SonyCSLParis/vqpc-gan

- **Adversarial audio synthesis.** Applications of GANs to audio synthesis have mainly focused on speech tasks [6, 11, 12]. The first GAN to synthesize musical audio was WaveGAN [13]. Recently, GANSynth [4] surpassed WaveNet [8] baselines on the task of audio synthesis of musical notes using pitch as conditional information. Follow-up works applied similar architectures to drum sound synthesis [5, 14]. DrumGAN [5] performs high-quality synthesis of a variety of drum sounds conditioned on continuous timbral features. Other works employed GANs for Mel-spectrogram inversion [12], audio domain adaptation [15] or audio enhancement [16]. In this work, we adapt the Wasserstein GAN (WGAN) described in previous works [4, 5] to a sequence generation scheme by conducting two major architectural changes that we present in Sec. 3.
- **Time-Series GAN.** Early approaches to adversarial time-series generation used RNNs for both the generator and discriminator’s architecture [17, 18]. C-RNN-GAN [18] generates musical data recurrently, taking as input a noise vector and the previous step’s generated data. While this approach relied only on the binary adversarial feedback for learning, TimeGAN [19] incorporated an autoregressive supervised loss to better capture step-wise temporal dynamics of the training data. To maintain global and local consistency, the generator is conditioned on static and sequential random vectors. Similarly, GAN-TTS [6] synthesizes variable-length speech by conditioning the generator on sequential linguistic and pitch features, as well as a global random vector and a speaker ID. We take inspiration from these works and condition the generator on *static* and *dynamic* prior information. The *static* information is represented by a random vector and a pitch class, whereas the *dynamic* information is captured by a sequence of discrete tokens learned through self-supervised training.
- **Contrastive Predictive Coding (CPC)** [20] is a self-supervised framework used to learn general features from an unlabeled dataset of sequences by contrasting *positive* and *negative* examples in a so-called pretext task. CPC has been actively studied for speech tasks [20, 21, 22], where it was shown to improve the performance of ASR systems when used as front-end in replacement of spectrograms [21]. Introducing a VQ bottleneck to the CPC improved the system’s performance by discarding irrelevant information [22, 23]. In contrast to previous works exploiting VQCPC for discriminative downstream tasks [20, 21, 22], recent approaches explore small codebook sizes to learn compact, discrete representations of symbolic music from which to generate variations of any music piece [7]. We follow a similar strategy in this work and use VQCPC to condition a GAN on such discrete codes for synthesizing variable-length audio.

3. VQCPC-GAN

In this section, we briefly describe Contrastive Predictive Coding (CPC) and a variant for discrete representations, the Vector-Quantized CPC (VQCPC) [7], as well as the two building blocks of VQCPC-GAN, the VQCPC encoder and the GAN.

3.1. Background

Contrastive Predictive Coding (CPC) is a self-supervised representation learning technique for extracting compact, low-dimensional sequences of latent codes from high-dimensional signals [20].

Given an input sequence $\mathbf{x} = [x_1, \dots, x_L]$ with length L , an encoder f_{enc} maps each element x_i into a real-valued embedding vector $z_i = f_{\text{enc}}(x_i) \in R^{d_z}$. Next, an autoregressive model f_{ar} summarizes past and present context of the embeddings $z_{\leq t}$ into a single context vector $h_t = f_{\text{ar}}(z_{\leq t}) \in R^{d_h}$.

The encoder and autoregressive model are trained to minimize the Information Noise Contrastive Estimation (InfoNCE) loss. Minimizing the InfoNCE loss is equivalent to maximizing the mutual information between the context vector h_t and the encoding of future subsequences $z_{t+k} = f_{\text{enc}}(x_{t+k}), \forall k \in [1, K]$, where K is the number of future predictions [20]. Formally, given an entry of the dataset \mathbf{x} , the model has to identify the encoding obtained from the true x_{t+k} , so-called *positive example*, from those obtained from a set of so-called *negative examples*, drawn from a proposal distribution $p(x_{t+k})$. Defining \mathcal{S} as the set containing $N - 1$ negative examples as well as the single positive example x_{t+k} , the InfoNCE loss is written as

$$\mathcal{L}_{\text{NCE}}(x_t) = - \sum_{k=1}^K E_{\mathcal{S}} \left[\log \frac{f_k(x_{t+k}, h_t)}{\sum_{s \in \mathcal{S}} f_k(s, h_t)} \right], \quad (1)$$

where $E[\cdot]$ denotes expectation and $f_k(a, b) := \exp(f_{\text{enc}}(a)^T W_k b)$ is a simple log-bilinear model with the W_k being k trainable $d \times d$ matrices.

VQCPC [7] introduces a Vector Quantization (VQ) [10] bottleneck on top of the CPC encoder to obtain a discrete latent representation of the subsequences:

$$\text{VQCPC}(x_i) := z^q(f_{\text{enc}}(x_i)) \in \mathcal{C}, \quad (2)$$

where the *codebook* \mathcal{C} is a set of C centroids partitioning the embedding space R^{d_c} .

3.2. The VQCPC encoder

The VQCPC encoder f_{enc} is a stack of 4 convolutional blocks operating frame-by-frame. Each block is composed of a 1D CNN (in time) with a kernel size of 1 and number of channels (512, 512, 256, d_z) respectively for each block in the stack. Each CNN, except the last one, is followed by a ReLU. As opposed to [24], there is no projection head. VQ is trained with a squared \mathcal{L}_2 loss with a commitment component [10]. We choose a codebook \mathcal{C} containing $C = 16$ centroids, and where $d_c = d_z = 32$. The codebook size is chosen relatively small, enforcing an information bottleneck that only lets through the most salient information needed to discriminate between positive and negative examples [7]. The autoregressive model f_{ar} is a 2-layer GRU with a hidden size of 256 and an output size of 512, and we use its output at timestep t as the context vector h_t to predict $K = 5$ timesteps into the future. The overall training objective is the VQ and the InfoNCE loss (1).

As mentioned earlier, an important choice in contrastive learning is the design of the negative sampling strategy, as this controls what features are represented by the encoding. Usually, the proposal distribution for the negative samples is chosen to be uniform over the training set [7, 20, 25]. However, in this work, we sample 16 negative examples in an intra-sequence fashion: given an audio excerpt \mathbf{x} , the negative examples are all drawn from a uniform distribution over \mathbf{x} (i.e., the same audio excerpt). This intra-sequence sampling forces the network to encode only information which varies *within a sample* (i.e., dynamic information such as onset, offset, loudness change, vibrato, tremolo, etc.), while ignoring static information like instrument type and pitch. This shows that VQCPC provides a convenient way to control what should be

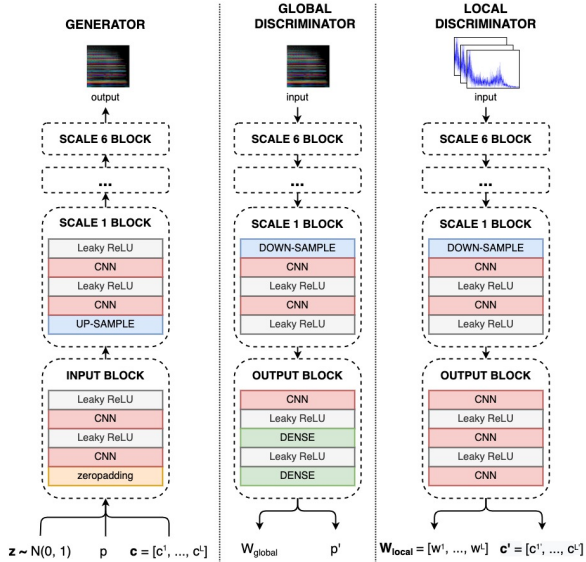


Figure 1: Proposed architecture for VQPC-GAN (see Sec. 3.3).

represented by the discrete representations. In this work, the remaining information (instrument type and pitch) is represented by the GAN’s input noise and the explicit pitch conditioning.

3.3. The GAN architecture

The proposed WGAN is inherited from DrumGAN [5]. We adapt the architecture to a sequential scheme by conducting two major changes. First, the input tensor to the generator G is a sequence containing *static* and *dynamic* information. The *static* information refers to the global context and accounts for the pitch class, a one-hot vector $p \in \{0, 1\}^{26}$ with 26 possible pitch values, as well as a noise vector $z \sim N(0, 1) \in \mathbb{R}^{128}$ sampled from a standard normal distribution with zero mean and unit variance $N(0, 1)$. The *dynamic* information provides local frame-level context and is composed of a sequence of discrete, one-hot vectors $\mathbf{c} = [c^1, \dots, c^L]$ where $c^l \in \{0, 1\}^{16}$ and L is the number of frames in the sequence. The tensor \mathbf{c} identifies a sequence of spectrogram clusters obtained by encoding real audio using VQPC (see Sec. 3.1). At training time, L is set to 32 frames, which corresponds to approximately 1 second of audio given the pre-processing parameters (see Sec. 4.1). The *static* vectors p and z are repeated across the sequence dimension L of the *dynamic* information \mathbf{c} , resulting in a tensor $\mathbf{v} \in \mathbb{R}^{L \times 160}$. This tensor is unsqueezed, reshaped to $(160 \times 1 \times L)$ and fed through a stack of convolutional and nearest-neighbour up-sampling blocks to generate the output signal $x = G_\theta(z; \mathbf{c}; p)$. In order to turn the input tensor into a spectrogram-like convolutional input, it is first zero-padded in the frequency dimension. As depicted in Fig. 1, the generator’s input block performs this zero-padding followed by two convolutional layers with ReLU non-linearity. Each scale block is composed of one nearest-neighbour up-sampling step at the input and two convolutional layers with filters of size (3, 3). The number of feature maps decreases from low to high resolution as {512, 256, 256, 256, 128}. We use Leaky ReLUs as activation functions and apply pixel normalization.

The second major change is the use of two discriminators (see Fig. 1). A local discriminator D_l , implemented in a fully convolutional manner, estimates W_{local} which is the Wasserstein distance [26] between real and generated distributions at a frame-level

(i.e. using batches of frames instead of batches of full spectrograms). Additionally, to encourage G to consider the conditional sequence of VQPC tokens, D_l performs an auxiliary classification task where each input spectrogram frame is assigned to a VQPC token c_l . We add an additional cross-entropy loss term for D_l ’s objective [27]. A global discriminator D_g with two dense layers in its output block estimates W_{global} over complete sequences of $L = 32$ spectrogram frames and predicts the pitch class. As in D_l , we add an auxiliary cross-entropy loss term to D_g ’s objective function for the pitch classification task [27].

3.4. Training Procedure

As in [5], training follows Progressive Growing of GANs (ProGAN) [2] where the architecture is built dynamically during training. The process is divided into scale iterations wherein each scale a new block is added to G and D (see Fig. 1). The models are trained 200k iterations per scale using batches of 30, 30, 20, 12, and 12, respectively. We employ Adam as the optimization method with a learning rate of 0.001 in both networks.

4. EXPERIMENTS

In this work, we employ a VQPC encoder (see Sec. 3.1) to learn discrete sequences of high-level features from a dataset of tonal sounds (see Sec. 4.1). As described in Sec. 3.3, we condition a GAN on such discrete sequential representations in order to perform audio synthesis. Variable-length audio is achieved by up/down-sampling, respectively for longer or shorter sounds, of the conditional VQPC sequence. In the following, we present the training dataset, the evaluation metrics and the baselines.

4.1. Dataset

We employ a subset of audio excerpts obtained from the NSynth dataset [28]. The subset contains over 50k single-note audios played by 11 instrument families. Among the available annotations, we make use of the pitch class. The (monophonic, 16kHz) audio clips are trimmed to 1 second. We only consider samples with a pitch in the MIDI range from 44 to 70 (103.83 - 466.16 Hz). For the evaluation, we use a 90/10% split of the data.

Previous works indicated that magnitude and Instantaneous Frequency of the Short-time Fourier Transform (STFT) are well-suited for structuring tonal audio information [29, 4, 30]. We, therefore, preprocess the data with this transform using an FFT size of 2048 bins and an overlap of 75%. For the VQPC encoder (see 3.2), we rely on the Constant-Q Transform (CQT) spanning 6 octaves with 24 bins per octave. We use a hop-length of 512 samples, for the output token sequence to match the temporal resolution of the data used to train the GAN.

4.2. Evaluation

Similarly to previous works on adversarial audio synthesis [4, 29, 5, 31], we evaluate our models using the following metrics:

- The **Inception Score (IS)** [32] is defined as the mean KL divergence between the conditional class probabilities $p(y|x)$, and the marginal distribution $p(y)$ using the class predictions of an Inception classifier. IS penalizes models (low IS) that generate examples that cannot be easily classified into a single class with high confidence as well as models whose examples belong to only a few of all possible classes. We train our own inception model to classify pitch and instrument on the NSynth dataset yielding two metrics: the Pitch Inception Score (PIS) and Instrument Inception Score (IIS) [29].

- The **Kernel Inception Distance (KID)** [33] is defined as the squared Maximum Mean Discrepancy (MMD) between embeddings of real and generated data in a pre-trained Inception-like classifier. The MMD measures dissimilarity between real and generated distributions therefore the lower the better. We compute the KID with the same Inception model as in IS.
- The **Fréchet Audio Distance (FAD)** [34] is computed as the difference between multivariate Gaussian distributions fitted to the output embedding of a pre-trained VGG-like model. A lower FAD indicates a smaller distance between real and generated distributions. We employ Google’s FAD implementation.³

4.3. Baselines

We compare the metrics described in Sec. 4.2 with a few baselines and include results scored by *real data* to delimit the range of each metric. Specifically, we compare with GANSynth [4], obtained from Google Magenta’s github,⁴ and two baselines that we train using the same architecture of VQPCP-GAN but removing the sequence generation scheme, i.e., without VQPCP conditioning nor the local D . We train the two baseline models, $WGAN_{1s}$ and $WGAN_{4s}$, on 1s and 4s-long audio excerpts, respectively, whereas GANSynth is originally trained on 4s audio excerpts. As mentioned early on in this section, we condition VQPCP-GAN on varying-length VQPCP sequences in order to generate audio with different duration. To do so, we just up/down-sample the VQPCP sequence accordingly to obtain the desired number of output frames. In particular for these experiments, we take the original VQPCP sequences of length 32 (i.e., 1s-long) and perform nearest-neighbour up-sampling by a factor of 4 to obtain 128 tokens (i.e., 4s-long).

5. RESULTS AND DISCUSSION

In this section, we present the results from the evaluation metrics described in Sec. 4.2. We informally validate these quantitative results by listening to the generated content and sharing our assessment.

5.1. Quantitative results

Table 5.2 presents the metrics scored by our proposed VQPCP-GAN and the baselines. Overall, our WGANs score closest to those of *real data* in most metrics, or even better in the case of the PIS. GANSynth follows closely and VQPCP-GAN obtains slightly worse results. VQPCP-GAN performs particularly good in terms of PIS, which suggests that the generated examples have an identifiable pitch content and that the distribution of pitch classes follows that of the training data. This is not surprising given that the model has explicit pitch conditioning, making it trivial to learn the specific mapping between the pitch class and the respective tonal content. Conversely, results are worse in the case of IIS, suggesting that the model failed to capture the timbre diversity existent in the dataset and that generated sounds cannot be reliably classified into one of all possible instrument types (i.e. mode failure). Turning now our attention to the KID, VQPCP-GAN scores results very similar to GANSynth and slightly worse than $WGAN_{1s}$. A low KID indicates that the Inception embeddings are similarly distributed for real and generated data. Our Inception classifier is trained on several discriminative tasks of specific timbral attributes, including pitch and instrument classification. Therefore, we can infer that similarities in

³https://github.com/google-research/google-research/tree/master/frechet_audio_distance

⁴<https://github.com/magenta/magenta/tree/master/magenta/models/gansynth>

duration (s)	IIS \uparrow		PIS \uparrow		KID ^a \downarrow		FAD \downarrow	
	1	4	1	4	1	4	1	4
<i>real data</i>	6.3	4.5	17.9	18.0	6.7	6.6	0.0	0.0
GANSynth [4]	-	4.1	-	19.7	-	7.0	-	2.1
$WGAN_{1s}$	4.5	-	19.0	-	6.8	-	0.8	-
$WGAN_{4s}$	-	4.5	-	20.1	-	6.9	-	1.0
VQPCP-GAN	3.0	2.9	18.5	17.2	7.3	7.1	5.6	5.4

^a $\times 10^{-4}$

Table 1: IIS, PIS, KID, and FAD (Sec. 4.2), scored by VQPCP-GAN and baselines. The metrics are computed over 25k samples.

such embedding space indicate shared timbral and tonal characteristics, from a statistical point of view, between real and generated audio data. This trend is not as evident in the case of the FAD, where VQPCP-GAN obtains considerably worse results than the baselines, particularly in the case of $WGAN_{1s}$. This could indicate the existence of artefacts as FAD was found to correlate well with several artificial distortions [34].

To wrap up: despite the architectural changes introduced for sequential generation, VQPCP-GAN exhibits results comparable to GANSynth, the SOTA on adversarial audio synthesis of tonal sounds, as well as two strong baselines $WGAN_{1,4s}$ trained on 1 and 4-second long audio respectively. Notably, our $WGAN_{4s}$ baseline scores better results than GANSynth in all metrics. In the following section, we informally validate these quantitative results by sharing our assessment when listening to generated audio material.

5.2. Informal listening

The accompanying website contains audio examples generated under different settings (e.g. latent interpolations, pitch scales, generation from MIDI files) and different duration (0.5, 1, 2 and 4 seconds). Synthesis of variable-length audio is achieved by up/down-sampling of the conditional VQPCP sequence. Overall, we find the results discussed in Sec. 5.1 to align well with our perception. The range of instruments is narrow, and only a few from the most homogeneous and populated classes in the dataset can be identified (e.g., mallet, guitar, violin), hence the low IIS. In the *pitch scale* examples, we can perceive that the pitch content responds nicely to the conditional signal, and it is consistent across the generation time span, which explains the higher PIS. Although we eventually obtain some artifacts when using certain VQPCP token combinations as conditional input, the overall quality is acceptable. This is aligned with having a low FAD but a KID comparable to the baselines.

6. CONCLUSION

In this work, we presented VQPCP-GAN, an adversarial model capable of performing variable-length sound synthesis of tonal sounds. We adapted the WGAN architecture found in previous works [4, 5] to a sequential setting by conducting two major architectural changes. First, we condition G on *dynamic* and *static* information captured, respectively, by a sequence of discrete tokens learned through VQPCP, and a global noise z . Additionally, we introduce a secondary fully-convolutional D that discriminates between real and fake data distributions at a frame level and predicts the VQPCP token associated with each frame. Results showed that VQPCP-GAN can generate variable-length sounds with controllable pitch content while still exhibiting results comparable to previous works generating audio with fixed-duration. We provide audio examples in the accompanying website. As future work, we plan on investigating hierarchical VQPCP tokens to condition the GAN on longer-term, compact representations of audio signals.

7. ACKNOWLEDGMENTS

The authors thank Gaëtan Hadjeres for his guidance and support.

8. REFERENCES

- [1] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. C. Courville, and Y. Bengio, "Generative adversarial nets," in *NeurIPS*, Montreal, Quebec, Canada, Dec. 2014.
- [2] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of gans for improved quality, stability, and variation," in *ICLR*, May 2018.
- [3] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of stylegan," in *CVPR*, June 13-19 2020.
- [4] J. Engel, K. K. Agrawal, S. Chen, I. Gulrajani, C. Donahue, and A. Roberts, "Gansynth: Adversarial neural audio synthesis," in *ICLR*, May 2019.
- [5] J. Nistal, S. Lattner, and G. Richard, "Drumgan: Synthesis of drum sounds with timbral feature conditioning using generative adversarial networks," in *ISMIR*, 2020.
- [6] M. Binkowski, J. Donahue, S. Dieleman, A. Clark, E. Elsen, N. Casagrande, L. C. Cobo, and K. Simonyan, "High fidelity speech synthesis with adversarial networks," in *ICLR*, Apr. 2020.
- [7] G. Hadjeres and L. Crestel, "Vector quantized contrastive predictive coding for template-based music generation," *CoRR*, 2020.
- [8] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," in *ISCA Speech Synthesis Workshop*, Sept. 2016.
- [9] Y. Fan, Y. Qian, F. Xie, and F. K. Soong, "TTS synthesis with bidirectional LSTM based recurrent neural networks," in *INTERSPEECH*, Sept. 2014.
- [10] A. van den Oord, O. Vinyals, and K. Kavukcuoglu, "Neural discrete representation learning," in *NeurIPS*, Long Beach, CA, USA, Dec. 2017.
- [11] J. Kong, J. Kim, and J. Bae, "Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis," in *NeurIPS*, Dec. 2020.
- [12] K. Kumar, R. Kumar, T. de Boissiere, L. Gestein, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. C. Courville, "Melgan: Generative adversarial networks for conditional waveform synthesis," in *NeurIPS*, Vancouver, BC, Canada, Dec. 2019.
- [13] C. Donahue, J. McAuley, and M. Puckette, "Adversarial audio synthesis," in *ICLR*, May 2019.
- [14] J. Drysdale, M. Tomczak, and J. Hockman, "Adversarial synthesis of drum sounds," in *DAFX*, 2020.
- [15] E. Hosseini-Asl, Y. Zhou, C. Xiong, and R. Socher, "A multi-discriminator cyclegan for unsupervised non-parallel speech domain adaptation," in *INTERSPEECH*, Hyderabad, India, Sept. 2018.
- [16] A. Biswas and D. Jia, "Audio codec enhancement with generative adversarial networks," in *ICASSP*, Barcelona, Spain, May 2020.
- [17] L. Yu, W. Zhang, J. Wang, and Y. Yu, "Seqgan: Sequence generative adversarial nets with policy gradient," in *AAAI Conference on Artificial Intelligence*, San Francisco, California, USA, Feb. 2017.
- [18] O. Mogren, "C-RNN-GAN: continuous recurrent neural networks with adversarial training," *CoRR*, 2016.
- [19] J. Yoon, D. Jarrett, and M. van der Schaar, "Time-series generative adversarial networks," in *NeurIPS*, Vancouver, BC, Canada, Dec. 2019.
- [20] A. van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *CoRR*, 2018.
- [21] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," in *INTERSPEECH*, Graz, Austria, Sept. 2019.
- [22] A. Baevski, S. Schneider, and M. Auli, "vq-wav2vec: Self-supervised learning of discrete speech representations," in *ICLR*, Addis Ababa, Ethiopia, Apr. 2020.
- [23] B. van Niekerk, L. Nortje, and H. Kamper, "Vector-quantized neural networks for acoustic unit discovery in the zerospeech 2020 challenge," in *INTERSPEECH*, Shanghai, China, Oct. 2020.
- [24] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
- [25] O. J. Hénaff, A. Srinivas, J. D. Fauw, A. Razavi, C. Doersch, S. M. A. Eslami, and A. van den Oord, "Data-efficient image recognition with contrastive predictive coding," *CoRR*, vol. abs/1905.09272, 2019. [Online]. Available: <http://arxiv.org/abs/1905.09272>
- [26] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of wasserstein gans," in *NeurIPS*, Long Beach, CA, USA, Dec. 2017.
- [27] A. Odena, C. Olah, and J. Shlens, "Conditional image synthesis with auxiliary classifier gans," in *ICML*, Sydney, NSW, Australia, Aug. 2017.
- [28] J. Engel, C. Resnick, A. Roberts, S. Dieleman, M. Norouzi, D. Eck, and K. Simonyan, "Neural audio synthesis of musical notes with wavenet autoencoders," in *ICML*, Aug. 2017.
- [29] J. Nistal, S. Lattner, and G. Richard, "Comparing representations for audio synthesis using generative adversarial networks," in *EUSIPCO*, Amsterdam, NL, Jan. 2021.
- [30] T. Bazin, G. Hadjeres, P. Esling, and M. Malt, "Spectrogram inpainting for interactive generation of instrument sounds," *CoRR*, 2021.
- [31] C. Gupta, P. Kamath, and L. Wyse, "Signal representations for synthesizing audio textures with generative adversarial networks," *CoRR*, 2021.
- [32] T. Salimans, I. J. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved techniques for training gans," in *NeurIPS*, Barcelona, Spain, Dec. 2016.
- [33] M. Binkowski, D. J. Sutherland, M. Arbel, and A. Gretton, "Demystifying MMD gans," in *ICLR*, Vancouver, BC, Canada, Apr. 2018.
- [34] K. Kilgour, M. Zuluaga, D. Roblek, and M. Sharifi, "Fréchet audio distance: A metric for evaluating music enhancement algorithms," *CoRR*, 2018.