NEURAL PERCUSSIVE SYNTHESIS PARAMETERISED BY HIGH-LEVEL TIMBRAL FEATURES

António Ramires¹, Pritish Chandna¹, Xavier Favory¹, Emilia Gómez^{1,2}, Xavier Serra¹

¹Music Technology Group, Universitat Pompeu Fabra, Barcelona, Spain

²Joint Research Centre, European Commission, Seville, Spain

ABSTRACT

We present a deep neural network-based methodology for synthesising percussive sounds with control over high-level timbral characteristics of the sounds. This approach allows for intuitive control of a synthesizer, enabling the user to shape sounds without extensive knowledge of signal processing. We use a feedforward convolutional neural network-based architecture, which is able to map input parameters to the corresponding waveform. We propose two datasets to evaluate our approach on both a restrictive context, and in one covering a broader spectrum of sounds. The timbral features used as parameters are taken from recent literature in signal processing. We also use these features for evaluation and validation of the presented model, to ensure that changing the input parameters produces a congruent waveform with the desired characteristics. Finally, we evaluate the quality of the output sound using a subjective listening test. We provide sound examples and the system’s source code for reproducibility.

Index Terms— Wave-U-Net, Percussive Sound Synthesis, Generative Models, Music Information Retrieval, Creative Interfaces

1. INTRODUCTION

Percussion is one of the main components in music and is normally responsible for a song’s rhythm section. Classic percussion instruments create sound when struck or scraped; however new electronic instruments were developed for generating these sounds either through playing prerecorded samples or through synthesising them. These are called drum machines and became very popular for electronic music [1]. However, these early drum machines did not provide much control over the generation of the sounds. With the developments in digital audio technology and computer music, new

drum machines were hand-designed using expert knowledge on synthesis techniques and electronic music production.

With the success of deep learning, several innovative generative methodologies have been proposed in the recent years. These methodologies include Generative Adversarial Networks (GANs) [2], Variational Autoencoders (VAEs) [3] and autoregressive networks [4, 5]. In the audio domain, such methodologies have been applied for singing voice [6], instrumental sounds [5] and drum sound generation [7]. However, in the case of percussive sounds, the proposed methods only allow the user to navigate in non-intuitive high dimensional latent spaces.

The aim of our research is to create a single-event percussive-sound synthesizer that can be intuitively controlled by users, despite their sound design knowledge. This requires both a back end of a generative model that is able to map the user controls to the output sound and a front end user interface. In this paper, we propose a generative methodology based on the Wave-U-Net architecture [8]. Our method maps high-level characteristics of sounds to the corresponding waveforms. The use of these features is aimed at giving the end-user intuitive control over the sound generation process. We also present a dataset of 10 000 percussive one-shot sounds collected from Freesound [9], curated specially for this study.

The source code for our model is available online¹, as are sound examples², showcasing the robustness of the models.

2. GENERATIVE MODELS FOR AUDIO

In the audio domain, several generative models have been proposed over the recent years. In the context of music, generative models have shown success specially in creating pitched instrumental sounds, when conditioned on musical notes. A pioneering work on this field was NSynth [5]. This synthesizer is based on the wavenet vocoder [10] an autoregressive architecture, which, while capable of generating high quality sounds is very resource intensive. Several alternate architectures have been used for the generation of musical

This work is partially funded by the European Unions Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No765068, MIP-Frontiers. This work is partially supported by the Towards Richer Online Music Public-domain Archives (TROMPA) project. The TITANX used for this research was donated by the NVIDIA Corporation.

¹https://github.com/pc2752/percussive_synth

²https://pc2752.github.io/percussive_synth/

notes, based on GANs [11, 12], VAEs [13, 7, 14], Adversarial AEs [15] and AEs with Wavenet [5].

For percussive sound synthesis, the most relevant work is the Neural Drum Machine [7], which uses a Conditional Wasserstein Auto Encoder [16], trained on the magnitude component of the spectrogram of percussive sounds coupled with a Multi-Head Convolutional Neural Network for reconstructing the audio from the spectral representation. Principal Component Analysis is used on the low-dimensional representation learned by the AE to select the 3 most influential dimensions of the 64 dimensions of the embedding. These are provided to the user over a control interface. However these parameters controlled by the user are abstract and are not shown to be perceptually relevant or semantically meaningful.

In our case, we wish to directly map a chosen set of features to the output sound. The WaveNet [4] architecture has been shown to generate high quality waveforms conditioned on input features. However, the autoregressive nature of the model makes it resource extensive and the short nature of percussive sounds do not require the use of a long temporal model. Therefore, for our study, we decided to use the Wave-U-Net [8] architecture, which has been shown to effectively model waveforms in the case of source separation and follows a feedforward convolutional architecture, making it resource efficient. The model takes as input the waveform of the mixture of sources, downsamples it through a series of convolutional operations to generate a low dimensional representation and then upsamples it through linear interpolation followed by convolution to the output dimensions. There are concatenative connections between the corresponding layers of the upsampling and downsampling blocks. In our work, we adapt this architecture to fit the desired use case.

3. TIMBRAL FEATURES

For our end goal, we require semantically meaningful features that can allow for intuitive control of the synthesizer. In the field of Music Information Retrieval, a strong effort has been put on developing hand crafted features which can characterise sounds. These features enable users to retrieve sounds or music from large audio collections by automatically describing them according to their timbre, their mood, or other characteristics which are easy to understand by users. For our purpose, we need features pertaining to timbre. We understand timbre as pertaining to perceptual characteristics of sounds analogous to colour or quality. Control over such features would enable the user to intuitively shape sounds.

A set of such features have been proposed in [17], where recurrent query terms, related to timbral characteristics, used for searching sounds in large audio databases were identified. Regression models were developed by mapping user-collected ratings to timbral characteristics, which quantify semantic attributes. These are hardness, depth, brightness,

roughness, boominess, warmth and sharpness. The work proposes feature extractors pertaining to these query terms and we use an open-source implementation of the same³. For the rest of this paper, we refer to the 7 features extracted by these extractors as timbral features.

Another relevant characteristic which is commonly present in drum synthesizers and music makers are used to work with is the temporal envelope of the sound. This feature describes the energy of the sounds over time and is normally available to users in drum synthesizers as a set of *attack* and *decay* controls. We use an open-source implementation of the envelope algorithm described in [18], present in the Essentia library [19]. An attack time of 5 ms and a release time of 50 ms was used to generate a smooth curve which matched the sound energy over time. It must be noted that the timbral features described previously are summary features, i.e. have a single value for each sound while the envelope is time evolving and of the same dimensions as the waveform.

4. DATASET CURATION

We curated two datasets in order to train our models in different scenarios. The first consists of sounds taken from Freesound, a website which hosts a collaborative collection of Creative Commons licensed sounds⁴ [9]. We performed queries to the database with the name of percussion instruments as keywords in order to retrieve a set of percussive sounds, with a limit on effective duration of 1 s. We then conducted a manual verification of these sounds⁵: to select the ones that were containing one single event, and were of appreciable quality in the context of traditional electronic music production. This process created a dataset of around 10 000 sounds, containing instruments such as kicks, snares, cymbals and bells. The dataset is publicly available in a Zenodo repository⁶. For the rest of this paper, we refer to this dataset as FREESOUND.

A second dataset was created by aggregating about 5000 kick drum one-shot samples from our personal collections, originating mostly from commercial libraries. This type of sounds are often of high quality, annotated and contain only one event which makes it very handy to construct a dataset of isolated sounds, suiting our needs for training our model in a restricted context. We refer to this dataset as KICKS.

The aim of creating two datasets was to understand if our method could be applicable for synthesising a wide variety of percussion sounds, or if it was more appropriate to focus on synthesising only one type of sounds, in this case the kick drum.

³<https://github.com/AudioCommons/ac-audio-extractor>

⁴<https://freesound.org>

⁵We developed an annotation tool available at this repository <https://github.com/xavierfav/percussive-annotator>.

⁶<https://zenodo.org/record/3665275>

5. METHODOLOGY

We aim to model the probability distribution of the waveform x as a function of the timbral features fs and the time-domain envelope e . To this end we use a feedforward convolutional neural network as a function approximator to model $P(x|fs, e)$. We use a U-Net architecture, similar to the one used by [8], which has been shown to effectively model the waveform of an audio signal. Our network takes the envelope as input and concatenates to it the timbral features, fs , broadcast to the input dimensions, as done by [4]. As shown in Figure 1, downsampling is done via a series of convolutions with stride 2, to produce a low-dimension embedding. We use a filter length of 5 and double the number of filters after each 3 layers, starting with 32 filters. A total of 15 layers are used in the encoder, leading to an embedding of size 512. We upsample this low dimensional embedding sequentially to the output $\hat{x}$, using linear interpolation followed by convolution. This mirrors the approach used by [6, 8] and has been shown to avoid high frequency artefacts which appear while upsampling with transposed convolutions. As with the U-Net, there are connections between the corresponding layers of the encoder and decoder, as shown in Figure 1.

Fig. 1. The proposed architecture, with $K = 15$ layers.

We initially used a simple reconstruction loss function, shown in equation 1 to optimise the network.

$$\mathcal{L}_{recon} = \mathbb{E}[\|\hat{x} - x\|_1] \quad (1)$$

While this resulted in a decent output, we noticed that the network was able to reproduce the low frequency components of the desired sound, but lacked details in high frequency components. To rectify this, we added a short time fourier transform ($STFT$) based loss, similar to [20]. This loss is shown in equation 2.

$$\mathcal{L}_{stft} = \mathbb{E}[\|STFT(\hat{x}) - STFT(x)\|_1] \quad (2)$$

The final loss of this network is shown in equation 3.

$$\mathcal{L}_{final} = \mathcal{L}_{recon} + \lambda \mathcal{L}_{stft} \quad (3)$$

Fig. 2. A sample of the input envelope and features and the output waveforms for the various models for the KICK dataset

Where λ is the weight given to the high frequency component of the reconstruction.

6. EXPERIMENTS

6.1. Data Pre-processing

All sound were downsampled to a sampling rate of 16 kHz and silences were removed from the beginning and end of the sounds. Following this, we calculated the timbral features and envelope described in section 3 and then zero-padded at the end of the sound to 16 000 samples. The features were normalised using min-max normalisation, to ensure that the inputs were within the range 0 to 1

6.2. Training the network

The network was trained using the Adam optimiser [21] for 2500 epochs with a batch size of 16. We use 90 % of the data for training and 10 % for evaluation. The $STFT$ used for the $\mathcal{L}_{stft}$ loss function is calculated over 1024 samples and a hopsize of 512. With the given sampling rate, this led to a frequency resolution of 16.125 Hz per bin. We evaluate the model with three losses: the $\mathcal{L}_{recon}$ loss, henceforth referred to as WAVE; the $\mathcal{L}_{final}$, referred to as FULL; and a version with only the high frequency components of the $STFT$ for the $\mathcal{L}_{stft}$, referred as HIGH. This last model uses $STFT$ components above 650 Hz or 40 bins as traditional kick synthesizers model a kick sound via a low frequency sinusoid, generally below 650 Hz with some high frequency noise. We use $\lambda = 0.5$ for our experiments.

6.3. Evaluation

The proposed models need to be evaluated in terms of the perceived audio quality and the coherence of timbral features between the input and the output. A preliminary assessment of the quality of reconstruction can be made by looking at the output waveforms, shown in Figure 2 for a sample from the test set of the KICKS dataset. Although the reconstruction

seems to be visually accurate for the three models, the perceived quality of the audio is a subjective metric that cannot be judged by simply looking at the plots. We can objectively assess the coherence of the timbral features used as input to the model. More importantly, we want to assess that a change in these features leads to a corresponding change in the output.

To this end, we vary each individual timbral feature while maintaining the other features constant. We then check the accuracy of the output waveform via the same feature extractors used for training. For each individual feature, we set values of *low*, corresponding to 0.2 over the normalised scale, *mid*, corresponding to 0.5 and *high*, corresponding to 0.8. The respective outputs for such models are termed $\hat{x}_{low}^i$, $\hat{x}_{mid}^i$ and $\hat{x}_{high}^i$ and their corresponding features are f_{low}^i , f_{mid}^i and f_{high}^i for the i^{th} feature. For coherent modelling, the models should follow the order $f_{high}^i > f_{mid}^i > f_{low}^i$. We assess the accuracy of this order in three tests, E1, which checks the condition $f_{high}^i > f_{low}^i$, E2, which checks $f_{high}^i > f_{mid}^i$ and E3, which checks $f_{mid}^i > f_{low}^i$ over all values of i . The accuracy of the models over these tests is shown in Table 1 and a feature wise summary is shown in Table 2.

Dataset	Model	Accuracy		
		E1	E2	E3
FREESOUND	WAVE	0.601	0.569	0.552
	HIGH	0.649	0.601	0.657
	FULL	0.825	0.758	0.780
KICKS	WAVE	0.805	0.722	0.722
	HIGH	0.876	0.789	0.769
	FULL	0.920	0.814	0.798

Table 1. Objective verification of feature coherence across models and datasets.

It can be seen that the FULL model, followed by the HIGH, are the most efficient at mapping the input features to the output waveform in terms of feature coherence, but all three models do maintain this coherence to a high degree.

Feature	FREESOUND			KICKS		
	E1	E2	E3	E1	E2	E3
Boominess	0.98	0.82	0.98	0.96	0.86	0.95
Brightness	0.99	0.99	1.00	0.99	0.98	0.84
Depth	0.94	0.65	0.94	0.99	0.89	0.94
Hardness	0.64	0.66	0.59	0.85	0.61	0.79
Roughness	0.63	0.59	0.57	0.84	0.80	0.62
Sharpness	0.63	0.77	0.45	0.90	0.91	0.54
Warmth	0.92	0.79	0.91	0.88	0.61	0.87

Table 2. Objective verification of the accuracy on feature coherence for the best performing models for each dataset.

While feature coherence is maintained for features like boominess, brightness, depth and warmth for the full dataset,

the models are less consistent in terms of hardness, roughness and sharpness, particularly true for the FREESOUND dataset.

Given the absence of a suitable baseline system, we decided to use an online AB listening test that compared the models amongst themselves and a reference for subjective evaluation of quality. The participants of the test were presented with 15 examples each from both datasets. Each example had two options, A and B from two of the models used for the dataset, along with a reference ground truth audio. There were 5 examples each from each of the 3 pairs. The participant was asked to choose the audio clip which was closest in quality to the reference audio. There were 35 participants in the listening test, the results of which are shown in Figure 3.

Fig. 3. Results of the listening test, displaying the user preference between loss functions for each of the datasets.

A clear preference for the HIGH model can be seen, especially for the KICKS dataset. This can be attributed partly to the choice of cutoff frequency used in the model and partly to the diversity of sounds in the FREESOUND dataset. We note the difficulty in assessing audio quality over printed text and encourage the user to visit our demo page and listen to the audio samples for assessment.

7. CONCLUSIONS AND FUTURE WORK

In this work, we proposed a method using a feedforward convolutional neural network based on the Wave-U-Net [8] for synthesising percussive sounds conditioned on semantically meaningful features.

Our final aim is to create a system that can be controlled using high-level parameters, being semantically meaningful characteristics that correspond to concepts casual music makers are familiar with. To this end, we use hand crafted features designed by MIR experts and curate and present a dataset for the purpose of modelling percussive sounds. Via objective evaluation, we were able to verify that the control features do indeed modify the output waveform as desired and quality assessment was done via an online listening test.

Future work will focus on developing an interface for interacting with the synthesizer, which will allow to evaluate the approach in its context of use, with real users.

8. REFERENCES

- [1] Zainab Hasnain, “How the roland tr-808 revolutionized music,” Apr 2017.
- [2] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, “Generative adversarial nets,” in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [3] Diederik P. Kingma and Max Welling, “Auto-encoding variational bayes,” in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- [4] 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 *The 9th ISCA Speech Synthesis Workshop, Sunnyvale, CA, USA, 13-15 September 2016*, 2016, p. 125.
- [5] Jesse Engel, Cinjon Resnick, Adam Roberts, Sander Dieleman, Mohammad Norouzi, Douglas Eck, and Karen Simonyan, “Neural audio synthesis of musical notes with wavenet autoencoders,” in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*. 2017, ICML’17, pp. 1068–1077, JMLR.org.
- [6] Pritish Chandna, Merlijn Blaauw, Jordi Bonada, and Emilia Gomez, “Wgansing: A multi-voice singing voice synthesizer based on the wasserstein-gan,” 2019.
- [7] Cyran Aouameur, Philippe Esling, and Gaëtan Hadjeres, “Neural drum machine : An interactive system for real-time synthesis,” 2019.
- [8] Daniel Stoller, Sebastian Ewert, and Simon Dixon, “Wave-u-net: A multi-scale neural network for end-to-end audio source separation,” in *Proceedings of the 19th International Society for Music Information Retrieval Conference, ISMIR 2018, Paris, France, September 23-27, 2018*, 2018, pp. 334–340.
- [9] Frederic Font, Gerard Roma, and Xavier Serra, “Freesound technical demo,” in *Proceedings of the 21st ACM international conference on Multimedia*. ACM, 2013, pp. 411–412.
- [10] Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerrv-Ryan, et al., “Natural tts synthesis by conditioning wavenet on mel spectrogram predictions,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [11] Chris Donahue, Julian J. McAuley, and Miller S. Puckette, “Synthesizing audio with generative adversarial networks,” 2018.
- [12] Jesse Engel, Kumar Krishna Agrawal, Shuo Chen, Ishaan Gulrajani, Chris Donahue, and Adam Roberts, “GANSynth: Adversarial neural audio synthesis,” in *International Conference on Learning Representations*, 2019.
- [13] Philippe Esling and Adrien Bitton, “Bridging audio analysis, perception and synthesis with perceptually-regularized variational timbre spaces,” in *Proceedings of the 19th International Society for Music Information Retrieval Conference (ISMIR 2018)*, 2018.
- [14] Adrien Bitton, Philippe Esling, and Axel Chemla-Romeu-Santos, “Modulated variational auto-encoders for many-to-many musical timbre transfer,” 2018.
- [15] Adrien Bitton, Philippe Esling, Antoine Caillon, and Martin Fouilleul, “Assisted sound sample generation with musical conditioning in adversarial auto-encoders,” 2019.
- [16] Ilya O. Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Schölkopf, “Wasserstein auto-encoders,” in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018.
- [17] Andy Pearce, Tim Brookes, and Russell Mason, “Timbral attributes for sound effect library searching,” in *Audio Engineering Society Conference: 2017 AES International Conference on Semantic Audio*, Jun 2017.
- [18] Udo Zölzer, *Digital audio signal processing*, vol. 9, Wiley Online Library, 2008.
- [19] Dmitry Bogdanov et al., “Essentia: an audio analysis library for music information retrieval,” in *International Society for Music Information Retrieval Conference (ISMIR’13)*, Curitiba, Brazil, 04/11/2013 2013, pp. 493–498.
- [20] Abhimanyu Sahai, Romann Weber, and Brian McWilliams, “Spectrogram feature losses for music source separation,” 2019.
- [21] Diederik P. Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.