



This is an electronic reprint of the original article. This reprint may differ from the original in pagination and typographic detail.

Wright, Alec; Välimäki, Vesa; Juvela, Lauri

Adversarial Guitar Amplifier Modelling with Unpaired Data

Published in: ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)

DOI: 10.1109/ICASSP49357.2023.10094600

Published: 10/06/2023

Document Version Peer-reviewed accepted author manuscript, also known as Final accepted manuscript or Post-print

Please cite the original version:

Wright, A., Välimäki, V., & Juvela, L. (2023). Adversarial Guitar Amplifier Modelling with Unpaired Data. In ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 1-5). Article 10094600 (Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing). IEEE. https://doi.org/10.1109/ICASSP49357.2023.10094600

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

ADVERSARIAL GUITAR AMPLIFIER MODELLING WITH UNPAIRED DATA

Alec Wright¹, Vesa Välimäki^{*,1}, Lauri Juvela^{1,2}

¹Acoustics Lab, Dept. Information & Communications Eng. Aalto University, Espoo, Finland ²Neural DSP Technologies, Helsinki, Finland

ABSTRACT

We propose an audio effects processing framework that learns to emulate a target electric guitar tone from a recording. We train a deep neural network using an adversarial approach, with the goal of transforming the timbre of a guitar, into the timbre of another guitar after audio effects processing has been applied, for example, by a guitar amplifier. The model training requires no paired data, and the resulting model emulates the target timbre well whilst being capable of real-time processing on a modern personal computer. To verify our approach we present two experiments, one which carries out unpaired training using paired data, allowing us to monitor training via objective metrics, and another that uses fully unpaired data, corresponding to a realistic scenario where a user wants to emulate a guitar timbre only using audio data from a recording. Our listening test results confirm that the models are perceptually convincing.

Index Terms— Audio systems, deep learning, generative adversarial networks, music, nonlinear systems, unsupervised learning.

1. INTRODUCTION

When recording musical instruments for music production, audio effects are an essential component of the resulting timbre. This is especially true of electric guitar, where the timbre, or as it is often referred to by guitarists, *tone*, imparted by certain amplifiers and effects pedals is highly sought-after. Digital emulation of analog audio devices such as these is known as Virtual Analog (VA) modelling [1].

A related but more challenging problem is to emulate the timbre of an electric guitar directly from a recording. This corresponds to a practical and very exciting task: how to imitate the guitar tone heard on a commercial music recording using your own electric guitar? We refer to this as modelling from *unpaired data*, as the unprocessed guitar signal used to generate the target guitar tone is unavailable.

The usual paradigm of VA modelling is the emulation of certain analog devices, however this is insufficient for this purpose for two reasons. Firstly, the specific instrument and devices used to craft the guitar tone on a famous recording might not be known. Even if the required equipment is available, recreating the recording setup is a time consuming process that requires expert knowledge.

Secondly, the desired timbre includes both the musical instrument itself, as well as any processing applied to it. This means that any emulation of a recorded guitar tone must also account for differences between the timbre of your guitar, and the timbre of the target guitar, before effects processing is applied.

The problem thus corresponds more closely to that of timbre transfer, but in the specific case where the input and target instrument are both guitars. The problem can be divided into two sub-problems: the guitar-to-guitar transformation, and modelling of the

guitar signal chain. Previous work has approached guitar-to-guitar transformation using linear filtering, for example, to process an electric guitar such that it sounds like an acoustic guitar [2].

The guitar signal chain modelling problem involves the emulation of all the devices, other than the instrument itself, used to create the recording. These might include a guitar amplifier, speaker cabinet, and other effects such as compression or distortion. Modelling these types of devices has been widely studied, and broadly falls into two approaches, "white-box", where equations describing the system behaviour are derived [3, 4, 5], or "black-box", where data collected from the system is used to fit a generic model [6, 7, 8]. Recently, deep learning methods based on neural networks have become a popular choice [9]. These include architectures for modelling a wide range of effects [10], feedforward [11], and recurrent [12] models for guitar amplifiers and time-varying effects [13], as well as models using DSP components such as linear filters [14, 15, 16].

Whilst the aforementioned approaches are all viable for the task of imitating a target electric guitar timbre, they all require us either to know what devices were used during recording, or to have access to *paired data* with both the target timbre and the unprocessed signal taken directly from the guitar used during the recording. In the case where paired training data is impossible, or unavailable, one solution is to formulate the problem as domain transfer, similar to *Image Style Transfer*, where the style of an input image is transformed to match a target style, whilst retaining the input image content [17, 18].

A related problem in audio is *Audio Style Transfer*, which seeks to transfer the style of some input audio to a target style, whilst retaining the content of the input audio [19]. In the speech domain, an example of this is voice conversion, which can be achieved, for example, using non-parallel speech data to train a Generative Adversarial Network (GAN) with a cycle-consistency loss [20].

For musical applications the problem is frequently referred to as *timbre transfer* [21, 22] or *tone transfer* [23]. This can be achieved by applying image style transfer techniques to a time-frequency representation of the input audio, and then re-synthesising the audio in the waveform domain [21], or by learning a high level representation of the input and using that as input to a synthesizer [24]. Recent work has also proposed style transfer of audio effects, to impart the production style of one recording onto another [25].

In this paper, we present a method to process a signal recorded directly from an electric guitar, or *input timbre*, such that its timbre matches that of a guitar from a recording, or *target timbre*. We solve this domain transfer problem using an unsupervised approach, in which a black-box feedforward convolutional model is trained adversarially using unpaired data. The distinction between supervised and unsupervised modelling is shown in Fig. 1.

The rest of this paper is structured as follows. Sec. 2 discusses the modelling approach, training objective and data. Sec. 3 describes two experiments conducted to validate the proposed approach. Sec. 4 discusses the results and Sec. 5 concludes the paper.

^{*}This research is part of the activities of the Nordic Sound and Music Computing Network—NordicSMC (NordForsk project no. 86892)

^{© 2023} IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.



Fig. 1. (a) Supervised black-box modelling is based on *paired* audio data $\{x_i, y_i\}_{i=0}^N$, where the target audio y_i is obtained by processing the input audio x_i with the target device. When paired data is unavailable, we propose to use (b) *unpaired* data, made up of examples of a *source timbre* $\{x_i\}_{i=0}^N \in X$ and examples of a *target timbre* $\{y_j\}_{j=0}^M \in Y$, where neither the content nor the timbre contained in x_i match those contained in y_j .

2. METHOD

The proposed method uses an adversarial training approach. We define domains in the space of guitar timbre, where a domain is defined by the sounds produced by a particular combination of a guitar, pick-ups, amplifier, and any other audio effects. Note that the domain is defined purely by the guitar timbre, and not the content of the guitar playing. If we define domains, X and Y, and we sample audio data, x, where $x \in X$, and y, where $y \in Y$, our objective is to learn a mapping, G, that converts audio from domain X to domain Y:

$$G(x) = \hat{y}.\tag{1}$$

The ultimate objective is to transform the tone such that \hat{y} is perceptually indistinguishable from y. To train G we use a discriminator model D trained to identify examples from the target timbre domain.

Previous work on adversarial domain-to-domain translation [18, 20] has applied a cycle-consistency criterion to avoid a specific form of mode collapse, which results in G simply ignoring the content of the input x and synthesising unrelated content from the target domain. However, in our present experiments the generator model G did not exhibit this problem, likely due to the constrained expressive capability of the feedforward convolutional neural network used, which learns to directly apply a transformation to the input signal in the time domain.

2.1. Generator

The generator model G used in this work is a feedforward variant of the WaveNet architecture [26]. A non-causal feedforward WaveNet variant was first proposed for speech denoising [27]. A causal version was later applied to guitar amplifier modelling [11], and this architecture is used as the generator throughout this work. The model has two main components, a stack of dilated causal convolutional layers, and a linear post-processor. The post-processor is a fully connected layer that takes the outputs of each of the convolutional layers as input.

All generator models used in this work consist of two stacks of nine dilated convolutions, with each stack starting at a dilation of one and increasing by a factor of two with each subsequent convolutional layer. Each convolutional layer has a kernel size of three, and uses

Table 1. Convolutional layer parameters for proposed spectral domain discriminator.

Layer #	1	2	3	4	5	6	7
Kernel Size	10	21	21	21	21	5	3
Out Channels	32	128	512	1024	1024	1024	1
Groups	1	8	32	64	64	1	1

the same gated activation function as the original WaveNet [26]. The receptive field of this model is 2045 samples, or about 46.4 ms at a 44.1-kHz sample rate. It should also be noted that previous work has shown that a C++ implementation of this model is capable of running in real-time on a modern desktop computer [28].

2.2. Discriminator

The input to the proposed discriminator D is a time-frequency representation of the audio, as proposed in [29]. The discriminator consists of a stack of 1D convolutional layers, and the frequency bins of the time-frequency representation are provided as channels to the first layer of the discriminator. Subsequent layers use grouped convolutions to reduce computational cost. The hyperparameters for each layer are shown in Table 1. All layers use weight normalization, and all layers except the final output layer are followed by a Leaky ReLU activation function with negative slope of 0.2. Four different time-frequency representations were trialled, either a magnitude spectrogram, a magnitude mel-spectrogram. For all melspectrograms, 160 mel bands were used, the maximum frequency was set to Nyquist and the minimum frequency was set to 0 Hz.

Additionally, a multi-scale version of the spectral domain discriminator was trialled which included three sub-discriminators, each operating on time-frequency representations obtained using different window sizes. In the case of the single spectral-domain discriminator, a window size of N = 1024 was used, as for the multiscale spectral discriminator, window sizes of [512, 1024, 2048] were used. In all cases, the hop size was set to N/4.

2.3. Training Objective

The generator and discriminator of the GAN were trained using the hinge loss objective [30], identical to that which was used in Mel-GAN [31], as follows:

$$\mathcal{L}(D) = \mathbb{E}_{y}[\max(0, 1 - D(y))] + \mathbb{E}_{x}[\max(0, 1 + D(G(x)))]$$
(2)

and

$$\mathcal{L}(G) = \mathbb{E}_x[-D(G(x))],\tag{3}$$

respectively, where x is a guitar audio waveform used as input to the generator and y is a guitar audio waveform taken from the target set. The training scheme is shown in Fig. 2. During training, both the input and target guitar datasets are split into two-second segments before processing by G or D. The models were trained using the Adam optimizer [32] with a batch size of 5.

2.4. Data

The data used throughout this work was taken from a guitar dataset originally proposed for the task of automatic transcription [33]. We use the fourth subset of the dataset, which consists of 64 short musical pieces played at fast and slow tempi. The pieces were recorded on two different electric guitars, a Career SG and an Ibanez 2820.



Fig. 2. Training setup for (a) the Generator and (b) Discriminator. Generator inputs, taken from the input domain X, are processed to emulate the timbre, but not content, of the target domain Y.

We processed the recordings to remove leading and trailing silence, as well as the 2 bar count in at the beginning of each piece. Additionally, it was noted that clipping is present in some of the samples in the Career SG dataset, so examples where excessive clipping was observed were removed.

After the pre-processing, there was approximately 40 min of audio from the Ibanez 2820 guitar, and 30 min from the Career-SG. To create the datasets used during our experiments, the guitar audio was processed by a guitar amplifier plugin. To test the robustness of our modelling approach, a separate dataset was created for three different plugin settings, hereafter referred to as 'Clean', 'Light Distortion', and 'Heavy Distortion'. The amount of harmonic distortion introduced increases from a relatively small amount in the 'Clean' setting, to an extremely distorted guitar tone found in the 'Heavy Distortion' case.

3. EXPERIMENTS

Our proposed modelling approach was tested on two different problems. In the first scenario the signal contained in both the input and target datasets is recorded from the same guitar. In this case, the specific instrument in each dataset is identical, and the modelling task is to recreate the effects processing applied to the input signal. Although not realistic, this scenario is relevant because it allows the use of supervised metrics to evaluate our unsupervised training method.

In the second (more realistic) scenario, the input dataset is the unprocessed audio recorded from one guitar, and the target dataset is audio that has been recorded from a different guitar, after audio effects processing has been applied to it. In this case, the modelling task implicitly includes transforming the tone of one guitar into another, as well as recreating the effects processing applied to that guitar. The two experiments are depicted in Fig. 3.

3.1. Experiment 1: Single Guitar

The dataset was split to ensure that the input and target datasets do not contain any of the same guitar content, by dividing it into twosecond segments, with subsequent segments being sent alternately to either the input or target training dataset. This ensures that the spectral content of the unprocessed guitar in both datasets is similar, but that the actual content is different.

Unsupervised models trained with the spectral domain discriminator configurations introduced in Sec. 2 were tested, as well as an unsupervised baseline that used the MelGAN discriminator [31]. As a supervised baseline, the same generator model was trained in a supervised manner, to minimise the Error-to-Signal ratio loss function, \mathcal{E}_{ESR} , with high-pass filter pre-emphasis [11].

As in this case the guitar is the same in both datasets, as shown in Fig. 3(a), a ground truth, or reference, for how the input guitar should sound after the effects processing is applied is available. This allows a validation loss to be calculated over a held-out validation set, that consists of paired input/output guitar audio.



Fig. 3. (a) In the *Single Guitar* experiment, the input and output audio are produced using the same guitar, but with unpaired data, whereas (b) in the *Mismatched Guitar* experiment, the input and output audio are generated with different guitars.

For validation loss metrics, we use both the linear and log scaled multi-scale magnitude spectrogram loss described in [24], which we refer to as \mathcal{E}_{ms} and \mathcal{E}_{lms} respectively. In addition to this, we also present the L1 distance between the output and target mel magnitude spectrograms, again using both linear and log scaling, which we will refer to as \mathcal{E}_{mel} and \mathcal{E}_{lmel} respectively.

Additionally, a MUSHRA [34] style listening test was carried out. Participants were presented with audio clips that were processed by the target plugin, as well as the various neural network models. An anchor was also included, which was created by processing the input with a tanh nonlinearity, as well as a hidden reference. Participants were asked to rate each test condition out of 100, based on perceived similarity to the reference.

Results are shown in Table 2. In each case the training was run for 400k iterations, and the validation loss was used to select the best performing model. For the spectral domain discriminators, the validation loss used to determine the best performing model was chosen depending on the form of the input provided to the model, for example, if the input to the discriminator was a log-mel spectrogram, then the training iteration where the lowest validation \mathcal{E}_{lmel} was achieved was selected as the best performing model. In each case, our experiments included both multi-scale and single-scale spectral domain discriminators, however, for brevity, only the results for the best performing of the two are included in Table 2.

3.2. Experiment 2: Mismatched Guitar

For the scenario in Experiment 2, the guitar used to create each dataset is different, as shown in Fig. 3(b). This means that objective metrics are unavailable. As such, we conducted a listening test, in which participants were presented with a reference, consisting of a few seconds of guitar playing from the target timbre domain. The participants were then asked to rate a number of test conditions, which consisted of the next few seconds of the same piece of music, but performed on the guitar from the input timbre domain. The test conditions all consisted of processed versions of the same guitar audio. It was impossible to include a hidden reference in the test, as it does not exist.

Two baselines were created, both having access to some ground truth information. The first baseline was created by processing the input guitar with the same effects plugin that was used to create the target guitar timbre, this baseline is referred to as the "plugin-only" timbre. This corresponds to a simplified solution in which the guitarto-guitar transformation is not included, but the signal chain applied to the guitar is identical to that used to produce the reference.

The second baseline, referred to as the "EQ+plugin", was created by applying a linear equalization (EQ) matching to the unprocessed input guitar tone, with the EQ target being the target guitar

Table 2. Objective and subjective results for the Single Guitar experiment. For validations losses, bold indicates best performing unsupervised model. For the listening test result bold indicates best performing of all models and 95% confidence intervals are shown.

		Validation Loss				Listening					
Model		$\mathcal{E}_{\mathrm{ms}}$	\mathcal{E}_{lms}	\mathcal{E}_{mel}	\mathcal{E}_{lmel}	$\mathcal{E}_{\mathrm{ESR}}$	Test				
Target Tone: Clean											
Supervised		5.12	0.76	0.57	0.12	0.003	81±4.1				
MelGAN		37.5	1.47	2.75	0.17	2.38	71±4.8				
Spectral I	Domain										
Input	# Disc.										
Spect.	1	39.2	3.27	3.39	0.39	2.55	32±4.7				
Mel	1	40.0	1.51	2.88	0.28	1.27	46±4.4				
Log Spect.	3	44.1	0.81	3.76	0.18	2.71	82±4.5				
Log Mel	3	46.9	0.93	4.07	0.19	1.04	83±3.9				
Target Tone: Light Distortion											
Supervised		2.57	0.81	0.28	0.09	0.001	93±3.0				
MelGAN		25.2	2.18	1.32	0.18	2.51	73±5.4				
Spectral Domain											
Input	# Disc.										
Spect.	1	32.5	4.26	2.39	0.45	1.49	35±4.0				
Mel	1	34.4	4.12	2.57	0.48	2.43	34±4.0				
Log Spect.	1	45.3	1.11	4.51	0.23	2.18	81±4.8				
Log Mel	3	38.1	1.17	3.36	0.21	2.50	88.7±3.9				
Target Tone: Heavy Distortion											
Supervised		6.33	2.53	0.60	0.19	0.03	57±4.6				
MelGAN		22.4	2.49	1.81	0.22	2.04	92±2.8				
Spectral Domain											
Input	# Disc.										
Spect.	1	28.9	4.14	2.70	0.37	2.33	54±5.7				
Mel	1	25.5	7.15	2.36	0.60	0.86	28±3.4				
Log Spect.	1	32.1	2.52	3.25	0.29	3.17	81±4.8				
Log Mel	3	24.5	2.55	2.21	0.23	2.37	85±3.8				

before effects processing was applied. This EQ-matched version of the input guitar was then processed by the effects plugin used to create the reference timbre. Notice that this processing is impossible to achieve in a practical setting, but is used here in lieu of an ideal reference. A low-quality anchor was also included in the listening test, which consisted of the input guitar processed by a tanh nonlinearity.

The test conditions consisted of three unsupervised models, selected based on the results of the first experiment. These were trained using the MelGAN discriminator, or the spectral domain discriminator with either log spectrogram or log mel-spectogram input. In both cases, the multi-scale version of the spectral domain discriminator was used. The models were trained for 250k iterations. The results of the listening test are shown in Fig. 4.

The listening tests for both experiments were completed by twelve participants. Three participants identified as female and nine as male. The participants all had experience completing listening tests, and their mean age was 30.8 years. For the first listening test, four participants were removed in post screening as they rated the hidden reference less than 90 in more than 15% of the trials.

4. DISCUSSION

The objective results for Experiment 1, shown in Table 2, indicate that the feedforward WaveNet model trained in a supervised fashion performs better than the unsupervised models on all the proposed metrics. Generally, of the unsupervised models, those trained with the MelGAN discriminator tend to perform better in the objective loss metrics. However, the results from the listening tests for Experiment 1, also shown in Table 2, indicate that there is no clear best



Fig. 4. MUSHRA scores with 95% confidence intervals for the (a) Clean, (b) Light Distortion and (c) Heavy Distortion guitar tone settings that were modelled in Experiment 2.

performing model between the supervised and unsupervised training approach. The results do clearly show, however, that models trained using a spectral domain discriminator with linear scaled spectrogram as input perform poorly. The supervised model of the 'Heavy Distortion' tone also reveals a misalignment between the time-domain ESR loss metric and perception. For all target tones, at least one unsupervised model achieved a score of 80 or higher, indicating a perceptual match between Good and Excellent on the MUSHRA scale.

The results of the listening test for Experiment 2, shown in Fig. 4, indicate that the unsupervised models are competitive with our proposed baselines. For the 'Clean' and 'Light Distortion' case, the MelGAN model performs poorly. One possible explanation for this is that during training for the first experiment it was observed that the MelGAN produced some oscillation and instability as training went on, as the spectral discriminator models tended to quickly plateau and then remain stable. As no validation loss was available to monitor the training for the mismatched guitar case, it was not possible to select the best performing model after training.

The experimental results show that the proposed framework is capable of producing perceptually convincing models of nonlinear effects processing using unpaired data. The limitations of the proposed method are that for each desired input target pair, a new generator model must be trained. The amount of training time required for a model to achieve a perceptually convincing result was not investigated extensively, but typically took between 4 and 10 hours of training on a GPU. Further work is also required to test the method when applied to real-world use cases, for example when a smaller amount of data from the target guitar timbre is available.

5. CONCLUSION

This work shows for the first time how the guitar timbre heard on a music recording can be imitated by another guitar, using an unsupervised method based on a GAN framework. We formulated the problem as domain transfer, and proposed a spectral domain discriminator. We validated our method through two listening tests and showed that the models produced are perceptually convincing. Audio samples are available at our demonstration page¹.

 $[\]mathbf{1}_{\texttt{https://ljuvela.github.io/adversarial-amp-modeling-demo}$

6. REFERENCES

- V. Välimäki, S. Bilbao, J. O. Smith, J. S. Abel, J. Pakarinen, and D. Berners, "Virtual analog effects," in *DAFX: Digital Audio Effects*, U. Zölzer, Ed., pp. 473–522. Wiley, Chichester, UK, second edition, 2011.
- [2] M. Karjalainen, H. Penttinen, and V. Välimäki, "Acoustic sound from the electric guitar using DSP techniques," in *Proc. IEEE ICASSP*, Istanbul, Turkey, June 2000, pp. 773–776.
- [3] M. Karjalainen and J. Pakarinen, "Wave digital simulation of a vacuum-tube amplifier," in *Proc. IEEE ICASSP*, Toulouse, France, May 2006, pp. 153–156.
- [4] D. T. Yeh, "Automated physical modeling of nonlinear audio circuits for real-time audio effects—Part II: BJT and vacuum tube examples," *IEEE Trans. Speech Audio Process.*, vol. 20, no. 4, pp. 1207–1216, 2012.
- [5] O. Kröning, K. Dempwolf, and U. Zölzer, "Analysis and simulation of an analog guitar compressor," in *Proc. Int. Conf. Digital Audio Effects (DAFx)*, 2011, pp. 205–208.
- [6] A. Novak, L. Simon, P. Lotton, and J. Gilbert, "Chebyshev model and synchronized swept sine method in nonlinear audio effect modeling," in *Proc. Int. Conf. Digital Audio Effects* (*DAFx*), Graz, Austria, Sep. 2010, pp. 423–426.
- [7] F. Eichas and U. Zölzer, "Black-box modeling of distortion circuits with block-oriented models," in *Proc. Int. Conf. Digital Audio Effects (DAFx)*, Brno, Czech Republic, Sep. 2016, pp. 39–45.
- [8] S. Orcioni, A. Terenzi, S. Cecchi, F. Piazza, and A. Carini, "Identification of Volterra models of tube audio devices using multiple-variance method," *J. Audio Eng. Soc.*, vol. 66, no. 10, pp. 823–838, Oct. 2018.
- [9] T. Vanhatalo, P. Legrand, M. Desainte-Catherine, et al., "A review of neural network-based emulation of guitar amplifiers," *Appl. Sci.*, vol. 12, no. 12, pp. 5894, 2022.
- [10] M. A. Martínez Ramírez and J. D. Reiss, "Modeling nonlinear audio effects with end-to-end deep neural networks," in *Proc. IEEE ICASSP*, Brighton, UK, May 2019, pp. 171–175.
- [11] E.-P. Damskägg, L. Juvela, E. Thuillier, and V. Välimäki, "Deep learning for tube amplifier emulation," in *Proc. IEEE ICASSP*, Brighton, UK, May 2019, pp. 471–475.
- [12] A. Wright, E.-P. Damskägg, and V. Välimäki, "Real-time black-box modelling with recurrent neural networks," in *Proc. Int. Conf. Digital Audio Effects (DAFx)*, Birmingham, UK, Sep. 2019, pp. 173–180.
- [13] A. Wright and V. Välimäki, "Neural modeling of phaser and flanging effects," *J. Audio Eng. Soc.*, vol. 69, no. 7/8, pp. 517– 529, Jul. 2021.
- [14] B. Kuznetsoz, J. D. Parker, and F. Esqueda, "Differentiable IIR filters for machine learning applications," in *Proc. Int. Conf. Digital Audio Effects (DAFx)*, 2020, pp. 297–303.
- [15] S. Nercessian, A. Sarroff, and K. J. Werner, "Lightweight and interpretable neural modeling of an audio distortion effect using hyperconditioned differentiable biquads," in *Proc. IEEE ICASSP*, Toronto, Canada, June 2021, pp. 890–894.
- [16] A. Wright and V. Välimäki, "Grey-box modelling of dynamic range compression," in *Proc. Int. Conf. Digital Audio Effects* (*DAFx*), Vienna, Austria, 2022, pp. 304–311.

- [17] Y. Jing, Y. Yang, Z. Feng, J. Ye, Y. Yu, and M. Song, "Neural style transfer: A review," *IEEE Trans. Vis. Comput. Graph.*, vol. 26, no. 11, pp. 3365–3385, Nov. 2020.
- [18] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, 2017, pp. 2223–2232.
- [19] E. Grinstein, N. Q. K. Duong, A. Ozerov, and P. Pérez, "Audio style transfer," in *Proc. IEEE ICASSP*, Calgary, Canada, 2018, pp. 586–590.
- [20] T. Kaneko and H. Kameoka, "CycleGAN-VC: Non-parallel voice conversion using cycle-consistent adversarial networks," in *Proc. Eur. Signal Process. Conf.*, 2018, pp. 2100–2104.
- [21] H. Sicong, L. Qiyang, A. Cem, B. Xuchan, O. Sageev, and B. G. Roger, "TimbreTron: A wavenet(cycleGAN(CQT(audio))) pipeline for musical timbre transfer," in *Proc. Int. Conf. Learning Representations (ICLR)*, New Orleans, LA, 2019.
- [22] D. K. Jain, A. Kumar, L. Cai, S. Singhal, and V. Kumar, "ATT: Attention-based timbre transfer," in *Proc. Int. Joint Conf. Neural Networks (IJCNN)*, Glasgow, UK, 2020, pp. 1–6.
- [23] M. Carney, C. Li, E. Toh, N. Zada, P. Yu, and J. Engel, "Tone transfer: In-browser interactive neural audio synthesis," in *Proc. ACM IUI Workshops*, Apr. 2021.
- [24] J. Engel, L. Hantrakul, C. Gu, and A. Roberts, "DDSP: Differentiable digital signal processing," in *Proc. Int. Conf. Learning Representations (ICLR)*, Addis Ababa, Ethiopia, 2020.
- [25] C. J. Steinmetz, N. J. Bryan, and J. D. Reiss, "Style transfer of audio effects with differentiable signal processing," *J. Audio Eng. Soc.*, vol. 70, no. 9, pp. 708–721, Sep. 2022.
- [26] A. van den Oord, S. Dieleman, H. Zen, et al., "WaveNet: A generative model for raw audio," arXiv preprint arXiv: 1609.03499, Sep. 2016.
- [27] D. Rethage, J. Pons, and X. Serra, "A Wavenet for speech denoising," in *Proc. IEEE ICASSP*, Calgary, Canada, Apr. 2018, pp. 5069–5073.
- [28] E.-P. Damskägg, L. Juvela, and V. Välimäki, "Real-time modeling of audio distortion circuits with deep learning," in *Proc. Int. Sound and Music Computing Conf. (SMC)*, Malaga, Spain, May 2019, pp. 332–339.
- [29] L. Juvela, B. Bollepalli, X. Wang, et al., "Speech waveform synthesis from MFCC sequences with generative adversarial networks," in *Proc. IEEE ICASSP*, Calgary, Canada, Apr. 2018, pp. 5679–5683.
- [30] J. H. Lim and J. C. Ye, "Geometric GAN," arXiv preprint arXiv:1705.02894, 2017.
- [31] K. Kumar, R. Kumar, T. de Boissiere, et al., "MelGAN: Generative adversarial networks for conditional waveform synthesis," Advances in Neural Inform. Process. Syst., vol. 32, 2019.
- [32] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learning Representations* (*ICLR*), San Diego, CA, 2015.
- [33] C. Kehling, J. Abeßer, C. Dittmar, and G. Schuller, "Automatic tablature transcription of electric guitar recordings by estimation of score- and instrument-related parameters," in *Proc. Int. Conf. Digital Audio Effects (DAFx)*, Erlangen, Germany, Sep. 2014, pp. 219–226.
- [34] ITU, "BS.1534: Method for the subjective assessment of intermediate quality levels of coding systems," Recommendation ITU-R BS.1534-3, Geneva, Switzerland, 2015.