

DREAMSOUND: DEEP ACTIVATION LAYER SONIFICATION*Federico Nicolás Cámara Halac*Ohio State University
1813 North High Street
Columbus, Ohio

camarahalac.1@osu.edu

*Matías Delgado*Oxford University
Woodstock Road
Oxford, OX2 6GG

matias.delgado@maths.ox.ac.uk

ABSTRACT

Deep learning (DL) in audio signal processing has received much attention in the last four years, and it is still a growing field. Some attempts were made to translate deep networks from image to audio applications, and we discuss some of the issues that arise. In the present paper, we present DreamSound, a creative adaptation of Deep Dream to sound addressed from two approaches: *input manipulation*, and *sonification design*. Our chosen model is YAMNet, a pre-trained deep network for sound classification. We present the original Deep Dream *activation maximization* function in relation to three filter-based creative variations that were implemented. We find that more interesting results are achieved with this filter-based approach, but that there is still room for experimentation.

1. INTRODUCTION

Deep learning (DL) in audio signal processing has received much attention in the last four years, and it is still a growing field.¹ The Music Information Retrieval (MIR) community incorporated DL via large-scale audio datasets that came as early as 2015 [13, 14, 15] and most MIR problems were outperformed by DL thenceforth.²

There exists some work at the intersection of DL and sonification [18, 19], and more work in musical contexts applying DL to audio synthesis. Raw audio synthesis with DL has its origins in speech synthesis, first with the deep convolutional neural network WaveNet [20], subsequently with recurrent neural networks (RNN) [21, 22]. Their heavy computational requirements were a disadvantage for both training and performance. Further optimizations with parallel computations emerged [23, 24, 25, 26], but it was not until Generative Adversarial Networks (GANs) that music benefited from raw audio synthesized with DL [27, 28, 29, 30,

31, 32, 33].³

These ground-breaking deep networks were tailored for short-length raw audio files of around one second, with a sample rate accurate enough for speech comprehension (16kHz). Longer or higher sample rate audio files were difficult to address due to high computation demands. A major breakthrough comes from the WaveNET-like autoencoder in [15], which was created as a new, data-driven synthesis technique for longer length audio files, extending GANs towards musical application. Several hybrid synthesizer models have then been trained [40], and the first GAN synthesizers appeared: GANSynth [41], trained on the musical instrument sound dataset NSynth [15], and EnvGAN [42], trained on the environmental sound generation ESC dataset [13]. In sum, the evolution of audio synthesis has been fruitful, to the point that Google’s Magenta Lab has developed tools and plugins accessible to a wider musical audience [43].

Historically, most DL development appeared first on images and then on sound, namely because of the availability of large-scale datasets, namely, ImageNet [44]. The literature reflects that one inspiration for sonic GANs came from its deep convolutional version DCGAN [45], an adversarial image generator. The problem of translating networks from images to sounds is mentioned in [46] and [1], and an interesting discussion can be read in [2]. Of interest here is Deep Dream [47], an image generating architecture using layer *activation maximization*⁴ of a deep pre-trained model for image classification [48].

In the present paper, we present DreamSound, a creative adaptation of [47] to sound using *activation maximization* as well as *timbre style transfer*. Audio examples can be found in the code repository.⁵ DreamSound advances work on DL and audio synthesis, because it is aimed at sonifying YAMNet layers [49], a novel network previously trained on sound. Further, DreamSound proposes several creative approaches in relation to layer *activation maximization*. In section 2, we present our approach as input manipulation and in section 3 as a sonification design. Both sections discuss our implementation and findings in relation to previous work, as well and the limitations we have found. We conclude in section 4 that there is much work to be done in both the adap-

¹See [1] for a complete review, [2] for a published book, as well as [3] for the first dedicated workshop, and [4] and the repositories for DL4M [5] and the Aalto courses [6] for introductory tutorials. For the use of DL for symbolic music representations, see [7, 8, 9, 10, 11, 12]

²Despite the usefulness of DL in raw audio, e.g. reducing the gap between symbolic and audio classification [16], it can be argued that not all MIR problems benefit from it [17].

³Multiple variants to the original WaveGAN have appeared in the literature. See for example: conditional GAN [34], MelGAN [35, 36], WGANSing [37], TFGAN [32], DRUMGAN [38], SyleMelGAN [39], among others. In our previous work Kwgan (<https://github.com/fdch/kwgan>), we extended WaveGAN with conditionals to tailor an artistically relevant context.

⁴Activation maximization is a process that returns the inputs that with most confidence would cause a certain output.

⁵<https://github.com/fdch/dreamsound>



tation of image to audio algorithms and in the navigable feature space of deep models.

The research in DreamSound can be understood from two approaches: *input manipulation* and *sonification design*. In turn, these approaches are described with examples of previous work and our implementation is briefly discussed.

2. PREVIOUS WORK

2.1. Deep Dream


Deep Dream is a project described in a blog post [47] and a tutorial [50]. It is geared towards image generation using deep layer *activation maximization*, on a deep network [48] that was trained for image classification with the ImageNet [44] dataset. At the heart of the algorithm, a *gradient ascent* (see 3.2) takes place between a *loss* and an input. Loss, in this context, refers to the activations of a layer given some input. The *gradient* is said to ascend because there is incremental manipulation on the input to a next iteration of the same routine. Thus, in time, the input is ‘steered’ in an additive way towards a particular *activation* or class of the model’s class space. Briot et al [2] refer to this type of increments *input manipulation*, since the manipulation occurs at the input, not the output of the architecture. In DreamSound, there is an incremental manipulation of an initially random or specific sound content towards a targeted feature space.

The output of Deep Dream can be understood as a psychedelic image generator that creates the *pareidolia* effect, i.e., the psychological phenomenon in which the mind, in response to a stimulus, perceives similar patterns where none exist [2]. Essentially, the effect itself directly depends on the activations of a layer within a deep neural network. Thus, the combination of the neural background and the psychedelic aspect gave way for the *dream* in the name, raising further questions about the nature of artificial networks and their otherness that extends the limits of this paper.


2.2. Deep Dream in Sound

There is existing work in the application of the Deep Dream effect [47] to sound. In [51], Balke and Dittmar use a magnitude spectrogram as the color channels of an image (RGB) and apply *gradient ascent* between the spectrogram and a deep network trained on images. Subsequently, the output images were resynthesized into sound with Griffin and Lim’s method.⁶ Stamenovic adapted a similar approach in the python library tensorflow [53]. Finally, an important antecedent that steers away from Deep Dream is Herrmann’s recent work in the visualization and sonification of neural networks [19]. In contrast, Herrmann uses a network previously trained on sound to sonify the features that activate a selected layer yielding impressive results. His research further proves that while lower layers focalize on input localities, higher ones produce more upper-level sonic activity, e.g., rhythmic or harmonic patterns. In this sense, the results obtained from sonifying layer activations are informative on both input and excitations of a layer.

⁶For the Griffin and Lim’s method, see [52]



(a) Gradients (10 steps)



(a) “Speech”

Figure 1: Spectrograms of some examples. (a) Gradients output during a first after 10 steps. This was done with an initial implementation of DreamSound. (b) The Mel-scaled spectrograms displayed here belong to the 49th iteration, and are in order: (1) the output of DreamSound, (2) the difference between the original sound and the output, (3) the gradients, and (4) the inverse STFT of the phase-reconstructed hard-cut filter. Audio representations for all examples are available in the the code repository.

2.3. Images as Sound

The creative potential of these projects lives in their experimental approach towards spectral transformation. In all these cases, the authors perform layer *activation maximization* on a deep network trained on images, taking the gradient between the activations triggered by an image and the input. The input, in this case,

is the spectrogram of a sound, either magnitude [51, 53] or the scaleogram (constant-Q transform) [19]. Therefore, some issues arise due to the usage of images as sound, namely the sequential aspect of audio. For example, Purwins et al [1] note the difference between raw audio (one-dimensional time series signal) and two-dimensional images, and thus claim that audio signals have to be studied sequentially and audio-specific solutions need to be addressed. More specifically, Briot et al [2] borrow from crystallography the term *anisotropy* to describe one of the central difficulties when transcoding architectures dealing with images to sounds within spectrogram representations. *Anisotropy* means direction dependence. While the spatial correlation between pixels maintains independently of their direction of an image (*isotropy*), this is not the case when dealing with spectrogram images. In a spectrogram representation the horizontal axis represents one thing (e.g. time) and the vertical, another (e.g. frequency). Therefore, significant information is lost when translating algorithms from isotropic to anisotropic contexts. In DreamSound, we implemented an audio-specific solution using raw audio, with variations on an *activation maximization* function that performs filtering to further adjust the output.

3. DESIGN

While visualization has been the main channel to understand and perceptualize deep networks [54], there exist work on their sonification like Herman’s work mentioned above. *Feature inversion* is a technique that has been a proposed solution for understanding these architectures [55, 56], achieving important advancement in the interpretation of Machine Listening models. In [57], Mishra et al proved that temporal and harmonic structures are preserved in deep convolutional layers. Winters et al [18] advanced scientific and participant study-based work by sonifying the penultimate layer of a deep convolutional neural network previously trained for image classification. An interesting aspect of their results is the concept of a *sonification layer* that can further reduce the signal-to-noise ratio within a deep network.

Following the line of these examples, DreamSound can be understood as artistic sonification geared to output sounds using a deep convolutional neural network in an intrinsically coherent way. Three design approaches are mentioned, defined by three creative implementations of an *activation maximization* function. The examples here were made with a single-class python package implementation of DreamSound, available for import via `pip install dreamsound`⁷. The full details of the implementation can be expanded in a more verbose documentation and extend the limits of this paper.

3.1. The YAMNet Model

While there are many available trained models for sound classification, we have chosen a general sound source classifier YAMNet [49]. YAMNet is a pretrained deep neural network classifier that uses MobileNets [58], a depthwise-separable convolution architecture with a final activation layer that contains most relevant features for classification. The model was previously trained with 521 audio classes based on the Audio Set corpus [14], which are mostly non-musical, short fragments.

⁷<https://pypi.org/project/dreamsound>

3.2. Sonifying the gradients

In most machine learning models, *gradient descent* is the process by which the loss, that is, a function describing the activations of a layer, is minimized with respect to the input. The inverse happens in *gradient ascent*: the loss is maximized so that the input image increasingly excites the layers [50], arriving at *activation maximization*. The ‘gradients’ are referred to here as the gradient vector between a loss and its inputs in a model. We have found similarities in our sonifications of the gradients and in Herrmann’s sonification of the activations [19], which sound like filtered noise with dynamic spectral envelopes. Given our choice of the last layer of YAMNet, we find, like Herrmann, that the dynamic aspect of these envelopes resembles the rhythmic elements of the input sound. Visually, the spectrogram of the gradients looks like the inverted image of the spectrogram of the original sound, see Figure 1 (a).

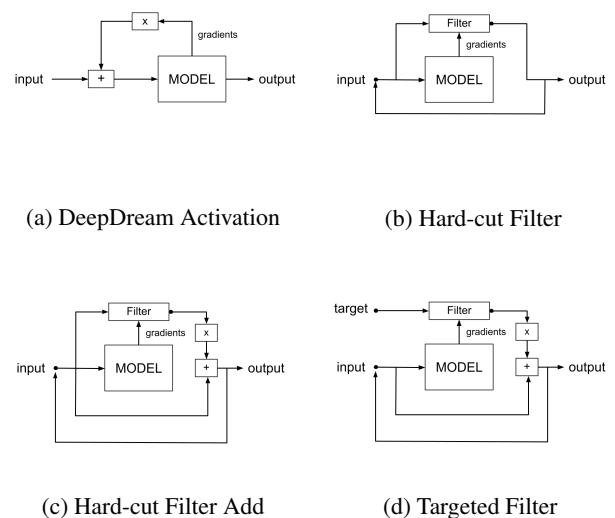


Figure 2: Activation Maximization Function types. The first type is directly adapted from DeepDream (a). From (b) to (d), a hard-cut filter is constructed from the magnitude spectra of one of the inputs, so that in (b) the gradients are filtered (convolved) by the original sound, and in (c) the output of this convolution is attenuated and added to the original sound. The difference in (d) comes with the use of another audio used as ‘target’ to construct the filters.

3.3. Deep Dream function

Following the Deep Dream function, the gradients obtained between the loss and the input are added with some attenuation to the input and fed back to the loop, as can be seen in Figure 2 (a). When maximizing the activation of a certain layer, the result is said to point to a direction in the class space of the classifier. In this case, the *activation maximization* is ‘steered’ in the *same* direction of the original sound’s class. In order to steer towards a different class of the model, we use a target. Therefore, the sound would sound more like itself or what the model considers its class to be.

However, while this function gives interesting results in the image domain, in the sound domain the additive aspect takes precedence. That is to say, the original sound and the gradients (Figure 1 b.3) are perceived as two superimposed independent layers. We recommend listening to the audio files for a better idea.

3.4. Filter-based function

Listing 1: Filter-based function with inputs x and y

```
X_mag, X pha = MAGPHASE(STFT(x))
Y_mag = ABS(STFT(y))
Y_mag_offset = NORMALIZE(Y_mag) - threshold
# hard cut based on sign
hard_cut = (SIGN(Y_mag_offset)+1) / 2
# apply the hard-cut filter
X_mag_cut = hard_cut * X_mag
# rephase
X_mag_rephased = COMPLEX_MUL(X pha, X_mag_cut)
# inverse stft
x_new = ISTFT(X_mag_rephased)
# add
output = x_new * step_size + y
return output
```

Therefore, a hard-cut filter (Figure 1 b.3) was designed to remove magnitudes that fall below a defined threshold and to let the remaining magnitudes through. Thus, by constructing this filter with the gradients and applying it to the original sound, the regions of the gradients where energy is present will let the original sound pass to the following iteration. In other words, the original sound is convolved with the gradients, and the order of the inputs for the convolution can be easily flipped to construct the filter with the original sound and convolve the gradients with it 2 (b).

The use of this function yielded more interesting sonic outcomes than the Deep Dream function, as can be seen in Figure 1 (b.1). The original sound and its gradients at that step seem to merge better than with the Deep Dream function. The difference between the original sound and the output can be seen in Figure 1 (b.2), and it evidences the elements that were removed from the original sound. Better results are achieved by an extension to this design (see 3.4). The additive element of (a) occurs between the output of the filter and the original sound, which only then is fed back into the loop, shown in Figure 2 (c). The audio examples using this function seem to be more expressive.

3.5. Target-based function

In Figure 2 (d), however, the filter is constructed with a new sound that we call ‘target’. This technique is a creative interpretation of *style transfer* that was originally designed for images in [59]. Gatys et al used a deep network to obtain content information (i.e., features) from one image and style (or feature correlations) from another. The notion of style represented then the style of an artist. In the case of audio transfer, style does not refer to musical style, but to timbre. There are other examples dealing with *timbre style transfer* [60, 61, 62], which are dealt with in [2], but in general, a gradient is used on an input and a target in order to synthesize a third, hybrid sound. In [63], Verma and Smith treat audio synthesis as a *style transfer* problem, and they use back-propagation to optimize the sound to conform to filter-output. With our target-based function, the gradients of the layer are convolved to the filter constructed by the target sound and then fed back into the loop.

4. CONCLUSIONS AND FUTURE WORK

We have presented a prototype adaptation of the Deep Dream project into sound using Tensorflow 2 and YAMNet. Our project began as a rapid and remote collaborative work within the Google Colab environment, and it developed into a python package. We have contributed some research at the intersection of deep learning and audio, and presented our results. The most interesting results occur with a filter-based *activation maximization* function built from the gradients obtained from an input and its loss. Future work consists of (1) adapting DreamSound to intake different length YAMNet layers, as well as other models. (2) Further explorations with an adversarial model such as Kwgan is expected.⁸ An interesting use of DreamSound is to further evaluate the classification capacity of YAMNet. Therefore, we intend to (3) use it as a data augmentation tool for this and other models. Finally, since DreamSound is a sound generator, we intend to (4) play it as a musical instrument in an artistic work.

5. REFERENCES

- [1] H. Purwins, B. Li, T. Virtanen, J. Schlüter, S.-Y. Chang, and T. Sainath, “Deep learning for audio signal processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 206–219, 2019.
- [2] J. Briot, G. Hadjeres, and F. Pachet, “Deep learning techniques for music generation - A survey,” *CoRR*, vol. abs/1709.01620, 2017. [Online]. Available: <http://arxiv.org/abs/1709.01620>
- [3] D. Herremans and C.-H. Chuan, “Proceedings of the first international workshop on deep learning and music,” 2017.
- [4] K. Choi, G. Fazekas, K. Cho, and M. Sandler, “A tutorial on deep learning for music information retrieval,” 2017.
- [5] “Deep learning for music (dl4m),” <https://github.com/ybayle/awesome-deep-learning-music>, 2017, accessed: 2021-06.
- [6] “Deep learning with audio,” <https://github.com/SopiMlab/DeepLearningWithAudio>, 2018, accessed: 2021-06.
- [7] L.-C. Yang, S.-Y. Chou, and Y.-H. Yang, “Midinet: A convolutional generative adversarial network for symbolic-domain music generation,” 2017.
- [8] J. Briot, A. C. da Silva Chear, R. Manzelli, V. Thakkar, A. Siahkamari, Á. Bruto, M. Simbólicos, ruído e desestruturado, and M. de Áudio Bruto, “An end to end model for automatic music generation: Combining deep raw and symbolic audio networks,” 2018.
- [9] R. Manzelli, V. Thakkar, A. Siahkamari, and B. Kulis, “Conditioning Deep Generative Raw Audio Models for Structured Automatic Music,” in *Proceedings of the 19th International Society for Music Information Retrieval Conference*. Paris, France: ISMIR, Sept. 2018, pp. 182–189. [Online]. Available: <https://doi.org/10.5281/zenodo.1492375>
- [10] H.-W. Dong and Y.-H. Yang, “Convolutional Generative Adversarial Networks with Binary Neurons for Polyphonic Music Generation,” in *Proceedings of the 19th International*

⁸<https://github.com/fdch/kwgan>

- Society for Music Information Retrieval Conference*. Paris, France: ISMIR, Sept. 2018, pp. 190–196. [Online]. Available: <https://doi.org/10.5281/zenodo.1492377>
- [11] G. Mittal, J. Engel, C. Hawthorne, and I. Simon, “Symbolic music generation with diffusion models,” 2021.
- [12] “Musegan papers,” <https://salu133445.github.io/musegan/papers>, accessed: 2021-06.
- [13] K. J. Piczak, “Esc: Dataset for environmental sound classification,” in *Proceedings of the 23rd ACM International Conference on Multimedia*, ser. MM ’15. New York, NY, USA: Association for Computing Machinery, 2015, p. 1015–1018. [Online]. Available: <https://doi.org/10.1145/2733373.2806390>
- [14] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *Proc. IEEE ICASSP 2017*, New Orleans, LA, 2017.
- [15] J. Engel, C. Resnick, A. Roberts, S. Dieleman, D. Eck, K. Simonyan, and M. Norouzi, “Neural audio synthesis of musical notes with wavenet autoencoders,” 2017.
- [16] S. Oramas, O. Nieto, F. Barbieri, and X. Serra, “Multi-Label Music Genre Classification from Audio, Text and Images Using Deep Features.” in *Proceedings of the 18th International Society for Music Information Retrieval Conference*. Suzhou, China: ISMIR, Oct. 2017, pp. 23–30. [Online]. Available: <https://doi.org/10.5281/zenodo.1417427>
- [17] H. Verma and J. Thickstun, “Convolutional composer classification,” in *Proceedings of the 20th International Society for Music Information Retrieval Conference*. Delft, The Netherlands: ISMIR, Nov. 2019, pp. 549–556. [Online]. Available: <https://doi.org/10.5281/zenodo.3527866>
- [18] W. R. Michael, A. Kalra, and B. N. Walker, “Hearing artificial intelligence: Sonification guidelines and results from a case-study in melanoma diagnosis,” 2019. [Online]. Available: <http://hdl.handle.net/1853/61501>
- [19] V. Herrmann, “Visualizing and sonifying how an artificial ear hears music,” in *Proceedings of the NeurIPS 2019 Competition and Demonstration Track*, ser. Proceedings of Machine Learning Research, H. J. Escalante and R. Hadsell, Eds., vol. 123. PMLR, 08–14 Dec 2020, pp. 192–202. [Online]. Available: <http://proceedings.mlr.press/v123/herrmann20a.html>
- [20] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” 2016. [Online]. Available: <https://deepmind.com/blog/article/wavenet-generative-model-raw-audio>
- [21] S. Mehri, K. Kumar, I. Gulrajani, R. Kumar, S. Jain, J. Sotelo, A. Courville, and Y. Bengio, “SAMPLERNN: An unconditional end-to-end neural audio generation model,” 2017.
- [22] N. Kalchbrenner, E. Elsen, K. Simonyan, S. Noury, N. Casagrande, E. Lockhart, F. Stimberg, A. van den Oord, S. Dieleman, and K. Kavukcuoglu, “Efficient neural audio synthesis,” 2018.
- [23] A. van den Oord, Y. Li, I. Babuschkin, K. Simonyan, O. Vinyals, K. Kavukcuoglu, G. van den Driessche, E. Lockhart, L. C. Cobo, F. Stimberg, N. Casagrande, D. Grewe, S. Noury, S. Dieleman, E. Elsen, N. Kalchbrenner, H. Zen, A. Graves, H. King, T. Walters, D. Belov, and D. Hassabis, “Parallel wavenet: Fast high-fidelity speech synthesis,” 2017.
- [24] L. Hantrakul, J. Engel, A. Roberts, C. Gu, and L. Hantrakul, “Fast and flexible neural audio synthesis,” in *Proceedings of the 20th International Society for Music Information Retrieval Conference*. Delft, The Netherlands: ISMIR, Nov. 2019, pp. 524–530. [Online]. Available: <https://doi.org/10.5281/zenodo.3527860>
- [25] R. Yamamoto, E. Song, and J.-M. Kim, “Parallel wavegan: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram,” 2020.
- [26] E. Song, R. Yamamoto, M.-J. Hwang, J.-S. Kim, O. Kwon, and J.-M. Kim, “Improved parallel wavegan vocoder with perceptually weighted spectrogram loss,” 2021.
- [27] B. Bollepalli, L. Juvela, and P. Alku, “Generative adversarial network-based glottal waveform model for statistical parametric speech synthesis,” *Interspeech 2017*, Aug 2017. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2017-1288>
- [28] T. Kaneko, H. Kameoka, N. Hojo, Y. Ijima, K. Hiramatsu, and K. Kashino, “Generative adversarial network-based postfilter for statistical parametric speech synthesis,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 4910–4914.
- [29] S. Pascual, A. Bonafonte, and J. Serra, “Segan: Speech enhancement generative adversarial network,” 2017.
- [30] C. Donahue, J. McAuley, and M. Puckette, “Adversarial audio synthesis,” 2018.
- [31] R. Prenger, R. Valle, and B. Catanzaro, “Waveglow: A flow-based generative network for speech synthesis,” in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 3617–3621.
- [32] Q. Tian, Y. Chen, Z. Zhang, H. Lu, L. Chen, L. Xie, and S. Liu, “TfGAN: Time and frequency domain based generative adversarial network for high-fidelity speech synthesis,” 2020.
- [33] J.-Y. Liu, Y.-H. Chen, Y.-C. Yeh, and Y.-H. Yang, “Unconditional audio generation with generative adversarial networks and cycle regularization,” *Interspeech 2020*, Oct 2020. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2020-1137>
- [34] C. Y. Lee, A. Toffy, G. J. Jung, and W. Han, “Conditional wavegan,” *CoRR*, vol. abs/1809.10636, 2018. [Online]. Available: <http://arxiv.org/abs/1809.10636>
- [35] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. C. Courville, “MelGAN: Generative adversarial networks for conditional waveform synthesis,” in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019. [Online]. Available: <https://proceedings.neurips.cc/paper/2019/file/6804c9bca0a615bdb9374d00a9fcba59-Paper.pdf>

- [36] W. Jang, D. Lim, and J. Yoon, “Universal melgan: A robust neural vocoder for high-fidelity waveform generation in multiple domains,” 2021.
- [37] P. Chandna, M. Blaauw, J. Bonada, and E. Gomez, “Wgansing: A multi-voice singing voice synthesizer based on the wasserstein-gan,” *2019 27th European Signal Processing Conference (EUSIPCO)*, Sep 2019. [Online]. Available: <http://dx.doi.org/10.23919/EUSIPCO.2019.8903099>
- [38] J. Nistal, S. Lattner, and G. Richard, “DrumGAN: Synthesis of drum sounds with timbral feature conditioning using generative adversarial networks,” in *Proceedings of the 21st International Society for Music Information Retrieval Conference*. Montreal, Canada: ISMIR, Oct. 2020, pp. 590–597. [Online]. Available: <https://doi.org/10.5281/zenodo.4245504>
- [39] A. Mustafa, N. Pia, and G. Fuchs, “Stylemelgan: An efficient high-fidelity adversarial vocoder with temporal adaptive normalization,” 2021.
- [40] O. McCarthy and Z. Ahmed, “Hooligan: Robust, high quality neural vocoding,” 2020.
- [41] J. Engel, K. K. Agrawal, S. Chen, I. Gulrajani, C. Donahue, and A. Roberts, “Gansynth: Adversarial neural audio synthesis,” 2019.
- [42] A. Madhu and S. K., “Envgan: Adversarial synthesis of environmental sounds for data augmentation,” 2021.
- [43] A. Roberts, J. Engel, Y. Mann, J. Gillick, C. Kayacik, S. Nørly, M. Dinculescu, C. Radebaugh, C. Hawthorne, and D. Eck, “Magenta Studio: Augmenting Creativity with Deep Learning in Ableton Live,” in *Proceedings of the 6th International Workshop on Musical Metacreation*. Charlotte, United States: MUME, June 2019, p. 7. [Online]. Available: <https://doi.org/10.5281/zenodo.4285266>
- [44] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, “ImageNet Large Scale Visual Recognition Challenge,” *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015.
- [45] A. Radford, L. Metz, and S. Chintala, “Unsupervised representation learning with deep convolutional generative adversarial networks,” 2015.
- [46] D. Rothman, “What’s wrong with cnns and spectrograms for audio processing?” <https://towardsdatascience.com/whats-wrong-with-spectrograms-and-cnns-for-audio-processing-311377d7cc6/>, accessed: 2021-06.
- [47] “Deepdream,” <https://ai.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>, 2015, accessed: 2021-06.
- [48] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” 2014.
- [49] M. Plakal and D. Ellis, “Yamnet,” <https://github.com/tensorflow/models/tree/master/research/audioset/yamnet>, 2020, accessed: 2021-06.
- [50] “Deepdream,” <https://www.tensorflow.org/tutorials/generative/deepdream>, accessed: 2021-06.
- [51] C. Dittmar and S. Balke, “Deepdreameffect,” <https://github.com/stefan-balke/hamr2015-deepdreameffect>, 2015, accessed: 2021-06.
- [52] D. Griffin and J. Lim, “Signal estimation from modified short-time fourier transform,” in *ICASSP ’83. IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 8, 1983, pp. 804–807.
- [53] M. Stamenovic, “Deep dreaming on audio spectrograms with tensorflow,” <https://github.com/markostam/audio-deepdream-tf>, 2016, accessed: 2021-06.
- [54] K. Simonyan, A. Vedaldi, and A. Zisserman, “Deep inside convolutional networks: Visualising image classification models and saliency maps,” 2014.
- [55] A. Mahendran and A. Vedaldi, “Understanding deep image representations by inverting them,” 2014.
- [56] A. Dosovitskiy and T. Brox, “Inverting visual representations with convolutional networks,” 2016.
- [57] S. Mishra, B. L. Sturm, and S. Dixon, “Understanding a Deep Machine Listening Model Through Feature Inversion,” in *Proceedings of the 19th International Society for Music Information Retrieval Conference*. Paris, France: ISMIR, Sept. 2018, pp. 755–762. [Online]. Available: <https://doi.org/10.5281/zenodo.1492527>
- [58] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, “Mobilenets: Efficient convolutional neural networks for mobile vision applications,” 2017.
- [59] L. A. Gatys, A. S. Ecker, and M. Bethge, “A neural algorithm of artistic style,” *CoRR*, vol. abs/1508.06576, 2015. [Online]. Available: <http://arxiv.org/abs/1508.06576>
- [60] D. Foote, D. Yang, and M. Rohaninejad, “Do androids dream of electric beats?” <https://audiostyletransfer.wordpress.com>, 2016, accessed: 2021-06.
- [61] D. Ulyanov and V. Lebedev, “Audio texture synthesis and style transfer,” <https://dmitryulyanov.github.io/audio-texture-synthesis-and-style-transfer>, 2016, accessed: 2021-06.
- [62] L. Wyse, “Audio spectrogram representations for processing with convolutional neural networks,” in *Proceedings of the 1st International Workshop on Deep Learning for Music*, May 2017, p. 37–41.
- [63] P. Verma and J. O. Smith, “Neural style transfer for audio spectrograms,” 2018.