
Articles

2021-12-01

A sinusoidal signal reconstruction method for the inversion of the mel-spectrogram

Anastasia Natsiou

Technological University Dublin, anastasia.natsiou@tudublin.ie

Sean O'Leary

Technological University Dublin, sean.oleary@tudublin.ie

Follow this and additional works at: <https://arrow.tudublin.ie/creaart>



Part of the [Engineering Commons](#)

Recommended Citation

Natsiou, A., & O'Leary, S. (2022). A sinusoidal signal reconstruction method for the inversion of the mel-spectrogram. Technological University Dublin. DOI: 10.21427/R27X-F169

This Article is brought to you for free and open access by ARROW@TU Dublin. It has been accepted for inclusion in Articles by an authorized administrator of ARROW@TU Dublin. For more information, please contact arrow.admin@tudublin.ie, aisling.coyne@tudublin.ie, gerard.connolly@tudublin.ie.



This work is licensed under a [Creative Commons Attribution-NonCommercial-Share Alike 4.0 License](#)
Funder: SFI

A sinusoidal signal reconstruction method for the inversion of the mel-spectrogram

Anastasia Natsiou
Technological University of Dublin
Dublin, Ireland
anastasia.natsiou@tudublin.ie

Seán O’Leary
Technological University of Dublin
Dublin, Ireland
sean.oleary@tudublin.ie

Abstract—The synthesis of sound via deep learning methods has recently received much attention. Some problems for deep learning approaches to sound synthesis relate to the amount of data needed to specify an audio signal and the necessity of preserving both the long and short time coherence of the synthesised signal. Visual time-frequency representations such as the log-mel-spectrogram have gained in popularity. The log-mel-spectrogram is a perceptually informed representation of audio that greatly compresses the amount of information required for the description of the sound. However, because of this compression, this representation is not directly invertible. Both signal processing and machine learning techniques have previously been applied to the inversion of the log-mel-spectrogram but they both caused audible distortions in the synthesised sounds due to issues of temporal and spectral coherence. In this paper, we outline the application of a sinusoidal model to the ‘inversion’ of the log-mel-spectrogram for pitched musical instrument sounds outperforming state-of-the-art deep learning methods. The approach could be later used as a general decoding step from spectral to time intervals in neural applications.

Index Terms—Sound reconstruction, mel-spectrogram, sinusoidal model, machine learning.

I. INTRODUCTION

Digitised audio signals can be represented in different forms depending on the application and the available resources. *Raw audio* is a representation of the time domain acoustic signal specified by the sampling rate and bit depth. This form is typically the end goal of sound synthesis algorithms. However, as we perceive sound in terms of time evolving spectral content, time-frequency representations are often appropriate. Such representations include the *spectrogram* based on *short time Fourier transform* (STFT), and the *log-mel-spectrogram*, which is a compressed form of spectrogram modelled after the frequency resolution of the human ear.

During the last years, many different approaches have been applied to sound synthesis. Basic signal processing methods, e.g. *concatenative synthesis* [1] and *statistical parametric synthesis* [2] were very influential in the field, especially in the area of speech synthesis. More recently, the rise of computational power and deep learning architectures have led to the development of high quality results for sound synthesis. WaveNet [3] is a state-of-the-art vocoder¹ based on

dilated convolutional neural networks. WaveNet’s high quality depends on synthesis directly in the time domain; every new sample is based on previous ones. Therefore, the system is computationally expensive and it needs a large amount of data in order to be trained.

Many deep learning algorithms have employed spectrogram like representations [4] [5] [6] [7] [8]. A major issue with this representation is that it is not directly invertible as it only describes the general spectral distribution over time. In this paper we use the log-mel-spectrogram to generate parameters for a signal processing synthesis model. We implemented a modification of the classic sinusoidal model generating amplitudes that approximate the spectral distribution of the log-mel-spectrogram; the phase calculation takes advantage of the continuous nature of quasi-sinusoidal oscillations over time. Thus, our system overcomes some of the issues of deep learning based inversion, such as that used in Tacotron 2 [9], being a less computationally expensive decoding step and shows the potential of producing higher quality results for the class of sounds under consideration.

The paper is organized as follows. In Section 2 we review previous works on sound reconstruction methods, and in Section 3 we describe the sinusoidal model and our modification for the inversion of the log-mel-spectrogram. Results are presented in Section 4, and conclusions in Section 5.

II. RELATED WORK

In many visual representations of sound, the phase part of a time-frequency representation is discarded. This is generally appropriate for sustained sounds as the ear is not sensitive to absolute phase. To reconstruct a pitched waveform from its given “phaseless” spectrogram, many signal processing-based and machine learning-based methods have been proposed.

Signal processing methods for sound generation focus directly on attempting to reconstruct the phase from the spectrogram. A much used solution for phase estimation is the Griffin-Lim algorithm [10] [11]. This method manages to estimate the parameters by minimizing a least square error criterion between the generated and the original spectrum. The algorithm is able to attain impressive results and it is used by many researchers even today [12]. Another significant vocoder for generating sound from a time-frequency spectrogram is STRAIGHT [13]. The key objective of this algorithm is to

¹Vocoder comes from the term ‘voice coder’ although it is also used for more general analysis/synthesis models e.g. ‘phase vocoder’. Here we are using it in the latter more general sense.

eliminate the periodicity of a speech signal and then use a sampling operation to generate the parameters. WORLD [14] is designed to reconstruct the sound by its fundamental frequency, the spectral envelope and an aperiodic parameter. Due to its quick processing and its low computational cost, WORLD is able to generate sound with adequate quality in real time.

Most of the models lack of sound naturalness, producing a "robotic" effect. Consequently, machine learning has been adopted widely. WaveNet [3] is an autoregressive and probabilistic neural network-based architecture applied on raw audio. Although WaveNet was initially proposed as a stand-alone application, in Tacotron 2 [9], it was conditioned on log-mel-spectrogram and used as a vocoder. Another implementation of deep learning for the inversion of the log-mel-spectrogram is WaveGlow [4]. This method models multiple invertible transformations to estimate and sample from the distribution of the training data. The generated sound achieved demonstrates the same quality as WaveNet but with less computational cost.

To take advantage of the strengths of both methodologies, some contemporary attempts include hybrid models. LPCNet [15] was proposed as an improvement to the WaveRNN [16]. A Linear Prediction Coefficient analysis on Bark-scale cepstral coefficients was applied to generate the training data of a neural network. This way, they managed to model the response of the vocal tract and improve the synthesised sound.

Although our approach clearly falls into the signal processing category, our short term plan includes the investment to a more hybrid method, using the method as a general decoding step of a neural architecture; an approach conceptually similar to the recent DDSP [17]. In DDSP, a differentiable autoencoder has been applied for the parameter estimation in a supervised or unsupervised manner. The model manages to synthesise monophonic sounds using the fundamental frequency, the loudness and a latent representation which forms the timbre. The computed parameters are then fed to a harmonic oscillator and a noise filtering component.

III. SINUSOIDAL MODEL

Although deep learning models demonstrated some results, their demand for a large amount of data as well as their high computational cost lead researchers to prior signal processing or hybrid methods. The sinusoidal model [18] constitutes a signal processing method that provides a parametric approximation of a signal. In the following section, after introducing the mel-scale, we will explain the original sinusoidal model, and then we will apply the effective inversion of the log-mel-spectrogram.

A. The mel-spectrogram

In 1937, a group of psychologists conducted a psychoacoustical experiment concerning the perceived distance in pitch between sounds with varying frequency [19]. The outcome was that the human ear has greater resolution at lower frequencies. This observation led to the development of the *mel-scale* (mel

from 'melody'). The conversion of the frequency in hertz (f) to the mel-scale is illustrated in Eq.1.

$$mel = 2595 \log_{10} \left(1 + \frac{f}{700} \right) \quad (1)$$

The mel-spectrogram demonstrates a compressed form of sound in the time-frequency domain. This nonlinear transformation constitutes the outcome of the Short Time Fourier Transform (STFT) after the application of mel-filters (a bank of bandpass filters with bandwidths modelled after the mel-scale).

B. Sinusoidal Model from Linear Spectrogram

The sinusoidal model [18] models signals as a linear combination of sinusoids, each having independent time varying amplitude and phase trajectories. The modeled signal is given by

$$y(t) = \sum_{l=1}^L A_s(t) \cos(\theta_s(t)) \quad (2)$$

where $\theta_s(t) = \int_0^t 2\pi f_l(\tau) d\tau + \phi_l$ is the time evolving phase of the s_{th} sinusoid ².

The basic steps of analysis include a *peak detection* of the spectral shape in each frame of the STFT, a *parameter estimation* using parabolic interpolation, and a *peak matching* across frames of the SFTF. The synthesis is then the addition of the sinusoids modelled from the time evolving trajectories.

C. Sinusoidal Model from Mel-Spectrogram

The following method uses the log-mel-spectrogram and an estimate of the fundamental frequency to synthesise the target sound. The log-mel-spectrogram constitutes a non-invertible representation since many frequencies can co-exist in the same mel-band. Furthermore, the mel-spectrogram encodes only the energy distribution over time-frequency of the given sound, therefore, it represents a phaseless representation. In this paper, we propose a sinusoidal model to model a harmonic source, using the log-mel-spectrogram to inform the spectral distribution. While not attempting to estimate the absolute phase of the original signal, this will preserve continuity in the oscillations of the signal. The approach outlined here can be seen as a source/filter approach - shaping a spectrally flat source with a time variant filter.

1) *Frequency Estimation*: For monophonic harmonic target signals, the sinusoid partials are described as integer multiples of the fundamental frequency [20]. Therefore, the harmonic frequencies can be calculated according to the Eq. 3 where f_n^1 is the fundamental frequency of the n_{th} frame, $i=[1, \text{floor} \left[\frac{f_s/2}{f_n^1} \right]]$.

$$f_m^i = i f_m^1 \quad (3)$$

where f_m^i is the frequency of the i_{th} sinusoid in the m_{th} frame and f_s is the sampling rate.

²Equation 2 is in continuous time, what follows will be in reference to frames of the STFT

In addition, in order to guarantee no distortions of the synthesised signal based on fluctuations of the fundamental frequency, a 6% tolerance is adopted between adjacent frames.

2) *Amplitude Computation*: To estimate the amplitude of each sinusoid, the energy of every frame of the log-mel-spectrogram needs to be preserved. Conversely, having the mel-spectrogram and the information of the mel-filter banks, the amplitude of each sinusoid can be approximated using a source filter model. As this project demonstrates an initial attempt of this idea, a filter equivalent to the rectangular one is applied. Finally, because a single frequency can be present in more than one filter, the amplitude computed is the average of those estimated in overlapping filters.

3) *Phase Reconstruction*: The phase of each estimated frequency peak is calculated by the phase of the previous frame and the frequencies on both the previous and the current frame interpolated by the hop size as it is expressed in Eq. 4.

$$\theta_n^i = \theta_{n-1}^i + 2\pi T \frac{f_n^i + f_{n-1}^i}{2} \quad (4)$$

where θ_n^i is the phase of the i th sinusoid in the n th frame, T is the hop size multiplied by f_s .

IV. EXPERIMENTS

To evaluate the performance of the modified sinusoidal model, we examined the inversion of the mel-spectrogram for independent pitches of a variety of instruments. The results are compared to a pre-trained WaveNet conditioned on the log-mel-spectrogram and to a pre-trained WaveNet autoencoder [21]. The sounds used in this experiment can be found here: <https://anastasianat57.github.io/sinusoidalreconstruction/>.

A. Dataset

For the conducted experiments, we used a subsample of the NSynth dataset [22], trying to be as diverse as possible. For each instrument, three different music notes were captured; a low, a medium and a high pitch, in a range of 80Hz to 2100Hz. The sounds could belong in different categories as per their velocity or acoustic quality.

B. Description of the experiments

The sinusoidal model approach as well as WaveNet are utilised in this paper in order to reconstruct the original music notes from their the log-mel-spectrogram. As we explained before, the models are based on two completely different technologies, therefore the comparison occurs between signal processing and deep learning approaches independently for the conversion of a frequency representation to time domain. The log-mel-spectrogram was created after the raw audio of the original sound with 80 bins, originating from an STFT of 1024 window length, 256 hop size and 1024 FFT size. More details about the parameters of both models are given in the following sections.

1) *WaveNet*: The sinusoidal model is compared with an already pre-trained WaveNet on speech [23]. These results can be considered as baseline because both the models are conditioned on the same log-mel-spectrogram.

2) *WaveNet Autoencoder*: The main comparison can be considered between our approach and a WaveNet autoencoder pre-trained on the NSynth dataset [24]. Although the network is not conditioned on the log-mel-spectrogram, the original tests [21] demonstrated that the representation generated from the autoencoder can outperform a baseline using spectrograms. However, we did not consider their baseline as an interesting comparison to our method.

3) *Sinusoidal Model*: For the sinusoidal model, a normalized blackman window is used to address the expectations of the amplitude computation and the yin algorithm searches for fundamental frequencies between 80Hz and 3000Hz.

C. Evaluation

After reconstructing the sample from its original representation, the models are evaluated using spectral convergence. For this purpose, the synthesised sounds are aligned with the original music note and the error is calculated by the Eq. 5, where S represents the power spectrogram of the original, \hat{S} the power spectrogram of the synthesised and n the FFT size of the power spectrogram.

$$SC = \sqrt{\frac{\sum_{i=1}^m \sum_{j=1}^n |S(i, j) - \hat{S}(i, j)|}{n + m}} \quad (5)$$

Finally, the two aligned sounds are normalized by a total energy calculated using the energy of the original and the energy of the generated sound as it is demonstrated in the Eq. 6.

$$energy = \sqrt{\frac{\sum_{i=1}^m S(i, j)}{\sum_{i=1}^m \hat{S}(i, j)}} \quad (6)$$

D. Results

The results, aggregating them by instrument, are illustrated in Fig. 1. Using the spectral convergence, one can identify that the modified sinusoidal model is more precise than the WaveNet autoencoder in all the cases. We encourage the reader to listen to the samples provided. The musical note synthesised by the sinusoidal model sounds more natural, presenting phase coherence.

Another advantage of the modified sinusoidal model over deep learning approaches is the low computational cost. The synthesis of a four second musical note using a pre-trained WaveNet model needs significantly more time than the synthesis of the same sound using the sinusoidal model on an average laptop. In cases where the model needs to be trained again, the memory needed and the computational cost exceed the capabilities of an average laptop.

V. CONCLUSION

Deep learning architectures have been proven very effective in sound synthesis applications. WaveNet is a state-of-the-art signal reconstruction model based on dilated convolutional neural networks and it has been used for the inversion of the log-mel-spectrogram, a non-invertible audio representation. While deep learning models do show promise, they can be

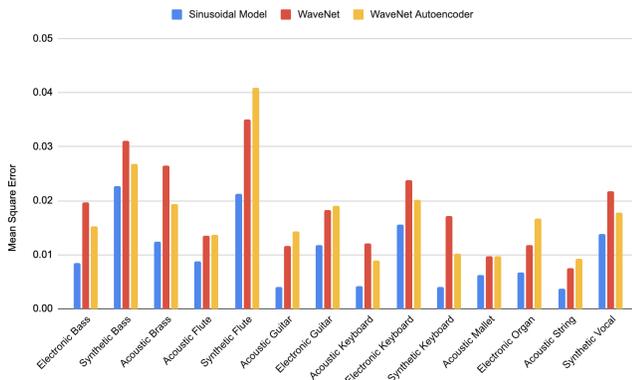


Fig. 1. Mean Square error of the power spectrogram between the original and the generated sound for both sinusoidal model and WaveNet

difficult to interpret and are computationally inefficient at the final synthesis step. Comparing it with a signal processing method, WaveNet autoencoder presents lack of continuity and phase coherence while the requirements for computational power are severe. In the current paper, we proposed a signal processing method based on a modified sinusoidal model for the inversion of the log-mel-spectrogram. The results seem promising, achieving more natural synthesised musical notes than WaveNet with a significantly lower computational cost. Although the outcome of the sinusoidal model illustrated high quality, the current version of the model is not appropriate for the generation of non harmonic signals. With the suitable extraction of the parameters, this model could be used for any sound. This research work illustrates an early investigation of marrying signal processing with deep learning. An extension of the presented model could be used as the last step of a deep learning architecture for decoding spectrogram-based representations for a more general class of single source sounds.

ACKNOWLEDGMENT

This work was funded by Science Foundation Ireland through the SFI Centre for Research Training in Machine Learning (18/CRT/6183)

REFERENCES

- [1] A. Hunt and A. Black, "Unit selection in a concatenative speech synthesis system using a large speech database," in *1996 IEEE International Conference on Acoustics, Speech, and Signal Processing Conference Proceedings*, vol. 1, (Atlanta, GA, USA), pp. 373–376, IEEE, 1996.
- [2] H. Zen, K. Tokuda, and A. W. Black, "Statistical parametric speech synthesis," *Speech Communication*, vol. 51, pp. 1039–1064, Nov. 2009.
- [3] A. v. d. 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," *arXiv:1609.03499 [cs]*, Sept. 2016. arXiv: 1609.03499.
- [4] R. Prenger, R. Valle, and B. Catanzaro, "WaveGlow: A Flow-based Generative Network for Speech Synthesis," *arXiv:1811.00002 [cs, eess, stat]*, Oct. 2018. arXiv: 1811.00002.
- [5] J. Engel, K. K. Agrawal, S. Chen, I. Gulrajani, C. Donahue, and A. Roberts, "GANSYNTH: ADVERSARIAL NEURAL AUDIO SYNTHESIS," p. 17, 2019.

- [6] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brebisson, Y. Bengio, and A. Courville, "MelGAN: Generative Adversarial Networks for Conditional Waveform Synthesis," *arXiv:1910.06711 [cs, eess]*, Dec. 2019. arXiv: 1910.06711.
- [7] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, "DiffWave: A Versatile Diffusion Model for Audio Synthesis," *arXiv:2009.09761 [cs, eess, stat]*, Mar. 2021. arXiv: 2009.09761.
- [8] S. Vasquez and M. Lewis, "MelNet: A Generative Model for Audio in the Frequency Domain," *arXiv:1906.01083 [cs, eess, stat]*, June 2019. arXiv: 1906.01083.
- [9] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. J. Skerry-Ryan, R. A. Saurous, Y. Ajiomyriannakis, and Y. Wu, "Natural TTS Synthesis by Conditioning WaveNet on Mel Spectrogram Predictions," *arXiv:1712.05884 [cs]*, Feb. 2018. arXiv: 1712.05884.
- [10] D. Griffin and Jae Lim, "Signal estimation from modified short-time Fourier transform," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 32, pp. 236–243, Apr. 1984.
- [11] D. Griffin and Jae Lim, "A new model-based speech analysis/Synthesis system," in *ICASSP '85. IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 10, (Tampa, FL, USA), pp. 513–516, Institute of Electrical and Electronics Engineers, 1985.
- [12] Y. Wang, R. J. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. Le, Y. Ajiomyriannakis, R. Clark, and R. A. Saurous, "Tacotron: Towards End-to-End Speech Synthesis," *arXiv:1703.10135 [cs]*, Apr. 2017. arXiv: 1703.10135.
- [13] H. Kawahara, I. Masuda-Katsuse, and A. de Cheveigné, "Restructuring speech representations using a pitch-adaptive time-frequency smoothing and an instantaneous-frequency-based F0 extraction: Possible role of a repetitive structure in sounds," *Speech Communication*, vol. 27, pp. 187–207, Apr. 1999.
- [14] M. Morise, F. Yokomori, and K. Ozawa, "WORLD: A Vocoder-Based High-Quality Speech Synthesis System for Real-Time Applications," *IEICE Transactions on Information and Systems*, vol. E99.D, no. 7, pp. 1877–1884, 2016.
- [15] J.-M. Valin and J. Skoglund, "LPCNET: Improving Neural Speech Synthesis through Linear Prediction," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, (Brighton, United Kingdom), pp. 5891–5895, IEEE, May 2019.
- [16] N. Kalchbrenner, E. Elsen, K. Simonyan, S. Noury, N. Casagrande, E. Lockhart, F. Stimberg, A. v. d. Oord, S. Dieleman, and K. Kavukcuoglu, "Efficient Neural Audio Synthesis," *arXiv:1802.08435 [cs, eess]*, June 2018. arXiv: 1802.08435.
- [17] J. Engel, L. Hantrakul, C. Gu, and A. Roberts, "DDSP: Differentiable Digital Signal Processing," *arXiv:2001.04643 [cs, eess, stat]*, Jan. 2020. arXiv: 2001.04643.
- [18] R. McAulay and T. Quatieri, "Speech analysis/Synthesis based on a sinusoidal representation," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 34, pp. 744–754, Aug. 1986.
- [19] J. Volkman, S. S. Stevens, and E. B. Newman, "A Scale for the Measurement of the Psychological Magnitude Pitch," p. 2.
- [20] J. W. Beauchamp, ed., *Analysis, synthesis, and perception of musical sounds: the sound of music*. Modern acoustics and signal processing, New York: Springer, 2007. OCLC: ocm80145578.
- [21] J. Engel, C. Resnick, A. Roberts, S. Dieleman, D. Eck, K. Simonyan, and M. Norouzi, "Neural Audio Synthesis of Musical Notes with WaveNet Autoencoders," *arXiv:1704.01279 [cs]*, Apr. 2017. arXiv: 1704.01279.
- [22] "The NSynth Dataset." Library Catalog: magenta.tensorflow.org.
- [23] "An open source implementation of WaveNet vocoder: https://r9y9.github.io/wavenet_vocoder."
- [24] "Wavenet autoencoder: <https://github.com/magenta/magenta>."