A context encoder for audio inpainting

Andrés Marafioti, Nathanaël Perraudin, Nicki Holighaus, and Piotr Majdak

Abstract—We study the ability of deep neural networks (DNNs) to restore missing audio content based on its context, i.e., inpaint audio gaps. We focus on a condition which has not received much attention yet: gaps in the range of tens of milliseconds. We propose a DNN structure that is provided with the signal surrounding the gap in the form of time-frequency (TF) coefficients. Two DNNs with either complex-valued TF coefficient output or magnitude TF coefficient output were studied by separately training them on inpainting two types of audio signals (music and musical instruments) having 64-ms long gaps. The magnitude DNN outperformed the complex-valued DNN in terms of signalto-noise ratios and objective difference grades. Although, for instruments, a reference inpainting obtained through linear predictive coding performed better in both metrics, it performed worse than the magnitude DNN for music. This demonstrates the potential of the magnitude DNN, in particular for inpainting signals that are more complex than single instrument sounds.

I. INTRODUCTION

Locally degraded or even lost information is encountered in various audio processing tasks. Some examples are corrupted audio files, lost information in audio transmission (referred to as packet-loss in the context of voice-over-IP transmission), and audio signals locally contaminated by noise. Restoration of lost information in audio has been referred to as audio inpainting [1], audio inter-/extrapolation [2], [3], or waveform substitution [4]. Reconstruction is usually aimed at providing a coherent and meaningful information while preventing audible artifacts so that the listener remains unaware of any occurred problem. Successful algorithms are limited to deal with a particular class of audio signals [5], or they focus on a specific duration of the problematic signal parts [6], and/or they exploit a-priori information about the problem [7].

In this work, we explore a new machine-learning algorithm with respect to the reconstruction of lost parts of audio signals, i.e., *gaps*. From all possible classes of audio signals, we limit the reconstruction to instrumental music, i.e., mix of sounds from musical instruments organized in time. We focus on gaps of *medium* durations, that is, in the range of tens of milliseconds. We assume that gaps are separated in time, such

Manuscript received on October 2018; revised on April 2019.

that the local audio information surrounding the gap, namely, the *context*, is reliable and can be exploited.

The proposed algorithm is based on an unsupervised feature-learning algorithm driven by context-based sample prediction. It relies on a DNNs with convolutional and fully connected layers (FCLs) trained to generate TF representations of sounds being conditioned on contextual TF information. We call the algorithm *context encoder*, as introduced for images [8] in analogy to auto encoders [9]. Our context encoder aims at studying the general ability of DNNs to accurately inpaint audio in the range of tens of milliseconds from limited but reliable context in order to determine factors with the largest potential for future improvement and details requiring a more sophisticated method.

A. Related deep-learning techniques

Deep learning excels in classification, regression, and anomaly detection tasks [9] and it has also shown good results in generative modeling with techniques such as variational auto encoders [10] and generative adversarial networks [11]. Unfortunately, for audio synthesis only the latter has been studied, applying it to generate snippets of sound [12]–[14]. In order to obtain meaningful results, state-of-the-art audio synthesis requires sophisticated networks [15], [16]. While these approaches directly predict audio samples based on the preceding samples, in the speech-synthesis field, synthesis of audio in domains other than time such as spectrograms [17], and mel-spectrograms [18], [19] have been proposed. In the field of speech transmission, DNNs have been used to achieve packet loss concealment [20].

The synthesis of *musical* audio signals using deep learning, however, is even more challenging [21]. A music signal is comprised of complex sequences ranging from short-term structures (any periodicity in the waveform) to long-term structures (like figures, motifs, or sections). In order to simplify the problem brought by long-range dependencies, music synthesis in multiple steps has been proposed including an intermediate symbolic representation like MIDI sequences [22], and features of a parametric vocoder [23].

While these contributions provide insights on the design of a neural network for audio synthesis, none of them addresses conditions in which some audio information has been lost, but the surrounding context is available.

B. Related audio-inpainting algorithms

The term "audio inpainting" was coined by Adler et al. to describe a large class of inverse problems in audio processing, while focussing their own study on the restoration of gaps in audio signals [1]. The general assumption for audio inpainting is that audio is represented in some domain as data and

Andrés Marafioti, Nicki Holighaus, and Piotr Majdak are with the Acoustics Research Institute, Austrian Academy of Sciences, Wohllebengasse 12–14, 1040 Vienna, Austria.

Nathanaël Perraudin is with the Swiss Data Science Center, ETH Zürich, Universitätstrasse 25, 8006 Zürich

Accompanying web page (sound examples, Matlab and Python code, color figures):

https://andimarafioti.github.io/audioContextEncoder/. We thank the reviewers and the editor for their review and their helpful suggestions. This work has been supported by Austrian Science Fund (FWF) project MERLIN (Modern methods for the restoration of lost information in digital signals;I 3067-N30). We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan X Pascal GPU used for this research.

some chunks of that data are corrupted yielding *gaps* in the representation.

The number and duration of the gaps as well as the type of corruption is manifold. For example, in declicking and declipping, corruptions may be frequent, but mostly confined to disconnected time-segments of only few milliseconds duration or less. We refer to inpainting such gaps as inpainting of *short* gaps. On the other hand, gaps on a scale of hundreds of milliseconds or even seconds may happen, e.g., when reading partially damaged physical media, in live music recordings, when unwanted noise originating from the audience needs to be removed, or in audio transmission with a total loss of the connection between transmitter and receiver lasting for seconds. We refer to inpainting such gaps as inpainting *long* gaps.

In contrast, we define *medium gaps* as those with tens of milliseconds duration, a scale on which the non-stationary characteristic of audio already becomes important, but the extrapolation of the missing information from short context surrounding the gap still seems feasible. Medium gaps may arise as a consequence of packet loss in audio transmission [5] or when short interruption happens while reading audio from partially damaged physical media. Interestingly, not much has been done for audio inpainting of medium gaps.

In contrast, for inpainting short gaps, various solutions have been proposed. [1] proposed a framework based on orthogonal matching pursuit (OMP), which has inspired a considerable amount of research exploiting TF sparsity [24]-[27] or structured sparsity [28]-[30]. Being tempted to extend these works to medium gap durations, one gets disappointed quite soon because for increasing gap durations (from the originally targeted of 10 ms to medium gap durations of around 50 ms), the reconstruction quality substantially decreases, see Fig. 1 in [27]. The degradation originates in the combination of the TF representation and the assumption of sparsity: TF sparse methods are ill-suited to restore gaps that approach or exceed the duration of the TF analysis and synthesis windows. This limitation is also valid, if less severe, for structured TF sparsity, rendering the sparsity-based methods as unsatisfactory for inpainting medium duration gaps. TF domain is popular for inpainting short gaps, e.g., interpolation of audio based on a Gabor regression model [6], or nonnegative matrix and tensor factorization [31]–[33]. More recently, a powerful framework has been proposed for various audio inverse problems [34] including time-domain audio inpainting, source separation [35], and declipping [36] even in a multichannel scenario [37]. All of these systems require valid audio data within a time-domain window, cf. [36], which makes them perfect for inpainting short gaps, but unsatisfactory for medium gap durations.

On the other hand, for inpainting long gaps, recent methods leverage repetition and determine the most promising reliable segment from uncorrupted portions of the input signal [5], [7]. Restoration is then achieved by inserting the determined segment into the gaps. These methods do not claim to restore the missing gap perfectly, they aim at *plausibility*. For example, a method based on MFCC feature similarity has been proposed for packet loss concealment [5]. It explicitly targets a perceptually plausible restoration. Similarly, exemplarbased inpainting was proposed based on a graph encoding spectro-temporal similarities within an audio signal [7]. In both studies, gap durations were beyond several hundreds of milliseconds and their reconstruction needed to be evaluated in psychoacoustic experiments. Other examples for similar methods are [38]–[41]. While all these methods might be in general capable of inpainting gaps of medium duration, the target of the inpainting is always *plausible* instead of *accurate* reconstructions.

When restricting the inpainting to simple sounds such as musical instruments, linear prediction coding (LPC) [42] can be applied even for medium gap durations. While LPC may sound antiquated, it is particularly suitable for the instrument sounds as it models the way the sound is created by many instruments, i.e., by means of weighted sum of resonances. From the algorithmic perspective, LPC is simple but recursive, thus, allows to synthesize complex sound signals at a low computational power. Initially proposed for inpainting short bursts of lost samples [43], LPC-based inpainting algorithms model the signal as an acoustic source filtered by an all-pole filter. The model parameters are derived from the context and the missing signal part is synthesized by extrapolating the context into the gap. LPC-based methods work well for inpainting gaps for durations from 5 to 100 ms [3], [44]. LPC-based methods are particularly good in inpainting gaps consisting of many consecutive missing audio samples surrounded by reliable context [44]. In our experiments for medium gaps, the LPC-based algorithm [44] performed better than the latests reports on OMP-based algorithms [27]. As it seems, when it comes to inpainting medium gaps, the LPC-based method [44] seems to be the choice for a reference method.

The performance of LPC-based methods relies on the underlying assumption of signal stationarity. Deep-learning techniques, on the other hand, promise a more generalized signal representation. A combination of TF representation with deep-learning techniques may provide better inpainting whenever the lost data cannot be predicted by LPC. Thus, here, we propose to link deep-learning techniques with audio inpainting.

II. CONTEXT ENCODER

Our end-to-end system is presented in Fig. 1. We consider the audio signal s consisting of the gap s_g and the context signals before and after the gap, s_b and s_a , respectively (Fig. 1a). Given that convolutional networks applied directly on time-domain signals would require extremely large training datasets [45], we provide the network with TF coefficients. The TF coefficients are obtained from an invertible representation, namely, a redundant short-time Fourier transform (STFT) [46], [47]. Our network, inspired by the context encoder for image inpainting [8], is an encoder-decoder pipeline fed with TF coefficients of the context information, S_b and S_a (Fig. 1b). In order to study the general ability of DNNs to accurately inpaint audio in the range of tens of milliseconds, our network is comprised only of standard widely-used building blocks, i.e., convolutional layers, FCLs, and rectified linear units



Fig. 1. The end-to-end system. **a**) Audio signal in the time domain, s_g is the gap. **b**) Audio signal in the TF domain, S_b and S_a is the context before and after the gap, respectively. **c**) Reconstructed gap S_g' in the TF domain. **d**) Reconstruction S_g' merged with the stripped context S_b' and S_a' in the TF domain. **e**) Reconstructed signal in the time domain, including the inpainted gap, s_g' .

(ReLUs).¹ The network predicts TF coefficients of the gap S_g' (Fig. 1c), which are then merged with the stripped TF coefficients of the context, (Fig. 1d), in order to synthesize the reconstruction in the time domain, s' (Fig. 1e).

To study the effect of the phase of the reconstructed TF representations, we considered two equivalent networks with different outputs: (a) *complex network*, i.e., a network directly reconstructing the complex-valued TF coefficients which are then applied to the inverse STFT for the synthesis of the time-domain audio signal, and (b) *magnitude network*, i.e., a network reconstructing the magnitude coefficients only, which are then applied to a phase-reconstruction algorithm in order to obtain complex-valued TF coefficients required for the signal synthesis. From accurate TF magnitude information, phaseless reconstruction methods such as [48]–[50] are known to provide perceptually close, often indiscernible, reconstruction despite the resulting time-domain waveforms usually being rather different.

The software was implemented in Tensorflow [51] and is publicly available.²

A. Pre-processing stage

We use STFT, which enables a robust synthesis of the timedomain signal from the reconstructed TF coefficients.³ The STFT is determined by the analysis window, hop size a, and the number of frequency channels M. In our study, the analysis window was an appropriately normalized Hann window of length M and a was M/4, enabling perfect reconstruction by an inverse STFT with the same parameters and window.

The STFT is applied to the signal $s \in \mathbb{R}^{L}$ (containing L samples of audio) resulting in S, both of which consist of the context before and after the gap (containing L_c samples each) and the gap (containing L_g samples),

$$s = \begin{pmatrix} s_b \\ \mathbf{0}_{L_g imes 1} \\ s_a \end{pmatrix}$$
 and $S = (S_b, \mathbf{0}_{(M/2+1) imes N_g}, S_a)$,

¹Before fixing the network structure described in the remainder of this section, we experimented with different standard architectures, depths, and kernel shapes, out of which the current structure showed the most promise.

²www.github.com/andimarafioti/audioContextEncoder

³This is in contrast to machine-learning methods solving classification tasks, in which such a synthesis is not targeted.

where $s_b, s_a \in \mathbb{R}^{L_c}$, $N_g = (L_g - M)/a + 1$, and $S_b, S_a \in \mathbb{C}^{(M/2+1) \times N_c}$ with $N_c = L_c/a$. $\mathbf{0}_{R \times C}$ is a matrix with R rows and C columns containing only zeros.

Then, S_b and S_a are split into real and imaginary parts, resulting in four channels S_b^{Re} , S_b^{Im} , S_a^{Re} , S_a^{Im} , which are fed to the network.

B. Encoder

For the architecture of the encoder, [8] used the first five layers from [52] to process images. To adapt the design of our network to process TF coefficients, our encoder consists of six regular convolutional layers sequentially connected via ReLUs, after which batch normalization [53] is applied. Instead of using classical squared filters, we used rectangular filters to give the encoder more capacity on frequency over time in the TF representation. For M = 512, the resulting encoder architecture is shown in Figure 2.

The inputs S_b^{Re} , S_b^{Im} , S_a^{Re} , S_a^{Im} of the context information are treated as separate channels, thus, the network is required to learn how the channels interact and how to mix them. Because the encoder is comprised of only convolutional layers, the information can not reliably propagate from one end of the feature map to another. This is a consequence of convolutional layers connecting all the feature maps together, but never directly connecting all locations within a specific feature map [8].

C. Decoder

Similar to [8], the decoder begins with a FCL and a ReLU nonlinearity in order to spread the encoder's information among the channels. FCLs are computationally expensive; in our case it contains 38% of all the parameters of the network. All the subsequent layers are (de-)convolutional and, as for the encoder, connected by ReLUs with batch normalization. The first three layers use squared filters, the remaining two layers use rectangular filters to give the decoder more capacity on frequency over time in the output TF representation. Figure 3 shows the decoder architecture for M = 512 and a gap size $L_q = 1024$ samples.



Fig. 2. The encoder is a convolutional network with six layers followed by reshaping. The four channel TF input is encoded into a matrix of size of 2048. Gray rectangles represent the convolution filters with size expressed as (height, width). White cubes represent the signal.

The decoder does not only output the gap content, but also the TF coefficients connecting the gap with the context. Thus, the decoder output S_g' is larger than the original gap by M/a - 1 columns before and after the gap each, i.e., $S_g' \in \mathbb{C}^{(M/2+1)\times((L_g+M)/a-1)}$. In our example with $L_g = 1024, M = 512$ and a = M/4, shown in Fig. 3, every decoder output channel is of size 257×11 .

Note that the final layer depends on the network. For the complex network, the final layer has two outputs, corresponding to the real and imaginary part of the complex-valued TF coefficients. For the magnitude network, the final layer has a single output for the magnitude TF coefficients. We denote the output TF coefficients as S_{a}' .

D. Post-processing stage

The post-processing stage synthesizes the audio signal of the context and the inpainted gap. To this end, (M/a - 1) coefficients of the context extending into the gap are removed, yielding the stripped context, $S_b', S_a' \in \mathbb{C}^{(M/2+1)\times(N_c-M/a+1)}$. Then, the reconstructed TF coefficients from the decoder, S'_g , are inserted between the TF coefficient of the stripped context, S_b' and S_a' , yielding the sequence $S' = (S_b', S_g', S_a')$, having the same size as S. Stripping the context and insertion of the reconstruction directly in the TF domain prevents transitional artifacts between the context and the gap because synthesis by the inverse STFT introduces an inherent cross-fading.

For the complex network, the decoder output represents the real and imaginary parts of complex-valued TF coefficients $S_{q'}$ and the inverse STFT can be directly applied yielding s'.

For the magnitude network, the decoder output represents the magnitudes of the TF coefficients and the missing phase information needs to be estimated separately. First, the phase gradient heap integration algorithm proposed in [54] was applied to the magnitude coefficients produced by the decoder in order to obtain an initial estimation of the TF phase. Then, this estimation was refined by applying 100 iterations of the fast Griffin-Lim algorithm [48], [49]. We modified the version implemented in the Phase Retrieval Toolbox Library [55] to use the valid phase from the context at every iteration.⁴ The resulting complex-valued TF coefficients S_g' were then transformed into a time-domain signal s' by inverse STFT.

E. Loss Function

The network training is based on the minimization of the total loss of the reconstruction. To this end, the reconstruction loss is computed by comparing the original gap TF coefficients S_g with the reconstructed gap TF coefficients S_g' . Targeting an accurate reconstruction of the lost information, we optimize an adapted ℓ^2 -based loss instead of mixing the ℓ^2 -loss with an adversarial term [8]. For this type of network [56], the comparison can be done on the basis of the squared ℓ^2 -norm of the difference between S_g and S_g' , commonly known as mean squared error (MSE). The MSE would depend on the total energy of S_q , putting more weight on signals containing more energy. In order to avoid that, the normalized mean squared error (NMSE) can be used, which normalizes MSE by the energy of S_q . Compared to MSE, NMSE puts more weight on small errors when the energy of S_g is small. In practice, however, minor deviations from S_g are insignificant regardless of the content of S_q , and NMSE would be too sensitive.

Therefore, for the calculation of the loss function, we use a *weighted mix* between MSE and NMSE,

$$\mathbf{F}(S_g, S_g') = \frac{\|S_g - S_g'\|^2}{c^{-1} + \|S_g\|^2},\tag{1}$$

where the constant c > 0 controls the incorporated compensation for small amplitude. In our experiments, c = 5 yielded good results.

Finally, as proposed in [57], the total loss is the sum of the loss function and a regularization term controlling the trainable weights in terms of their ℓ^2 -norm:

$$\mathbf{T} = F(S_g, S_g') + \frac{\lambda}{2} \sum_i w_i^2, \qquad (2)$$

with w_i being weights of the network and λ being the regularization parameter, here, set to 0.01. The numerical

⁴The combination of these two algorithms provided consistently better results than separate application of either.



Fig. 3. The decoder architecture for the complex and magnitude network producing one and two channels of TF coefficients, respectively. All other conventions as in Figure 2.

optimizations were done using the stochastic gradient descent solver ADAM [58].

III. EVALUATION

The main objective of the evaluation was to investigate our networks' ability to adapt to audio signals. The evaluation is based on a comparison of the inpainting results to those obtained for the reference method, i.e., LPC-based extrapolation [44]. The inpainting quality was evaluated by means of objective difference grades (ODGs, [59]) and signal-to-noise ratios (SNRs) applied to the time-domain waveforms and magnitude spectrograms.

We considered two classes of audio signals: instrument sounds and music. The respective networks were trained on the targeted signal class, with an assumed gap size of 64 ms. Reconstruction was evaluated on the trained signal class and other signals for 64 ms gaps.

Additionally, we evaluated the effect of the gap duration by evaluating the magnitude network for 48 ms gaps.

A. Parameters

The sampling rate was 16 kHz. We considered audio segments with a duration of 320 ms, which corresponds to L = 5120 samples. For the STFT, the size of the window and the number of frequency channels M were fixed to 512 samples, and a was 128 samples.

Each segment was separated in a gap of 64 ms corresponding to $L_g = 1024$ of the central part of a segment and the context of twice of 128 ms, corresponding to $L_c = 2048$ samples. Consequently, N_c was 16, the input to the encoder was $S_b, S_a \in \mathbb{C}^{257 \times 16}$, and the output of the decoder was $S_g' \in \mathbb{C}^{257 \times 11}$.

B. Datasets

The dataset representing musical instruments was derived from the NSynth dataset [60]. NSynth is an audio dataset containing 305,979 musical notes from 1,006 instruments, each with a unique pitch, timbre, and envelope. Each example is four seconds long, monophonic, and sampled at 16 kHz.

The dataset representing music was derived from the free music archive (FMA, [61]). The FMA is an open and easily

accessible dataset, usually used for evaluating tasks in musical information retrieval. We used the small version of the FMA comprised of 8,000 30-s segments of songs with eight balanced genres sampled at 44.1 kHz. We resampled each segment to the sampling rate of 16 kHz.

The original segments in the two datasets were processed to fit the evaluation parameters. First, for each example the silence at the beginning and end was removed. Second, from each example, pieces of the duration of 320 ms were copied, starting with the first segment at the beginning of a segment, continuing with further segments with a shift of 32 ms. Thus, each example yielded multiple overlapping segments s. Then, the energy of the segments was evaluated and the ones that were completely silent were removed. Note that for a gap of 64 ms, the segment can be considered as a 3-tuple by labeling the first 128 ms as the context before the gap s_b , the subsequent 64 ms as the gap s_g , and the last 128 ms as the context after the gap s_a .

In order to avoid overfitting, the datasets were split into training, validation, and testing sets before segmenting them. For the instruments, we used the splitting proposed by [60]. The music dataset, was split into 70%, 20% and 10%, respectively. The statistics of the resulting sets are presented in Table I.

	Count	Percentage
Instruments training	19.4M	94.1
Instruments validation	0.9M	4.4
Instruments testing	0.3M	1.5
Music training	5.2M	70.0
Music validation	1.5M	20.0
Music testing	0.7M	10.0

TABLE I

SUBDIVISION OF THE DATASETS USED IN THE EVALUATION. COUNT IS THE AMOUNT OF EXAMPLES. PERCENTAGE IS CALCULATED WITH RESPECT TO THE FULL DATASET.

C. Evaluation metrics

The first metric was the SNR in dB,

$$\mathbf{SNR}(x, x') = 10 \ \log \frac{\|x\|^2}{\|x - x'\|^2} \tag{3}$$

calculated separately for each segment of a testing dataset. Then, we averaged SNRs across all segments of a testing dataset.

For the evaluation in the time domain, we used $\mathbf{SNR}(s_g, s_g')$, which is the SNR calculated on the gaps of the actual and reconstructed signals, s_g and s_g' , respectively. We refer to the average of this metric across all segments to as SNR in the time domain (SNR_{TD}).

The SNR was also calculated on the magnitude spectrograms in order to accommodate for perceptually less-relevant phase changes. We calculated $\mathbf{SNR}(|S_g|, |S_{rg}|)$, where S_{rg} represents the central 5 frames of the STFT computed from the restored signal s' and thus represents the restoration of the gap. In other words, we compute the SNR between the spectrograms of the original signal and the restored signal in the region of the gap. We refer to the average of this metric (across all segments of a testing dataset) to as SNR_{MS} , where MS stands for magnitude spectrogram. Note that SNR_{MS} is directly related to the spectral convergence proposed in [62].

Additionally, we computed the ODGs, which correspond to the subjective difference grade used in human-based audio test and is derived from the perceptual evaluation of audio quality (PEAQ, [59]). ODG range from 0 to -4 with the interpretation shown in Tab. II. We calculated the ODGs on signals of 2-s duration, with the inpainted gap beginning at 0.5-s. We used the algorithm implemented in [63].

ODG	Impairment
0	Imperceptible
-1	Perceptible, but not annoying
-2	Slightly annoying
-3	Annoying
-4	Very annoying
· ·	,

TABLE II INTERPRETATION OF ODGS.

D. Training

Both complex and magnitude networks were trained for the instrument and music dataset, resulting in four trained networks. Each training started with the learning rate of 10^{-3} . In the case of the magnitude network, the reconstructed phase was not considered in the training. Every 2000 steps, the training progress was monitored. To this end, signals from the validation dataset were inpainted and the weighted NMSE was calculated between the predicted and the actual TF coefficients of the gap. When converging, which usually happened after approximately 600k steps, the learning rate was reduced to 10^{-4} and the training was continued by additional 200k steps.⁵ Table III shows the SNR_{MS} calculated for the training, validation, and testing datasets. The similar values across subsets indicate no evidence for an overfitting.

-		Music			Instruments		
		Train	Valid	Test	Train	Valid	Test
Mag	Mean	7.6	7.8	7.8	22.1	21.9	21.9
	Std	4.2	4.0	4.3	9.9	10.2	10.0
Complex	Mean	4.9	5.1	5.4	17.8	18.3	18.2
	Std	4.0	4.2	4.5	10.5	10.3	10.1

TABLE III
OVERFITTING CHECK BY MEANS OF SNR_{MS} (in dB) calculated
BETWEEN GENERATED AND ORIGINAL TF-COEFFICIENTS WITHOUT THE
SYNTHESIS STED FOD 64 MS CADS

E. Reference method

We compared our results to those obtained with a reference method based on LPC. For the implementation, we followed [44], especially [44, Section 5.3]. In detail, the context signals s_b and s_a were extrapolated onto the gap s_g by computing their impulse responses and using them as prediction filters for a classical linear predictor. The impulse responses were obtained using Burg's method [64] and were fixed to have 1000 coefficients according to [2] and [65]. Their duration was the same as that for our context encoder in order to provide the same amount of context information. The two extrapolations were mixed with the squared-cosine weighting function. Our implementation of the LPC extrapolation is available online⁶.

Then, we evaluated the results produced by the reference method in the same way as we evaluated the results produced by the networks.

IV. RESULTS AND DISCUSSION

A. Ability to adapt to the training material

As a general rule, a trained neural network should perform well on the distribution that it learned from. As the instrument dataset is made of discrete in-tune instrument notes, each note can be considered as a sum of discrete frequencies arranged in time. If our network was able to adapt to the instrument sounds then it should perform on these frequencies better than on others.

To evaluate this, we probed our trained networks with stationary tones of various frequencies. The pure tones were directly synthesized as sine oscillations with a fixed frequency. The probes were generated within a logarithmic frequency range from 20 Hz to 8 kHz, linear phase shift range from 0 to π , and linear amplitude range from 0.1 to 1. The duration was 320 ms corresponding to 5120 samples at the sampling rate of 16 kHz.

Figure 4 shows the SNR_{MS} of the reconstruction obtained with the complex network. The abscissa shows notes, i.e., frequencies corresponding to the Standard pitch (with A corresponding to the frequency of 440 Hz). For the network trained on the instruments, the SNR_{MS} was large in the proximity of notes and decreased by more than 15 dB for frequencies between the notes. This shows that the network was able to better predict signals corresponding to the trained notes, indicating a good adaptation to the trained material.

Music contains more broadband sounds such as drums, breathing, tone glides, i.e., sounds with non-significant energy

⁵We also considered training on the instrument training dataset (800k steps) followed by a refinement with the music training dataset (300k steps). While it did not show substantial differences to the training performed on music only, a pre-trained network on music with a subsequent refinement to genre may show improvements for that genre.

⁶www.github.com/andimarafioti/audioContextEncoder



Fig. 4. SNR_{MS} for reconstruction of pure tones with the complex network trained on the instrument (black) and music (grey) dataset. SNR_{MS} are shown as a function of musical notes corresponding to the Standard pitch, i.e., the note A4 corresponds to the frequency of 440 Hz.

at frequencies between the Standard pitch being non-stationary even within the tested 320 ms. A network trained on music is expected to be less sensitive to predictions performed on Standard pitch only. Figure 4 shows the SNR_{MS} obtained for the reconstruction of pure tones with the network trained on the music. The SNR_{MS} fluctuations were smaller than those from the network trained on the instruments. This further supports our conclusion about the good ability of our network structure to adapt to various training materials.

B. Effect of the network type

The difference between the magnitude and complex networks both trained on instruments can be anticipated from the Figure 5, which shows the SNR_{MS} of the reconstructions of pure tones. As an average over frequency, the magnitude network provided an SNR_{MS} of 10.2 dB larger than that of the complex network. For the magnitude network, the SNR_{MS} was more or less similar for frequencies up to 200 Hz and decreased with frequency. For the complex network, the SNR_{MS} decrease started already at approximately 100 Hz and was much steeper than that of the magnitude network. Above the frequency of approximately 4 kHz, the complex network provided an extremely poor SNR_{MS} of 5 dB or less, indicating that the complex network had problems reconstructing the signals at higher frequencies. This is in line with [66], where neural networks were trained to reconstruct phases of amplitude spectrograms and their predictions were also poorer for higher frequencies.

Unfortunately, the problem of poor high-frequency reconstruction also persisted when predicting instrument sounds instead of pure tones. Figure 6 shows the spectrogram of an original sound from the instrument testing set (left panel) and of its reconstruction obtained from the complex network (center panel). The reconstruction clearly fails at frequencies higher than 4 kHz.

In order to further compare between the two network types, reconstructions of the testing datasets were performed. Table IV shows the SNR_{MS} and ODG of those predictions. The magnitude network resulted in consistently better results with an SNR_{MS} difference of 2.3 dB and 3.5 dB when tested



Fig. 5. SNR_{MS} for reconstruction of pure tones with the complex (black) and magnitude (grey) networks both trained to the instruments database. The thicker lines show averages over 25 surrounding frequency points.



Fig. 6. Magnitude spectrograms (in dB) of an exemplary signal reconstruction. Left: Original signal. Center: Reconstruction by the complex network. Right: Reconstruction by the LPC-based method. The gap was the area between the two red lines.

on music and instruments, respectively. Similarly, ODGs favor the magnitude network, although to a smaller extent. The comparison may appear flawed because the magnitude network has to predict only half of the features to be predicted by the complex network, at almost the same number of neurons. However, even doubling the size of the complex network would not yield significantly better predictions, as the link between the size of a DNN and its performance is not proportional [67].

In addition to the improvement in SNR_{MS} and ODG of the magnitude network over the complex network, the complex network predictions were observed to often be corrupted by clearly audible broadband noise⁷.

	Music			Instruments		
	Mag	Complex	LPC	Mag	Complex	LPC
Mean SNR _{MS}	7.7	5.4	6.3	22.4	18.5	30.5
Std SNR _{MS}	4.3	4.5	5.1	10.7	10.2	18.9
Mean ODG	-0.8	-1.0	-0.8	-1.6	-1.8	-0.3
Std ODG	0.4	0.2	0.2	1.0	0.9	0.3



 ${\rm SNR}_{\rm MS}$ (in DB) and ODGs of reconstructions of 64 ms gaps for the complex and magnitude networks, as well as for the LPC-based method.

⁷visit https://andimarafioti.github.io/audioContextEncoder/ for audio examples.

C. Comparison to the reference method

Table IV provides the SNR_{MS} and ODGs for the LPCbased reference reconstruction method. When tested on music, on average, our magnitude network outperformed the LPCbased method in terms of SNR_{MS} by 1.4 dB. When tested on instruments, our magnitude network underperformed the LPC by 8.6 dB, which was also reflected in poorer ODGs. Both SNRs and ODGs reveal a consistent picture. The LPC-based method seems to better inpaint instruments. The CE seems to be better or equivalent for inpainting music. This can be attributed to the better compliance of the instruments with the LPC, and a better universality of our CE.

In order to look more deeply into the differences between the two inpainting methods, we compared their abilities to inpaint frequency sweeps. A sweep represents a controlled frequency modulation, which violates the assumptions for the LPC and is not present in the data the CE was trained on. The signal consisted of a sum of five linear frequency sweeps with a 320-ms duration each, starting frequencies of 500, 2000, 3500, 5000 and 6500 Hz, and bandwidth of 500 Hz. Figure 7 shows the signal and the inpainting results. The gap inpainted by the LPC method (right panel) shows constant frequencies expanding into the gap causing a discontinuity in the gap's center. In contrast, the gap inpainted by the magnitude network (center panel) follows the frequency changes better at the price of noise appearing between the sweeps.

Other interesting examples are shown in Figure 8. The top row shows an example in which the magnitude network outperformed the LPC-based method. In this case, the signal is comprised of steady harmonic tones in the left side context and a broadband sound in the right side context. While the LPC-based method extrapolated the broadband noise into the gap, the magnitude network was able to foresee the transition from the steady sounds to the broadband burst, yielding a prediction much closer to the original gap, with a 13 dB larger SNR_{MS} than that from the LPC-based method.

On the other hand, the magnitude network did not always outperform the LPC-based method. The bottom row of Fig. 8 shows spectrograms of such an example. This signal had stable sounds in the gap, which were well-suited for an extrapolation, but rather complex to be perfectly reconstructed by the magnitude network. Thus, the LPC-based method outperformed the magnitude network yielding a 9 dB larger SNR_{MS}.



Fig. 7. Log-magnitude spectrograms (in dB) of an exponential frequency sweep. Left: Original signal. Center: Reconstruction by the magnitude network. Right: Reconstruction by the LPC-based method.



Fig. 8. Magnitude spectrograms (in dB) of exemplary signal reconstructions. Left: Original signal. Center: Reconstruction by the magnitude network. Right: Reconstruction by the LPC-based reference method. Top: Example with the magnitude network outperforming the reference by an SNR_{MS} of 13 dB. Bottom: Example with the magnitude network underperforming the reference by an SNR_{MS} of 9 dB.

Finally, Table V presents the SNR_{TD} of reconstructions of the instrument and music. Note that the SNR_{TD} provided for the magnitude network is for the sake of completeness only. The SNR_{TD} metric is highly sensitive to phase differences, which do not necessarily lead to perceptual differences and, for the magnitude network, is reconstructed with an accuracy of up to a constant phase shift. Thus, SNR_{TD} can remain low even in cases of very good reconstructions. Hence, here, we compare the performance of the complex network with that of the LPC-based method only.

For the music, on average, the complex network outperformed the LPC-based method providing a 0.3 dB larger SNR_{TD}. Given the large standard deviation, we performed a pair t-test on the SNR_{TD} which showed that the difference was statistically significant (p < 0.001). For the instruments, on average, the LPC-based reconstruction outperformed our network by 12 dB.

The excellent performance of the LPC-based method reconstructing instruments can be explained by the assumptions behind the LPC well-fitting to the single-note instrument sounds. These sounds usually consist of harmonics stable on a short-time scale. LPC extrapolates these harmonics preserving the spectral envelope of the signal. Nevertheless, the magnitude network yielded an SNR_{MS} of 22.4 dB, on average, demonstrating a good ability to reconstruct instrument sounds.

When applied on music, the performance in terms of SNR_{MS} of both methods was much poorer, with our network performing slightly but statistically significantly better than the LPC-based method. The better performance of our network can be explained by its ability to adapt to transient sounds and modulations in frequencies, sound properties that the LPC-based method is not suited to handle.

The gap duration of 64 ms is close to those tested in [27] when comparing various OMP methods. For 50 ms, their approaches showed SNR_{TD} below 2 dB and ODG values around -3 (see their Fig. 1 and 4). The LPC-based method showed average SNR_{TD} of 3.8 dB and ODGs of -0.8. This confirms our assumption that for the studied range, the LPC is better suited than the sparsity-based audio inpainting techniques.

	Music			Instruments		
	Complex	Mag	LPC	Complex	Mag	LPC
Mean	3.8	1.1	3.5	16.0	14.6	28.0
Std	4.1	3.9	5.0	9.7	10.8	19.1

TABLE V
SNR_{TD} (in dB) of reconstructions of 64 ms gaps for the
COMPLEX AND MAGNITUDE NETWORKS, AS WELL AS FOR THE
LPC-BASED METHOD

D. Effect of the gap duration

The proposed network structure can be trained with different contexts and gap durations. For problems of varying gap duration, a network trained to the particular gap duration might appear optimal. However, training takes time, and it might be simpler to train a network to single gap duration and use it to reconstruct any shorter gap as well.

In order to test this idea, we introduced gaps of 48 ms (corresponding to $L_g = 768$ samples) in our testing datasets. These gaps were then reconstructed by the magnitude network trained for 64 ms gaps. As this network outputs, at reconstruction time, a solution for a gap of length 64-ms, the 48-ms gaps needs to be enlarged. We tested three approaches to enlarge them: by discarding 16 ms forwards, 16 ms backwards, or 8 ms forwards and 8 ms backwards (centered).

Table VI shows SNR_{MS} obtained from averaging the reconstructions of the three types of gap enlargements. Also, the corresponding SNR_{MS} for the LPC-based method are shown. The results are similar to those obtained for larger gaps: for the instruments, the LPC-based method outperformed our network; for the music, our network outperformed the LPC-based method.

	Mu	isic	Instruments		
	Ours	LPC	Ours	LPC	
Mean	8.0	6.9	21.8	33.2	
Std	4.6	5.5	11.8	20.1	

TABLE VI SNR_{MS} (in dB) of reconstructions of 48 ms gaps for the magnitude network and the LPC-based method.

V. CONCLUSIONS AND OUTLOOK

We proposed a neural network architecture for inpainting *medium* gaps of audio. The study aims at showing general abilities of a neural network working on TF coefficients as a context encoder. The proposed network was able to adapt to the particular frequencies provided by the training material. It was able to reconstruct frequency modulations better than the LPC-based reference method and it was able to inpaint gaps shorter than the trained ones. For the reconstruction of complex signals like music, our network was able to

outperform the LPC-based reference method, in terms SNR calculated on magnitude spectrograms, and both methods were rated equally with ODG between *imperceptible* and *perceptible but not annoying*. LPC yielded better results when applied on more simple signals like instrument sounds. In general, our results suggest that standard DNN components and a moderately sized network can be applied to form audio-inpainting models, offering a number of angles for future improvement.

For example, we have analyzed two types of networks. The complex network works directly on the complex-valued TF coefficients. The magnitude network provides only magnitudes of TF coefficients as output and relies on a subsequent phase reconstruction. We observed clear improvement of the magnitude network over the complex network especially in reconstructing high-frequency content.

From our study, it follows that DNNs, when applied to inpainting audio gaps for medium durations, do not suffer from the restrictions of previous methods. Additionally, even for a simple DNN, the performance on complex signals is already on par with the state of the art. It also follows that by representing audio as TF coefficients, a generative network developed for image inpainting can be adapted to audio inpainting.

Generally, better results can be expected for increased depth of the network and the available context. Experiments with our method for longer medium-duration gaps and longer context can be easily implemented just by adapting the parameters of the network. Nevertheless, we expect technical limitations like computational power to be an issue for long contexts. Instead, a study of more efficient audio features will be required. Our STFT features, meant in this study as a reasonable first choice, provided a decent performance, however, in the future, we expect hearing-related features to provide better reconstructions. In particular, an investigation of Audlet frames, i.e., invertible time-frequency systems adapted to perceptual frequency scales, [68], as features for audio inpainting seem to offer intriguing opportunities.

In the future, instead of training on a very general dataset, improved performance can be obtained for more specialized networks trained to specific genres or instrumentation. Further, applied to a complex mixture and potentially preceded by a source-separation algorithm, our proposed architecture could be used jointly in a mixture-of-experts, [69], approach.

REFERENCES

- A. Adler, V. Emiya, M. G. Jafari, M. Elad, R. Gribonval, and M. D. Plumbley, "Audio inpainting," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 20, no. 3, pp. 922–932, March 2012.
- [2] I. Kauppinen, J. Kauppinen, and P. Saarinen, "A method for long extrapolation of audio signals," *Journal of the Audio Engineering Society*, vol. 49, no. 12, pp. 1167–1180, 2001.
- [3] W. Etter, "Restoration of a discrete-time signal segment by interpolation based on the left-sided and right-sided autoregressive parameters," *IEEE Transactions on Signal Processing*, vol. 44, no. 5, pp. 1124–1135, may 1996.
- [4] D. Goodman, G. Lockhart, O. Wasem, and W.-C. Wong, "Waveform substitution techniques for recovering missing speech segments in packet voice communications," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 34, no. 6, pp. 1440–1448, dec 1986.

- [6] P. J. Wolfe and S. J. Godsill, "Interpolation of missing data values for audio signal restoration using a gabor regression model," in *Proc. of ICASSP*, vol. 5. IEEE, 2005, pp. v–517.
- [7] N. Perraudin, N. Holighaus, P. Majdak, and P. Balazs, "Inpainting of long audio segments with similarity graphs," *IEEE/ACM Transactions* on Audio, Speech and Language Processing, vol. PP, no. 99, pp. 1–1, 2018.
- [8] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. Efros, "Context encoders: Feature learning by inpainting," in *Proc. of CVPR*, 2016.
- [9] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, http://www.deeplearningbook.org.
- [10] D. Kingma and M. Welling, "Auto-encoding variational bayes." in *Proc.* of *ICLR*, 2014.
- [11] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672– 2680.
- [12] C. Donahue, J. McAuley, and M. Puckette, "Adversarial audio synthesis," in *Proceedings of the 7th International Conference on Learning Representations*, 2019.
- [13] A. Marafioti, N. Perraudin, N. Holighaus, and P. Majdak, "Adversarial generation of time-frequency features with application in audio synthesis," in *Proc. of the 36th ICML*, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. Long Beach, California, USA: PMLR, 09–15 Jun 2019, pp. 4352–4362. [Online]. Available: http://proceedings.mlr.press/v97/marafioti19a.html
- [14] J. Engel, K. K. Agrawal, S. Chen, I. Gulrajani, C. Donahue, and A. Roberts, "Gansynth: Adversarial neural audio synthesis," in *Proceedings of the 7th International Conference on Learning Representations*, 2019.
- [15] S. Mehri, K. Kumar, I. Gulrajani, R. Kumar, S. Jain, J. Sotelo, A. Courville, and Y. Bengio, "SampleRNN: An unconditional end-toend neural audio generation model," in *Proc. of ICLR*, 2017.
- [16] 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," *CoRR*, vol. abs/1609.03499, 2016. [Online]. Available: http://arxiv.org/abs/1609.03499
- [17] Y. Saito, S. Takamichi, and H. Saruwatari, "Text-to-speech synthesis using STFT spectra based on low-/multi-resolution generative adversarial networks," in *Proc. of ICASSP*. IEEE, 2018, pp. 5299–5303.
- [18] J. Shen, R. Pang, R. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan, R. Saurous, Y. Agiomyrgiannakis, and Y. Wu, "Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions," in *Proc. of ICASSP*. IEEE, 2018.
- [19] Z. Jin, A. Finkelstein, G. J. Mysore, and J. Lu, "Fftnet: A real-time speaker-dependent neural vocoder," in *Proc. of ICASSP*. IEEE, 2018, pp. 2251–2255.
- [20] B.-K. Lee and J.-H. Chang, "Packet loss concealment based on deep neural networks for digital speech transmission," *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 24, no. 2, pp. 378–387, Feb. 2016. [Online]. Available: http://dx.doi.org/10.1109/TASLP.2015.2509780
- [21] S. Dieleman, A. v. d. Oord, and K. Simonyan, "The challenge of realistic music generation: modelling raw audio at scale," in *Proc. of NeurIPS*, 2018.
- [22] N. Boulanger-Lewandowski, Y. Bengio, and P. Vincent, "Modeling temporal dependencies in high-dimensional sequences: Application to polyphonic music generation and transcription," in *Proc. of ICML*, 2012.
- [23] M. Blaauw and J. Bonada, "A neural parametric singing synthesizer," in *Proc. of INTERSPEECH*, 2017.
- [24] A. Adler, V. Emiya, M. Jafari, M. Elad, R. Gribonval, and M. Plumbley, "A constrained matching pursuit approach to audio declipping," in *Proc.* of ICASSP. IEEE, may 2011.
- [25] I. Toumi and V. Emiya, "Sparse non-local similarity modeling for audio inpainting," in *Proc. of ICASSP*. Calgary, Canada: IEEE, Apr. 2018.
- [26] S. Kitić, N. Bertin, and R. Gribonval, "Sparsity and cosparsity for audio declipping: a flexible non-convex approach," in LVA/ICA 2015 - The 12th International Conference on Latent Variable Analysis and Signal Separation, Liberec, Czech Republic, Aug. 2015, p. 8. [Online]. Available: https://hal.inria.fr/hal-01159700
- [27] O. Mokrý, P. Záviska, P. Rajmic, and V. Veselý, "Introducing SPAIN (sparse audion inpainter)," *CoRR*, vol. abs/1810.13137, 2018. [Online]. Available: http://arxiv.org/abs/1810.13137
- [28] C. Gaultier, S. Kitić, N. Bertin, and R. Gribonval, "AUDASCITY: AUdio Denoising by Adaptive Social CosparsITY," in 25th European

Signal Processing Conference (EUSIPCO), Kos, Greece, Aug. 2017. [Online]. Available: https://hal.inria.fr/hal-01540945

- [29] K. Siedenburg, M. Kowalski, and M. Dörfler, "Audio Declipping with Social Sparsity," in *Proc. of ICASSP*. Florence, Italy: IEEE, May 2014, pp. AASP–L2. [Online]. Available: https://hal.archives-ouvertes. fr/hal-01002998
- [30] F. Lieb and H.-G. Stark, "Audio inpainting: Evaluation of time-frequency representations and structured sparsity approaches," *Signal Processing*, vol. 153, pp. 291–299, 2018.
- [31] J. Le Roux, H. Kameoka, N. Ono, A. De Cheveigne, and S. Sagayama, "Computational auditory induction as a missing-data model-fitting problem with bregman divergence," *Speech Communication*, vol. 53, no. 5, pp. 658–676, 2011.
- [32] P. Smaragdis, B. Raj, and M. Shashanka, "Missing data imputation for time-frequency representations of audio signals," *Journal of signal* processing systems, vol. 65, no. 3, pp. 361–370, 2011.
- [33] U. Şimşekli, Y. K. Yılmaz, and A. T. Cemgil, "Score guided audio restoration via generalised coupled tensor factorisation," in *Proc. of ICASSP.* IEEE, 2012, pp. 5369–5372.
- [34] C. Bilen, A. Ozerov, and P. Prez, "Solving time-domain audio inverse problems using nonnegative tensor factorization," *IEEE Transactions on Signal Processing*, vol. 66, no. 21, pp. 5604–5617, Nov 2018.
- [35] Ç. Bilen, A. Ozerov, and P. Pérez, "Joint audio inpainting and source separation," in *International Conference on Latent Variable Analysis and Signal Separation*. Springer, 2015, pp. 251–258.
- [36] —, "Audio declipping via nonnegative matrix factorization," in 2015 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA). IEEE, 2015, pp. 1–5.
- [37] A. Ozerov, Ç. Bilen, and P. Pérez, "Multichannel audio declipping," in *Proc. of ICASSP*. IEEE, 2016, pp. 659–663.
- [38] E. Manilow and B. Pardo, "Leveraging repetition to do audio imputation," in 2017 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA). IEEE, 2017, pp. 309–313.
- [39] B. Martin, P. Hanna, T. V. Thong, M. Desainte-Catherine, and P. Ferraro, "Exemplar-based assignment of large missing audio parts using string matching on tonal features." in *Proc. of ISMIR*, 2011, pp. 507–512.
- [40] R. C. Maher, "A method for extrapolation of missing digital audio data," *Journal of the Audio Engineering Society*, vol. 42, no. 5, pp. 350–357, 1994.
- [41] A. Lukin and J. Todd, "Parametric interpolation of gaps in audio signals," in Audio Engineering Society Convention 125. Audio Engineering Society, 2008.
- [42] T. E. Tremain, "The government standard linear predictive coding algorithm: Lpc-10," Speech Technology, pp. 40–49, Apr. 1982.
- [43] A. Janssen, R. Veldhuis, and L. Vries, "Adaptive interpolation of discrete-time signals that can be modeled as autoregressive processes," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 34, no. 2, pp. 317–330, 1986.
- [44] I. Kauppinen and K. Roth, "Audio signal extrapolation-theory and applications," in *Proc. DAFx*, 2002, pp. 105–110.
- [45] J. Pons, O. Nieto, M. Prockup, E. M. Schmidt, A. F. Ehmann, and X. Serra, "End-to-end learning for music audio tagging at scale," in *Proc. of ISMIR*, 2018.
- [46] M. Portnoff, "Implementation of the digital phase vocoder using the fast fourier transform," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 24, no. 3, pp. 243–248, 1976.
- [47] K. Gröchenig, Foundations of Time-Frequency Analysis, ser. Appl. Numer. Harmon. Anal. Birkhäuser, 2001.
- [48] D. Griffin and J. Lim, "Signal estimation from modified short-time fourier transform," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 32, no. 2, pp. 236–243, 1984.
- [49] N. Perraudin, P. Balazs, and P. L. Søndergaard, "A fast griffin-lim algorithm," in *Applications of Signal Processing to Audio and Acoustics* (WASPAA), 2013 IEEE Workshop on. IEEE, 2013, pp. 1–4.
- [50] Z. Průša, P. Balazs, and P. Søndergaard, "A noniterative method for reconstruction of phase from stft magnitude," *IEEE/ACM Transactions* on Audio, Speech and Language Processing, vol. 25, no. 5, pp. 1154– 1164, 2017.
- [51] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems,"

2015, software available from tensorflow.org. [Online]. Available: https://www.tensorflow.org/

- [52] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proc. of NIPS*, 2012, pp. 1097–1105.
- [53] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. of ICML*, 2015, pp. 448–456.
- [54] Z. Průša and P. L. Søndergaard, "Real-Time Spectrogram Inversion Using Phase Gradient Heap Integration," in *Proc. Int. Conf. Digital Audio Effects (DAFx-16)*, Sep 2016, pp. 17–21.
- [55] Z. Průša, "The Phase Retrieval Toolbox," in AES International Conference On Semantic Audio, Erlangen, Germany, June 2017.
- [56] H. Zhao, O. Gallo, I. Frosio, and J. Kautz, "Loss functions for image restoration with neural networks," *IEEE Transactions on Computational Imaging*, vol. 3, no. 1, pp. 47–57, March 2017.
- [57] A. Krogh and J. Hertz, "A simple weight decay can improve generalization," in Advances in neural information processing systems 4. Morgan Kaufmann, 1992, pp. 950–957.
- [58] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. of ICLR*, 2015.
- [59] I. Recommendation, "1387: Method for objective measurements of perceived audio quality," *International Telecommunication Union, Geneva*, *Switzerland*, 2001.
- [60] J. Engel, C. Resnick, A. Roberts, S. Dieleman, M. Norouzi, D. Eck, and K. Simonyan, "Neural audio synthesis of musical notes with wavenet autoencoders," in *Proc. of ICML*, 2017, pp. 1068–1077.
- [61] M. Defferrard, K. Benzi, P. Vandergheynst, and X. Bresson, "Fma: A dataset for music analysis," in 18th International Society for Music Information Retrieval Conference, 2017.
- [62] N. Sturmel and L. Daudet, "Signal reconstruction from stft magnitude: A state of the art," in *International conference on digital audio effects* (*DAFx*), 2011, pp. 375–386.
- [63] P. Kabal et al., "An examination and interpretation of itu-r bs. 1387: Perceptual evaluation of audio quality," TSP Lab Technical Report, Dept. Electrical & Computer Engineering, McGill University, pp. 1–89, 2002.
- [64] J. P. Burg, "Maximum entropy spectral analysis," 37th Annual International Meeting, Soc. of Explor. Geophys., Oklahoma City, 1967.
- [65] I. Kauppinen and J. Kauppinen, "Reconstruction method for missing or damaged long portions in audio signal," *Journal of the Audio Engineering Society*, vol. 50, no. 7/8, pp. 594–602, 2002.
- [66] S. Takamichi, Y. Saito, N. Takamune, D. Kitamura, and H. Saruwatari, "Phase reconstruction from amplitude spectrograms based on von-misesdistribution deep neural network," in *International Workshop on Acoustic Signal Enhancement (IWAENC)*, 2018, pp. 286–290.
- [67] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2016, pp. 770–778.
- [68] T. Necciari, N. Holighaus, P. Balazs, Z. Pra, P. Majdak, and O. Derrien, "Audlet filter banks: A versatile analysis/synthesis framework using auditory frequency scales," *Applied Sciences*, vol. 8, no. 1:96, 2018.
- [69] S. E. Yuksel, J. N. Wilson, and P. D. Gader, "Twenty years of mixture of experts," *IEEE transactions on neural networks and learning systems*, vol. 23, no. 8, pp. 1177–1193, 2012.