

DEEP AUDIO PRIORS EMERGE FROM HARMONIC CONVOLUTIONAL NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

Convolutional neural networks (CNNs) excel in image recognition and generation. Among many efforts to explain their effectiveness, experiments show that CNNs carry strong inductive biases that capture natural image priors. Do deep networks also have inductive biases for audio signals? In this paper, we empirically show that current network architectures for audio processing do not show strong evidence in capturing such priors. We propose *Harmonic Convolution*, an operation that helps deep networks distill priors in audio signals by explicitly utilizing the harmonic structure within. This is done by engineering the kernel to be supported by sets of harmonic series, instead of local neighborhoods for convolutional kernels. We show that networks using Harmonic Convolution can reliably model audio priors and achieve high performance in unsupervised audio restoration tasks. With Harmonic Convolution, they also achieve better generalization performance for sound source separation.

1 INTRODUCTION

Deep neural networks, in various forms and designs, have been proved extremely successful in both discriminative tasks such as image classification (He et al., 2016), machine translation (Sutskever et al., 2014), speech recognition (Hinton et al., 2012), and in generative tasks such as image and audio generation (Goodfellow et al., 2014; Oord et al., 2017). Recently, Lempitsky et al. (2018) showed that convolutional neural networks (CNNs) come with strong inductive biases to capture natural image priors purely by their structure. Such observation provides an intriguing perspective on the effectiveness of CNNs in the generative modeling of images.

However, unlike CNNs for image modeling, the design of deep neural networks for auditory signals has not yet converged, where temporal and spectral-temporal representations are all being actively explored. In addition, due to the success of the CNNs, audio processing networks usually adopts similar designs from their image counterparts, especially ones operating on spectrograms. To what degree are those designs well justified? Is it possible to tell which one might be more effective? Is it possible that audio signal modeling needs components that can not be found in designs for processing images?

As an attempt to answer the questions above, we investigate whether various audio processing networks can capture audio priors, similar to what CNNs do for images (Lempitsky et al., 2018). Recently, Michelashvili & Wolf (2019) reported that deep priors do exist for Wave-U-Net (Stoller et al., 2018), which can be exploited to perform unsupervised audio denoising. However, the definition of deep priors are slightly different between works by Michelashvili & Wolf (2019) and Lempitsky et al. (2018). Michelashvili & Wolf (2019) reported that the noisy signal causes more violent fluctuations on the spectrogram, which is then utilized as a prior for estimating the noise signal. On the contrary, we are more focused on the deep priors defined by Lempitsky et al. (2018). we try to investigate whether there is evidence that current deep networks for audio related tasks carry inductive biases for audio priors.

In this work, we first empirically show that current architectures for audio signal modeling do not show strong evidence for capturing audio priors. Specifically, we look at two general types of design: temporal CNNs (Michelashvili & Wolf, 2019; Stoller et al., 2018; Aytar et al., 2016) and spectrogram-based CNNs (Shen et al., 2018; Zhao et al., 2018). To test their prior modeling ability, we use the setup identical to (Lempitsky et al., 2018): the networks are initialized randomly, and

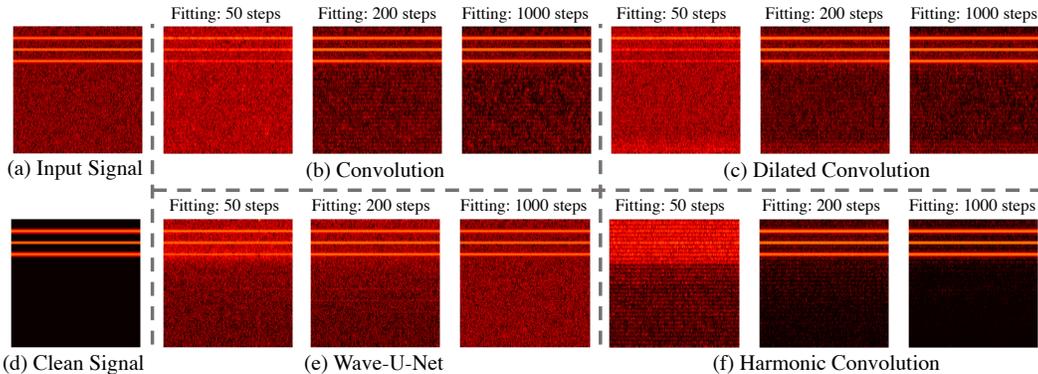


Figure 1: A simple illustrative case: fitting a harmonic sound under Gaussian noise. (a)(d) The corrupted signal and its clean version. The network only sees (a) and aims to fit it. (b)(c)(e) Both temporal and spectral-temporal convolution networks start with a very noisy output at 50 iterations, then fits the signal as well as the back ground noise at 200 iterations. At 1,000 iterations, the output is similar to the one at 200 iterations. This suggest that the network fits the noise and the signal without obvious distinctions. (f) Harmonic Convolution start with a noisy output as well, but at 200 steps, the network selectively fits the harmonic series instead. Note that the 1,000-step result is even cleaner than the 200-step one.

tasked to fit a single degenerated audio signal. If the network is capable of modeling the signal priors by construction, it would fit the signal faster than noise. We show an illustrative example in Figure 1, where no strong evidence is found for temporal or spectral-temporal CNNs, even with this simplest case.

What might be missing? As psychoacoustics experiments have shown (Moore et al., 1986; Popham et al., 2018), the structure of harmonic series are closely related to human perception. We therefore propose the *Harmonic Convolution*, an operation explicitly utilizes harmonic structures in audio signals. Then, with multiple experiments, we show that Harmonic Convolution does enable neural networks better model audio signal priors.

Finally, we show that Harmonic Convolution are useful in downstream applications, and prove its performance by comparing against various baselines. The most natural application is unsupervised audio restoration, where we aim to recover a clean signal from a degenerated audio signal corrupted by a high power Gaussian noise or aggressive quantization. In addition, we also demonstrate that networks with Harmonic Convolution achieve better generalization performances for supervised sources separation tasks.

In summary, our contributions are threefold. First, we show that current audio processing architectures do not model audio signal priors naturally. Second, we propose an operation called Harmonic Convolution, serving as an effective inductive bias for neural networks to model audio signal priors. Finally, we demonstrate that networks with Harmonic Convolution achieve state-of-art performances in unsupervised audio restoration tasks as well as improve the generalization ability on supervised musical source separation tasks.

2 MOTIVATION

In this section, we first give a brief review on deep image priors. We also provide a short survey on current popular network architecture designs for processing audio signals. Then, we show a motivating toy example, where current architectures fail to model the signal priors, even when the signal is purely stationary and harmonic. Finally, we provide a heuristic analysis on why such convolution-based approach may not capture the audio priors, using local signal statistics.

2.1 DEEP PRIORS

Lempitsky et al. (2018) first proposed the notion of deep priors on images. Specifically, they show that given a corrupted image x_0 , a deep neural network f_θ , parameterized by θ , can serve as a natural regularization without any pertaining. Formally, the deep prior method optimizes

$$\min_{\theta} E(f_\theta(z); x_0),$$

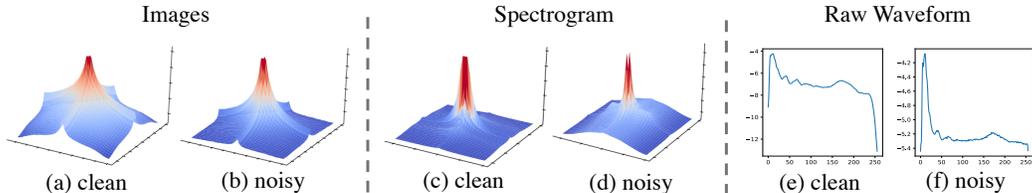


Figure 2: The frequency statistics for images, spectrograms and raw waveform. All magnitudes in this figure is on log scale. The DC component is not shown for visualization purposes. (a) clean image patch statistics. the $1/f^2$ law can be observed. (b) noisy image patch statistics. (c) clean speech spectrogram statistics. Note that this is more similar to (b) instead of (a). (d) noisy speech spectrogram statistics. Note that the frequency dimension and the temporal dimension have different behaviours. (e) clean speech statistics. The power distribution spreads to higher frequencies. (f) noisy speech statistics.

where $E(\cdot; \cdot)$ is a task-specific data term, and z is a fixed random noise vector. Lempitsky et al. (2018) showed that searching in the parameter space of the neural network f_θ is surprisingly effective: after several iterations, the optimized network parameters θ^* gives the restored image x by forwarding the noise vector, i.e. $x = f_{\theta^*}(z)$. This fact suggests that CNNs might be well suited for modeling images, where its structure and operation provides strong inductive biases.

2.2 CURRENT NETWORK DESIGNS FOR AUDIO PROCESSING

The network architectures for audio signal processing fall into two broad categories. The first one is to directly apply 1D convolutions on the raw audio signals (Michelashvili & Wolf, 2019; Stoller et al., 2018; Aytar et al., 2016). For instance, Wave-U-Net (Stoller et al., 2018) is a 1D adaptation of the U-Net architecture, which utilizes skip connections and multiscale processings to model signal properties at different time scales. The other category is characterized by performing 2D convolutions on spectrograms (Shen et al., 2018; Zhao et al., 2018). A common practice is to first extract a spectral-temporal representation from the audio signal, then apply 2D convolutions.

2.3 A MOTIVATING EXAMPLE

Inspired by the deep image prior experiments, we would like to see if the deep architectures above possess the proper inductive biases to model audio signals. To this end, we first test the above architectures to reconstruct a simple signal: a stationary signal composed of 1,000Hz, 2,000Hz and 3,000Hz sinusoidal waves. The corrupted version of this signal is generated by adding a stationary Gaussian noise with a standard deviation of 0.1. As can be seen from Figure 1, all the methods start with very noisy fittings at very early iterations (50 in the figure). Then, they start to fit the signal and the noise at a similar speed, rendering a noisy signal at 200 steps. At 1,000 steps, the network would fit to a noisy signal, with slightly less noise than the input signal.

2.4 A HEURISTIC ANALYSIS

Here we provide a heuristic analysis on why plain convolution-based networks would fail for modeling audio signal priors by structure. To begin with, We assume that the deep prior phenomenon for images indicates that the CNNs are well suited to model natural image statistics (Torralba & Oliva, 2003). In Figure 2, we show side by side the frequency statistics of clean natural images, clean speech spectrograms, clean speech signals and their noisy versions by adding a Gaussian noise to such that corrupted images and speech share the similar Noise-to-Signal ratio. The image statistics are computed by randomly sampling 1,000 images from ImageNet (Russakovsky et al., 2015). The audio spectrogram statistics is produced by randomly sampling 1,000 speech signals from the JL-Speech dataset (Ito, 2017), computing their spectrograms using the Short-Time-Fourier-Transform, and calculate the spatial frequency distribution of the spectrograms as if they are images. The 1D frequency distribution of the audio signals is calculated by the same speech signals, interpreting audio clips like 1D image patches. Note that the $1/f^2$ law for natural image statistics (Torralba & Oliva, 2003) do not hold up for either spectrograms or raw waveforms.

3 APPROACH

As shown in Figure 1, the architectures above, despite their success under supervised training settings, do not show strong evidence for having deep audio priors. In light of such facts, we aim to

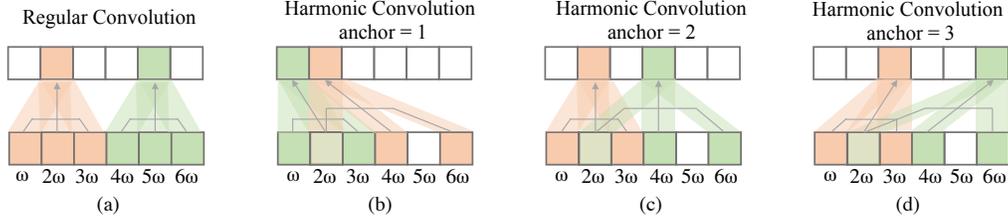


Figure 3: Illustrations for regular convolutions and Harmonic Convolution. (a) regular convolution kernels are supported by a local region, the shape of the support is transnational invariant. (b) Harmonic Convolution with anchoring parameter as 1. This interprets the output frequency as the fundamental frequency for the harmonic series. (c) With an anchoring of 2, the output frequency is the second lower frequency of a harmonic series. (d) Similarly, anchoring = 3 interprets output location as the third harmonics.

introduce new components for audio signal modeling. We start with harmonic structures, the most common patterns for audio signals. Subsequently, we propose Harmonic Convolution, an operation that treats harmonic structures explicitly.

3.1 HARMONIC CONVOLUTION

Harmonic Convolution is designed to explicitly exploit harmonic structures as an inductive bias for auditory signal modeling. Specifically, Harmonic Convolution operates on the discrete Short-Time Fourier Transform (STFT) of a given audio signal, utilizing the spectral-temporal structure of this representation. For audio signal $x[t]$, its discrete STFT spectrogram $X[\omega, \tau]$ is given by

$$X[\omega, \tau] = \left\| \sum_{t=-\infty}^{\infty} x[t]w[t - \tau]e^{-j\omega t} \right\|^2, \quad (1)$$

where $w(\cdot)$ is a locally supported windowing function and j denotes the imaginary unit. Regular 2D convolutions on the spectrogram $X[\omega, \tau]$ with a kernel function K is defined as:

$$(X * K)[\hat{\omega}, \hat{\tau}] = \sum_{\omega=-\Omega}^{\Omega} \sum_{\tau=-T}^T X[\hat{\omega} - \omega, \hat{\tau} - \tau]K[\omega, \tau], \quad (2)$$

where the kernel K is supported on $[-\Omega, \Omega] \times [-T, T]$. Note that regular convolution aggregates information in a $2\Omega \times 2T$ window on local regions of X . To utilize harmonic structures, we modify this information aggregation scheme to align with harmonics. Specifically, Harmonic Convolution is defined as an operation mapping $X(\omega, \tau)$ to $Y(\hat{\omega}, \hat{\tau})$, where

$$Y(\hat{\omega}, \hat{\tau}) = \sum_{k=1}^K \sum_{\tau=-T}^T X[k\hat{\omega}, \hat{\tau} - \tau]K[k, \tau]. \quad (3)$$

Note Harmonic Convolution interprets the frequency dimension of the kernel as weights for K harmonic series at each target frequency location $\hat{\omega}$, where regular convolutions interpret the kernel as weights for a local neighborhood at target spectral temporal locations. Figure 3 shows an illustration for both plain convolutions and Harmonic Convolution.

3.2 ANCHORS AND MIXING

As indicated by Equation 3, the output at target frequency $\hat{\omega}$ is a weight sum of its N harmonic series, starting from $\hat{\omega}$ to $N\hat{\omega}$. Note that there also exists other possible harmonic series that include $\hat{\omega}$. For example: $0.25\hat{\omega}, 0.5\hat{\omega}, 0.75\hat{\omega}, \hat{\omega}, \dots$ is also a valid harmonic series, but Equation 3 never aggregates the information from frequencies lower than $\hat{\omega}$. To over come this problem, we add an extra parameter to Equation 3 called anchor, which indicates the order of harmonics at target frequency location $\hat{\omega}$. Specifically, given the anchoring parameter n , we modify Equation 3 as:

$$Y_n[\hat{\omega}, \hat{\tau}] = \sum_{k=1}^K \sum_{\tau=-T}^T X \left[\frac{k\hat{\omega}}{n}, \hat{\tau} - \tau \right] K[k, \tau]. \quad (4)$$

An illustration for the effect of different anchoring parameters are shown in Figure 3(b)(c)(d). In addition, we can make the output at the frequency location $\hat{\omega}$ depend on multiple anchoring parameters. To this end, we mix different Y_n using a weighted sum: $Y[\hat{\omega}, \hat{\tau}] = \sum_{n=1}^N w_n Y_n[\hat{\omega}, \hat{\tau}]$,

where N is the largest anchoring parameter, and w_n can be seen as learnable parameters, similar to convolution kernels K . Therefore, the final Harmonic Convolution is defined as:

$$Y[\hat{\omega}, \hat{\tau}] = \sum_{n=1}^N \sum_{k=1}^K \sum_{\tau=-T}^T w_n X \left[\frac{k\hat{\omega}}{n}, \hat{\tau} - \tau \right] K[k, \tau], \quad (5)$$

where the learnable parameters are the convolution kernel K and the weights w_n .

Implementation details. We implement Equation 5 using the Deformable Convolution operation introduced by Dai et al. (2017). For better efficiency, we factorize the 2D kernel $K[k, \tau]$ as the product of two 1D kernels, i.e. $K[k, \tau] = K_f[k]K_t[\tau]$. Anchoring is implemented using grouped Deformable Convolution (Dai et al., 2017), and the weighted sum mixing is implemented as an extra 1×1 convolution.

4 EXPERIMENTS

In experiments, we test Harmonic Convolution under the deep prior modeling setup introduced by Lempitsky et al. (2018), where the networks are asked to fit a corrupted signal. We define a network’s ability of modeling audio priors as the quality of the audio it produced during the fitting process. Under this definition, we show that networks equipped with Harmonic Convolution can model audio priors better than various baselines. As a by-product, Harmonic Convolution performs comparably well with several state-of-art methods on unsupervised audio restoration. Finally, we demonstrate that Harmonic Convolution improves the generalization performance for supervised sound separation.

4.1 EXPERIMENT SETUPS

We use the LJ-Speech (Ito, 2017) dataset and the MUSIC (Zhao et al., 2018) dataset. LJ-Speech is a speech dataset consisting of short audio clips of a single speaker reading passages. MUSIC is a video dataset of musical instrument solos crawled from Youtube. We only use their audio tracks for all the experiments.

For fair comparisons, we use the same UNet (Ronneberger et al., 2015) architecture with different operations, i.e., regular convolutions, dilated convolutions and the Harmonic Convolution. The details of our network architecture can be found in the appendix. We train all the networks using the Adam optimizer Kingma & Ba (2014) with a learning rate of 0.001 for all the experiments.

4.2 EXPERIMENTS ON DEEP AUDIO PRIORS

Following Lempitsky et al. (2018), we test networks’ ability of modeling audio priors by fitting a corrupted audio signal with a fixed random input and randomly initialized weights. If the network can produce a restored signal faster and with better quality, then we call this network having stronger ability to model audio priors.

Setup. The random input is drawn from a standard Gaussian distribution, and the weights are initialized by drawing from a zero-mean Gaussian distribution with a standard deviation of 0.02. We test networks that rely on 2D spectral-temporal operations, i.e. regular convolutions, dilated convolutions and the Harmonic Convolution, to fit the complex spectrogram of the corrupted signal. The corrupted signal is generated by adding a clean signal with a zero-mean Gaussian noise with a standard deviation of 0.1. The spectrogram is generated by taking the Short-Time-Fourier-Transform (STFT) of the corrupted signal, with a box filter of length 1,022. We use a hop length of 64 to provide enough overlapping, reducing the noise introduced by taking the inverse STFT transformation of the fitted result. When comparing against Wave-U-Net, we use the publicly available implementations provided by the authors of Michelashvili & Wolf (2019).

Results. We compare the fitting progress of the networks with Harmonic Convolution against various baselines, where the quality of the output signal is measured by Peak Signal-to-Noise Ratio on temporal domain at each step. As demonstrated in Figure 4, Harmonic Convolution produces a cleaner signal faster than other methods. The example for fitting is randomly sampled from the LJ-Speech dataset. More examples can be found at <https://anyms-sbms.github.io>.

In addition, we also conduct an ablation study using this fitting process, showing that the design of anchoring and mixing helps with modeling audio priors. As shown in Figure 4 (d) and (f), the with out the anchoring operation, the fitting speed is much slower. Using anchors without the mixing operation would lead to sub-optimal fitting results.

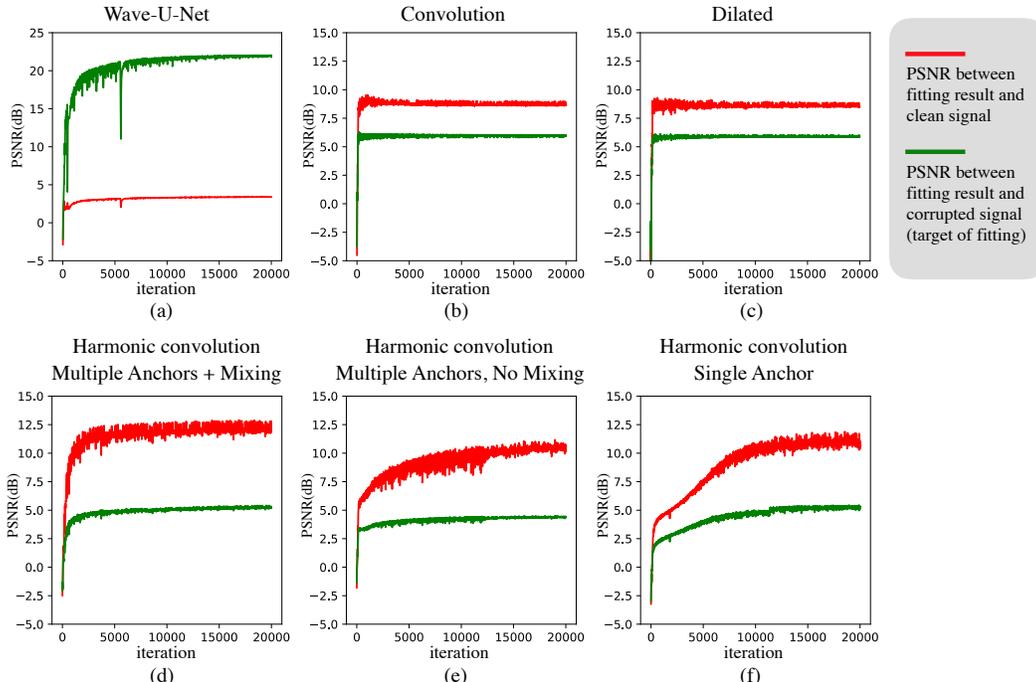


Figure 4: Experiments on deep audio prior using harmonic convolution, regular convolution, dilated convolution, and WaveUNet. Two PSNR scores are calculated at each point: comparing against the input noisy signal, measuring the fitting progress, and comparing against the ground truth clean signal, comparing the fitting quality. (a)(b)(c)(d) under the same setup, Harmonic Convolution is capable to produce significantly higher quality samples compared with (a),(b) and (c). (d)(e)(f) Ablation studies for anchoring and mixing. Without anchoring, the fitting speed is slower. With anchoring but without mixing, the final fitting quality is lower.

Methods	Speech				Music	
	CSIG	CBAK	COVL	PESQ	SSNR	SSNR
Wiener	1.00	1.37	1.00	1.08	0.01	0.68
Wavelet	1.00	1.92	1.01	1.16	2.06	6.50
DFL	1.00	1.51	1.00	1.03	-0.42	-
DNP+LSA	1.00	1.42	1.00	1.02	-3.73	4.26
DNP+Wiener	1.00	1.41	1.00	1.02	-3.33	4.74
Wave-U-Net	1.00	1.36	1.00	1.02	-4.62	3.99
Regular	1.17	2.15	1.10	1.09	4.13	5.07
Dilated	1.29	2.22	1.17	1.13	4.85	5.38
Harmonic	1.76	2.36	1.43	1.20	7.12	9.85

Methods	Speech					
	CSIG	CBAK	COVL	PESQ	SSNR	Human
Conv.	1.00	1.03	1.00	1.09	2.78	0.095
Dilated	1.00	1.02	1.00	1.10	2.38	0.115
Harmonic	1.00	1.01	1.00	1.09	2.05	0.79

Table 2: Quantitative results on the quantization audio restoration task.

Table 1: Quantitative results on the speech and music restoration tasks.

4.3 AUDIO RESTORATION

Similar to the experiments in Lempitsky et al. (2018), networks that models signal priors can be used to perform unsupervised restoration tasks. Here we examine the performance of the network with the Harmonic Convolution on restoring corrupted speech and music audios from random Gaussian noises and aggressive quantizations.

Setup. We conducted experiments on both LJ-Speech and MUSIC datasets. For each dataset, we randomly sample 200 audio clips and clip them to 3 seconds long for restoration tasks. For restoring from Gaussian noise, we set the noise to be zero-mean, with a standard deviation of 0.1. For recovering from quantization noise, we use 1s clips of the 200 randomly sampled speech signals and quantize them into 16 bins which uniformly covers the range from -1 to 1 .

Baselines. We compare with the following baselines:

- **Wiener:** Wiener filtering (Scalart et al., 1996) is an optimization-based methods utilizing a Signal-to-Noise Ratio (SNR) prior. We adopt the implementation where the SNR is estimated from the first 1024 samples of the signal.
- **Wavelet:** We use the MATLAB implementation for wavelet denoising. Wavelet denoising is based on the sparse prior of audio signals, which assumes the wavelet coefficients should be sparse for clean signals. We use the 8-tap symlet wavelets for this task.
- **DFL:** Deep feature losses (Germain et al., 2018) is a state-of-the-art supervised speech denoising approach using perceptual loss tailored for speech signals. When testing on LJ-Speech with Gaussian noise, we are testing its generalization performances under unseen settings.
- **DNP:** Deep network priors (Michelashvili & Wolf, 2019) is an unsupervised method for audio denoising using deep priors. Contrary to our method, the authors observed that during the fitting process, the injected noise varies more violently than the signal itself. Therefore, this property can be used to identify noise regions and provide an SNR estimate for traditional filtering methods such as LSA (Ephraim & Malah, 1985) or Wiener filtering (Scalart et al., 1996).

Metrics. For the speech restoration task, we adopt multiple quality metrics to measure the audio restoration results, including the mean opinion score (MOS) predictor of signal distortion (CSIG), the MOS predictor of background-noise intrusiveness (CBAK), the MOS predictor of overall signal quality (COVL), the perceptual evaluation of speech quality (PESQ), and the segmental Signal-to-Noise Ratio (SSNR). For the music restoration task, we only report results measured by the SSNR, since other metrics do not apply to the music signals. CSIG, CBAK, COVL and PESQ are on the scale of 1 to 5, where 5 being the best quality.

Results. Table 1 summarizes our results against the mentioned baselines over these metrics. Our method consistently outperforms all the baselines according to all measures by a considerable margin. This directly demonstrates that Harmonic Convolution can make neural networks more suited for modeling audio signals.

We also conduct an experiment to verify that Harmonic Convolution is not limited to the additive Gaussian noise case. In this experiment, we quantize 200 randomly sampled one-second speech signals into 16 bins, uniformly covering the range of $[-1, 1]$. The results are reported in Table 2. Since the scores for each metric are rather close, we conduct a carefully designed perceptual experiment with human listeners. For each model, we take 200 audio restorations from the same set for evaluation and each audio is evaluated by three independent Amazon Mechanical Turk (AMT) workers. To avoid cheating, we present the model results with random order, and ask the annotator to select the audio clip that has the best restoration results. The results show nearly 80% of the subjects vote for the results produced by the network using Harmonic Convolution.

4.4 GENERALIZED SOUND SEPARATION

Here we examine whether the Harmonic Convolution can improve generalizations of the supervised sound separation task, compared with regular and dilated convolutions.

Setup. To evaluate the generalization ability of sound separation networks, we select 5 musical instruments from the MUSIC dataset (Zhao et al., 2018): violin, cello, congas, erhu, and xylophone. Each category consists of 50 six-second solo audio clips. We aim to test the model’s ability to generalize to unseen music instruments mixtures. Specifically, the model is tasked with separating out the sound of a target instrument from a clip that also contains sounds of another instrument. During training, we avoid using clips that contain sounds of a selected holdout instrument class, so that the model has never ‘heard’ of the sound of that type of instrument before. We then test the separation performance of this model on mixtures made from sounds of the model’s target instrument and the holdout instrument. In particular, we train models to separate the sound of violins, and use three different hold out instruments to test its generalizability (congas, violin, and xylophone).

Implementations. We adopt the Mix-and-Separate framework (Zhao et al., 2018) for this task. We first generate a synthetic separation training set by mixing the audio signals from two different audio clips, and then trains a neural network to separate the sound of the target instrument, i.e. violin.

During training, we take a 1-second mixed audio clip as input, and transform it into spectrogram using the Short Time Fourier Transform (STFT) with a frame size of 1,022 and a hop size of 64. The spectrogram is then fed into a U-Net, whose architecture is described in Sec 4.1. The U-Net outputs a ratio mask, which is calculated as the ratio of the spectrogram between the sound of the

Methods	Guitar			Xylophone			Congas		
	SDR	SIR	SAR	SDR	SIR	SAR	SDR	SIR	SAR
Conv.	13.3	6.1	4.6	14.0	10.3	7.8	12.7	6.6	4.9
Dilated	14.6	7.0	5.7	13.8	12.6	9.2	12.8	7.8	6.0
Harmonic	15.1	7.9	6.7	13.7	14.0	9.9	13.2	7.6	6.1

Table 3: Quantitative results on the generalized sound separation task. The units are in dB.

target instrument and the input mixed sound. We use an L1 loss for training. To obtain the final separated audio waveform, an inverse STFT is applied using the previous STFT parameters. We use a 90:10 train-val split, and test the performance on the mixture between sounds of the target instrument and the sounds of the holdout instrument.

Results. We compare the performance of the proposed Harmonic Convolution against the regular and dilated convolutions used in previous works. We use the Signal-to-Distortion Ratio (SDR), Signal-to-Interference Ratio (SIR), and Signal-to-Artifact Ratio (SAR). metrics from the open-source `mir_eval` library Raffel et al. (2014) to quantify performances. Quantitative results are reported in Table 3. We observe that while all networks suffer when tested on mixtures under novel recording conditions, the Harmonic Convolution exhibits better generalization performances. This suggests that Harmonic Convolution not only can be used as a prior for unsupervised tasks, but also has the potential to be helpful for supervised tasks.

5 RELATED WORK

Deep priors. Our work is inspired by the recent paper on deep image priors (Lempitsky et al., 2018), which shows that the structure of CNNs imposes a strong prior to restore a single original image from the degraded image. The idea of deep priors has also shown to be useful in many applications, including semantic photo manipulation (Bau et al., 2019), image super-resolution (Shocher et al., 2018), and image decomposition (Gandelsman et al., 2018). While most prior papers focused on images, little work has explored deep priors on audio signals.

Deep learning for auditory signal modeling. Deep networks have gained remarkable success on the audio signal modeling, such as speech recognition (Hinton et al., 2012; Amodei et al., 2016), sound separation (Stoller et al., 2018; Zhao et al., 2018), audio denoising (Rethage et al., 2018; Germain et al., 2018), audio generation (Oord et al., 2017; Mehri et al., 2016), text to speech synthesis (Wang et al., 2017; Shen et al., 2018), and voice conversion (Hsu et al., 2017). A detailed survey can be found at (Purwins et al., 2019; Qian et al., 2019). However, it remains unclear if these architectures themselves capture the audio signal priors. Most related to our work is Michelashvili & Wolf (2019), where they used deep networks as a prior to estimate the SNR prior on the spectrogram and then used classical post-processing algorithms to perform the speech denoising. Note that the deep prior mentioned in works of Mehri et al. (2016) is different from the ones mentioned in Lempitsky et al. (2018). The former uses fitting time variances as noise indicating priors. We aim to find designs that bias towards clean audio signal.

Psychoacoustics. Harmonic structures are closely related to human perception of audio signals. The famous missing fundamental auditory illusion suggests that human can infer the missing fundamental frequency by only hearing its overtones (Todd & Loy, 1991). Moore et al. (1986) showed that shifts in harmonic components would be perceived as separate tones. More recently, Popham et al. (2018) showed that the harmonic structure plays an important role for human to solve the cock-tail problem, where inharmonicity would cause difficulties for human to track speakers for the cock-tail party problem. McPherson & McDermott (2018) showed that pitch perception is closely related to the harmonicity of the sound.

6 CONCLUSION

In this paper, we examined various architectures on deep audio prior modeling. We then proposed a novel operation called Harmonic Convolution, which can help networks better capture priors in audio signals. We showed that fitting a randomly-initialized network equipped with Harmonic Convolution is able to achieve high performance for unsupervised audio restoration tasks. We also showed Harmonic Convolution improves the generalization ability in sound separation.

REFERENCES

- Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al. Deep speech 2: End-to-end speech recognition in english and mandarin. In *ICML*, pp. 173–182, 2016. 8
- Yusuf Aytar, Carl Vondrick, and Antonio Torralba. Soundnet: Learning sound representations from unlabeled video. In *NIPS*, pp. 892–900, 2016. 1, 3
- David Bau, Hendrik Strobelt, William Peebles, Jonas Wulff, Bolei Zhou, Jun-Yan Zhu, and Antonio Torralba. Semantic photo manipulation with a generative image prior. *ACM Transactions on Graphics (TOG)*, 38(4):59, 2019. 8
- Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. *ICCV*, Oct 2017. doi: 10.1109/iccv.2017.89. URL <http://dx.doi.org/10.1109/ICCV.2017.89>. 5
- Yariv Ephraim and David Malah. Speech enhancement using a minimum mean-square error log-spectral amplitude estimator. *IEEE transactions on acoustics, speech, and signal processing*, 33(2):443–445, 1985. 7
- Yossi Gandelsman, Assaf Shocher, and Michal Irani. ”double-dip”: Unsupervised image decomposition via coupled deep-image-priors, 2018. 8
- Francois G. Germain, Qifeng Chen, and Vladlen Koltun. Speech Denoising with Deep Feature Losses. *arXiv e-prints*, Jun 2018. 7, 8
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, pp. 2672–2680, 2014. 1
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pp. 770–778, 2016. 1
- Geoffrey Hinton, Li Deng, Dong Yu, George Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Brian Kingsbury, et al. Deep neural networks for acoustic modeling in speech recognition. *IEEE Signal processing magazine*, 29, 2012. 1, 8
- Chin-Cheng Hsu, Hsin-Te Hwang, Yi-Chiao Wu, Yu Tsao, and Hsin-Min Wang. Voice conversion from unaligned corpora using variational autoencoding wasserstein generative adversarial networks. *arXiv preprint arXiv:1704.00849*, 2017. 8
- Keith Ito. The lj speech dataset. <https://keithito.com/LJ-Speech-Dataset/>, 2017. 3, 5
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pp. 1097–1105, 2012. 11
- Victor Lempitsky, Andrea Vedaldi, and Dmitry Ulyanov. Deep image prior. *CVPR*, Jun 2018. 1, 2, 3, 5, 6, 8
- Malinda J McPherson and Josh H McDermott. Diversity in pitch perception revealed by task dependence. *Nature human behaviour*, 2(1):52, 2018. 8
- Soroush Mehri, Kundan Kumar, Ishaan Gulrajani, Rithesh Kumar, Shubham Jain, Jose Sotelo, Aaron Courville, and Yoshua Bengio. Samplernn: An unconditional end-to-end neural audio generation model, 2016. 8
- Michael Michelashvili and Lior Wolf. Audio denoising with deep network priors, 2019. 1, 3, 5, 7, 8

- Brian CJ Moore, Brian R Glasberg, and Robert W Peters. Thresholds for hearing mistuned partials as separate tones in harmonic complexes. *The Journal of the Acoustical Society of America*, 80(2):479–483, 1986. 2, 8
- Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *ICLR*, 2017. 1, 8
- Sara Popham, Dana Boebinger, Dan PW Ellis, Hideki Kawahara, and Josh H McDermott. Inharmonic speech reveals the role of harmonicity in the cocktail party problem. *Nature communications*, 9(1):2122, 2018. 2, 8
- Hendrik Purwins, Bo Li, Tuomas Virtanen, Jan Schlüter, Shuo-Yiin Chang, and Tara Sainath. Deep learning for audio signal processing. *IEEE Journal of Selected Topics in Signal Processing*, 13(2):206–219, 2019. 8
- Kaizhi Qian, Yang Zhang, Shiyu Chang, Xuesong Yang, and Mark Hasegawa-Johnson. Zero-shot voice style transfer with only autoencoder loss. *ICML*, 2019. 8
- Colin Raffel, Brian McFee, Eric J Humphrey, Justin Salamon, Oriol Nieto, Dawen Liang, Daniel PW Ellis, and C Colin Raffel. mir_eval: A transparent implementation of common mir metrics. In *Proceedings of the 15th International Society for Music Information Retrieval Conference, ISMIR*, 2014. 8
- Dario Reithage, Jordi Pons, and Xavier Serra. A wavenet for speech denoising. In *ICASSP*, pp. 5069–5073, 2018. 8
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention MICCAI 2015*, pp. 234241, 2015. ISSN 1611-3349. doi: 10.1007/978-3-319-24574-4_28. URL http://dx.doi.org/10.1007/978-3-319-24574-4_28. 5
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 3
- Pascal Scalart et al. Speech enhancement based on a priori signal to noise estimation. In *ICASSP*, volume 2, pp. 629–632, 1996. 7
- Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerrv-Ryan, et al. Natural tts synthesis by conditioning wavenet on mel spectrogram predictions. In *ICASSP*, pp. 4779–4783, 2018. 1, 3, 8
- Assaf Shocher, Nadav Cohen, and Michal Irani. “Zero-Shot super-resolution using deep internal learning. In *CVPR*, pp. 3118–3126, 2018. 8
- Daniel Stoller, Sebastian Ewert, and Simon Dixon. Wave-u-net: A multi-scale neural network for end-to-end audio source separation. *arXiv preprint arXiv:1806.03185*, 2018. 1, 3, 8
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *NIPS*, pp. 3104–3112, 2014. 1
- Peter M Todd and D Gareth Loy. *Music and connectionism*. Mit Press, 1991. 8
- Antonio Torralba and Aude Oliva. Statistics of natural image categories. *Network: computation in neural systems*, 14(3):391–412, 2003. 3
- Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016. 11
- Yuxuan Wang, RJ Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, et al. Tacotron: Towards end-to-end speech synthesis. *arXiv preprint arXiv:1703.10135*, 2017. 8
- Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, and Antonio Torralba. The sound of pixels. In *ECCV*, pp. 570–586, 2018. 1, 3, 5, 7, 8

A NETWORK ARCHITECTURE

The Unet used by all methods consists of 5 blocks, each block contains two operation layers, which can be instantiated by regular convolutions, dilated convolutions or the Harmonic Convolution. The feature map is downsampled by average pooling with a stride of 2 and a window size of 2. Downsampling is performed after the first two operation layers. As in common designs of Unet, the upsampling is performed before the last two layers through bilinear upsampling with a scale of 2. Finally, we attach a 1×1 convolution layer after the UNet’s last layer to give the final output. Feature map sizes for all the layers are [input→input, input→35], [35→35, 35→70], [70→70, 70→70], [140→140, 140→35], [70→70, 70→35], then followed by a 1×1 regular convolution layer mapping the final 35 channels to the desired number of output channels.

In addition, we also keep the kernel sizes fixed for all different operations. For regular convolutions, we use 7×7 kernels for all the layers. For dilated convolutions, we use the same kernel size (7×7) with dilation of 3 for all the layers. For the Harmonic Convolution, we use a frequency kernel K_f of length 7 and a temporal kernel K_t also of length 7. We use 7 anchors ($N=7$ in Equation 5) for all the Harmonic Convolution operations. We use instance normalization (Ulyanov et al., 2016) and ReLU activations (Krizhevsky et al., 2012) for all the experiments.