


Hypernetworks build Implicit Neural Representations of Sounds

Filip Szatkowski^{*1,3}[0000-0001-8592-2001] 
Karol J. Piczak^{*2}[0000-0002-6115-0833],
Przemysław Spurek²[0000-0003-0097-5521],
Jacek Tabor^{2,5}[0000-0001-6652-7727], and
Tomasz Trzcíński^{1,2,3,4}[0000-0002-1486-8906]

¹ Warsaw University of Technology, Poland

{filip.szatkowski.dokt,tomasz.trzcinski}@pw.edu.pl

² Faculty of Mathematics and Computer Science, Jagiellonian University, Poland

{karol.piczak,przemyslaw.spurek,jacek.tabor}@uj.edu.pl

³ IDEAS NCBR

⁴ Tooploox

⁵ UES Ltd.

Abstract. Implicit Neural Representations (INRs) are nowadays used to represent multimedia signals across various real-life applications, including image super-resolution, image compression, or 3D rendering. Existing methods that leverage INRs are predominantly focused on visual data, as their application to other modalities, such as audio, is nontrivial due to the inductive biases present in architectural attributes of image-based INR models. To address this limitation, we introduce HyperSound, the first meta-learning approach to produce INRs for audio samples that leverages hypernetworks to generalize beyond samples observed in training. Our approach reconstructs audio samples with quality comparable to other state-of-the-art models and provides a viable alternative to contemporary sound representations used in deep neural networks for audio processing, such as spectrograms. Our code is publicly available at <https://github.com/WUT-AI/hypersound>.

Keywords: Machine Learning · Hypernetworks · Audio Processing · Implicit Neural Representations

1 Introduction

The field of Implicit Neural Representations (INRs) is a rapidly growing area of research, which aims to obtain functional, coordinate-based representations of multimedia signals through the use of neural networks. These representations are decoupled from the spatial resolution, so the signal can be resampled at any arbitrary frequency without incurring additional memory storage requirements. Applications of INRs include super-resolution [30, 22], compression [30, 22], 3D

* equal contribution

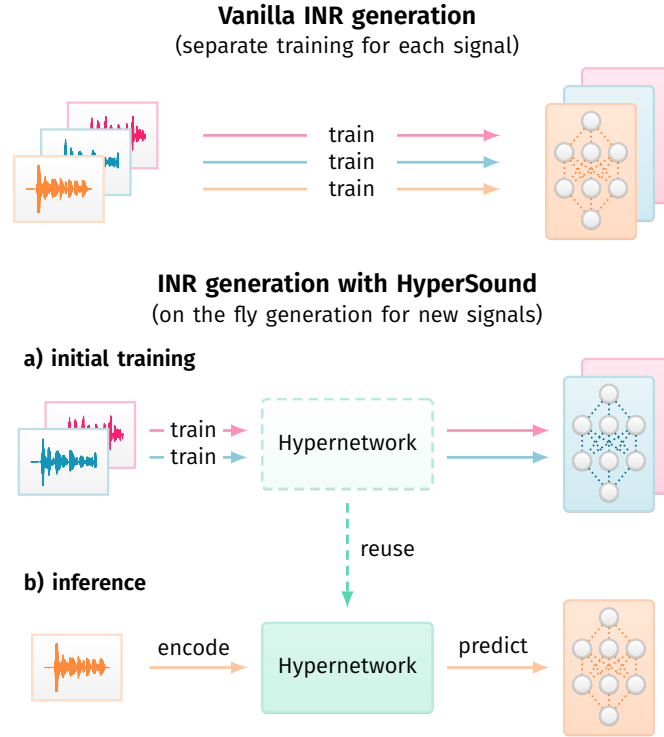


Fig. 1. We propose HyperSound, the first framework which produces Implicit Neural Representations for unseen audio signals, eliminating the need for retraining a model for each new example.

rendering [23], and missing data imputation [10]. Furthermore, INRs can be used to synthesise data [44], and they show promise as an alternative to the traditional feature extraction techniques used in neural networks [8].

INRs are a promising research direction in audio processing, yet so far their applications in this domain remain limited. As the size of INR networks can be significantly smaller than the size of the represented signal, they can be used as a compression method [34]. They also provide an alternative to super-resolution algorithms such as ProGAN [15], which often require hand-crafting complex architectures for selected resolution. In contrast, in INRs the signal resolution is given by the granularity of input coordinates, so a single network represents the signal for any resolution we might require while remaining constant in size. Moreover, as [24] show that INRs can be used as features in other algorithms, this representation can be an alternative for current deep audio processing models that primarily operate on features obtained from mel spectrograms or raw audio waves [27].

In practice, architectures typically used for INRs struggle with modeling high frequencies inherent in audio signals [28] and require training a new model for every

incoming data point, which is wasteful and greatly reduces their applicability. This issue is usually addressed by obtaining generalizable INRs through meta-learning methods, such as hypernetworks [12] – models that produce the weights for other models. The hypernetwork framework can be leveraged to obtain generalizable INRs for a variety of data types, including images [30, 22, 16], point clouds [22, 33, 36] 3D objects [14, 43], videos [22] and time series [10]. However, obtaining generalizable INRs for audio signals remains an unsolved task due to the high dimensionality and variance of this type of data, which makes the standard methods for this task insufficient.

In our work, we propose a method that learns a general recipe for creating INRs for arbitrary audio samples not present in the training dataset. We adapt the hypernetwork framework for sound processing by introducing a domain-specific loss function and hypernetwork weight initialization. Additionally, we investigate various variants of the hypernetwork framework in order to achieve improved compression while maintaining audio quality. To our knowledge, our method, named HyperSound, is the first application of hypernetwork-based INR generation to the audio domain.

The main contributions of our work are:

- We introduce HyperSound, to our knowledge the first INR method for audio processing that builds functional representations of unseen audio signals, which eliminates the need of retraining the INR for each new data point.
- We show a successful adaptation of the hypernetwork paradigm to the audio processing through domain-specific modifications such as the spectral loss function.
- We empirically evaluate various approaches of using hypernetworks to build INRs for audio generation, e.g. by leveraging SIREN architecture or Fourier mapping.

2 Related Works

2.1 Implicit Neural Representations (INRs)

Implicit Neural Representations (INRs) are continuous, coordinate-based approximations of data obtained through neural networks. The decoupling of representation from spatial resolution makes INRs attractive for a wide range of applications in signal processing. However, functions represented by standard MLP networks with ReLU activation function show a strong bias towards low frequencies [28], which in practice make them infeasible as representations of multimedia signals. To mitigate this bias, most INRs use either Fourier features [37] or SIREN networks [30] with periodic activation functions. Later works improve SIREN networks, introducing various forms of modulation to sinusoidal layers [22, 3]. INRs are currently used in many applications, such as super-resolution [30, 22], compression [34], 3D rendering [23], missing data imputation [10] or data generation [44]. However, typical INRs are trained for each data point, and obtaining a representation for new signals requires retraining the network from scratch.




Fig. 2. Overview of the HyperSound framework. We use a single hypernetwork model to produce distinct INRs based on arbitrary audio signals provided as input.

2.2 Generalizable INRs

To obtain a way of generating INRs for new data without the need for model retraining, various meta-learning approaches were proposed. [31] leverage gradient-based meta-learning algorithms to solve the task of learning an INR space for shapes. [5] employ a transformer meta-learner to obtain INRs for unseen examples. Many methods of obtaining generalizable INRs also involve the use of hypernetworks [12], where one network (hypernetwork) generates the weights for another network (target network). More broadly, hypernetworks are employed in many applications, including model compression [42], continual learning [39] or generating INRs. In particular, hypernetworks were successfully used to generate INRs for images [16, 30, 22], shapes [33, 30], videos [22] and time series [10]. To our knowledge, our work is the first application of hypernetworks for generating audio INRs.

2.3 Deep neural networks for audio processing

First successful attempts at raw waveform processing with deep neural networks were models such as WaveNet [26] and SampleRNN [21], but their autoregressive nature makes them slow and prone to accumulation of errors. Later architectures such as ParallelWaveNet [25], NSynth [9], MelGAN [17] or SING [7] proposed non-autoregressive architectures for audio generation. Recent autoencoder-based models such as RAVE [2] or SoundStream [41] are able to process high-resolution signals in an end-to-end fashion, producing audio of very good perceptual quality.

3 Method

Traditional digital representation of sound waves is a series of amplitude values sampled at regular time intervals, which approximates a continuous real function $x(t)$. Our goal is to obtain a meta-recipe for generating audio INRs approximating such functions. While the creation of INRs for particular audio samples can be successfully achieved with gradient descent, finding a general solution is much harder due to the inherent complexity of audio time series. Therefore, similar to [16], we model such functions with neural networks T (target networks),

parameterized by weights generated by another neural network H (hypernetwork). Our method, named HyperSound, can be described as

$$\begin{aligned}\boldsymbol{\theta}_x &= H(\mathbf{x}), \\ \hat{x}(t) &= T(t, \boldsymbol{\theta}_x).\end{aligned}\tag{1}$$

We show the overview of our method in Fig. 2. In the following sections, we describe the architecture of the models used as Implicit Neural Representations and of the hypernetwork. Then, we present the details required for the successful application of the hypernetwork framework to the audio domain.

3.1 Implicit Neural Representations

We consider two variants of INR networks, inspired by NeRF [23] and SIREN [30] papers. Both models have a single input and a single output.

We refer to a NeRF-inspired network as a Fourier-mapped multilayer perceptron (FMLP). The network consists of a positional encoding layer followed by dense layers with biases and ReLU activation functions. Positional embedding vectors $\boldsymbol{\gamma}$ are obtained as

$$\boldsymbol{\gamma}(t) = \left[\sin(2^0 \pi t), \cos(2^0 \pi t), \dots, \sin(2^{L-1} \pi t), \cos(2^{L-1} \pi t) \right]\tag{2}$$

where t is the time coordinate and L denotes the embedding size. We normalize input coordinates to a range of $[-1, 1]$ and use $L = 10$.

SIREN network consists of dense layers with a sinusoidal activation function. The output \mathbf{y}_i of i -th SIREN layer is defined as:

$$\mathbf{y}_i = \sin(\omega_i \mathbf{x}_i \mathbf{W}_i^T + \mathbf{b}_i),\tag{3}$$

where x_i is the input of the layer, \mathbf{W}_i and \mathbf{b}_i are layer weights and biases and ω_i is a scalar factor introduced in [30] to enable learning of higher signal frequencies. In practice, following the original paper, we set a distinct $\omega_0 \neq 1$ for the first layer and keep the same shared value of ω_i for all remaining layers.

[1] prove that FMLP networks are equal to single-layer SIREN networks. However, in our experiments, we see significant differences in the learning dynamics and representational capacity of both networks and thus explore both architectures in our experiments.

3.2 Hypernetwork

Typical audio recordings contain several thousands of samples, while hypernetworks generally use fully connected layers to produce target weights. Since we want to use raw audio recordings as input, using dense layers to process the signal directly leads to the explosion of the number of weights in the hypernetwork. To avoid this problem, we use a convolutional encoder to produce a latent, lower dimensional representation of the signal and then process this representation with

fully connected layers. The hypernetwork outputs weights θ that parameterise the target network.

We use an encoder based on SoundStream [41]. The fully connected part of the hypernetwork is composed of six layers with biases and ELU [6] activation, where the last layer produces the flattened weights of the target network. Input of the hypernetwork is 32 768 samples of audio at a sampling rate of 22 050 Hz, the size of a latent representation is 103×32 and the layer sizes of the dense part are 400, 768, 768, 768, 768, 400.

3.3 Hypernetwork adaptation for audio data

Audio signals contain inherent high-frequency content, and neural networks in general tend to focus on low-frequency features [28]. Using a hypernetwork to produce INR weights introduces additional complexity to the optimization process. During training, small updates to the hypernetwork weights might disproportionately affect the function represented by the output weights, which makes the training process unstable. Moreover, the outputs of the neural network will always contain some inherent noise, which affects the high-frequency content more. Finally, in the case of training INRs, hypernetworks tend to overfit the training data, as they usually require over-parametrization to be able to solve a difficult problem of function generation. Combined, those issues make the straightforward application of the hypernetwork framework to our problem impossible.

To enable hypernetworks to model audio data, we introduce STFT loss function, which counters the hypernetworks tendency to ignore high-frequencies. The optimization process is described in detail in Section 3.4. The spectral loss component is crucial to obtain the INRs of sufficient quality, as using only standard reconstruction loss in the time domain leads to the network collapsing to trivial solutions such as uniform noise or silence. Additionally, we notice that our method tends to overfit the training dataset and find that in order to make hypernetwork generalize to the unseen data we have to introduce random data augmentations to the training process.

3.4 Optimization for audio reconstruction

We train the hypernetwork in a supervised fashion using backpropagation. To obtain audio results that are more perceptually pleasant, we use a loss function that penalizes the reconstruction error both in time and frequency domains. Given an original recording \mathbf{x} and its reconstruction $\hat{\mathbf{x}}$ generated with a target network, we compute the loss function as:

$$L(\mathbf{x}, \hat{\mathbf{x}}) = \lambda_t L_t(\mathbf{x}, \hat{\mathbf{x}}) + \lambda_f L_f(\mathbf{x}, \hat{\mathbf{x}}), \quad (4)$$

where L_t is the L1 loss in time domain and L_f is a multi-resolution mel-scale STFT loss in frequency domain as introduced in ParallelWaveGAN [40]. We use coefficients λ_t and λ_f to control the tradeoff between both losses.

Table 1. Comparison of different variants of computing target network pre-activation layer output.

Network type	Instance-specific weights	Shared weights	Pre-activation layer output
standard	all layers	none	$\mathbf{y}_i = \mathbf{x}(\mathbf{W}_i^h)^T + \mathbf{b}_i^h$
additive	all layers	all layers	$\mathbf{y}_i = \mathbf{x}(\mathbf{W}_i^s + \mathbf{W}_i^h)^T + \mathbf{b}_i^s + \mathbf{b}_i^h$
multiplicative	all layers	all layers	$\mathbf{y}_i = \mathbf{x}(\mathbf{W}_i^s \cdot \mathbf{W}_i^h)^T + \mathbf{b}_i^s \cdot \mathbf{b}_i^h$
shared	some layers	some layers	$\mathbf{y}_i = \begin{cases} \mathbf{x}(\mathbf{W}_i^s)^T + \mathbf{b}_i^s, & \text{if layer shared} \\ \mathbf{x}(\mathbf{W}_i^h)^T + \mathbf{b}_i^h, & \text{otherwise} \end{cases}$

For the STFT loss, we use 128 mel bins and FFT sizes of 2048, 1024, 512, 256, 128 with matching window sizes and an overlap of 75%. We set $\lambda_t = 1$ and $\lambda_f = 1$. Using both loss functions together allows us to optimize both quantitative metrics through L1 part of the loss and the perceptual quality of generated audio through the STFT part. We base this intuition on results from [10], which show that the FFT loss is crucial for learning INRs for time series.

Neural networks are naturally biased towards low-frequency features, and we notice that the target networks obtained with HyperSound suffer from the same problem. Therefore, in some experiments we modify the STFT loss by frequency-weighting the Fourier transforms of \mathbf{x} and $\hat{\mathbf{x}}$. We compute the weightings w_{f_i} for frequency bins f_1, f_2, \dots, f_N as:

$$w_{f_i} = \frac{N \cdot i^p}{\sum_{j=0}^N j^p}, \quad (5)$$

where N is the number of frequency bins and $p \geq 0$ is a parameter that allows us to control the uniformity of this weighting. The weights are normalized so that they always sum to N , and the minimum value of p yields a uniform weighting (no bias towards high frequencies). During training, we gradually move from uniform weighting to target weighting in 500 epochs.

3.5 Alternative ways to leverage hypernetworks for INR generation

In the standard variant of our framework, the hypernetwork generates weights and biases for target network layers and other parameters, such as frequencies of positional encoding in FMLP or ω_i in SIREN, remain fixed. However, we also consider learning a shared, generic set of weights and biases for target network layers and combining them with instance-specific weights and biases generated by the hypernetwork in other ways. We refer to the three variants of target networks explored in our work as shared, additive and multiplicative. We show a comparison of all the approaches in Table 1.

With a shared variant of the target network, we explicitly learn a shared set of weights \mathbf{W}_i^s and biases \mathbf{b}_i^s for some layers, and use the hypernetwork only for generating instance-specific weights \mathbf{W}_i^h and biases \mathbf{b}_i^h for the remaining layers. Target networks’ outputs are then computed using either shared or instance-specific set of weights, depending on the layer.

In the additive and multiplicative variants of the target network, for each layer i we explicitly learn shared weights \mathbf{W}_i^s and biases \mathbf{b}_i^s . The hypernetwork learns to generate separate, instance-specific, weights \mathbf{W}_i^h and biases \mathbf{b}_i^h . Both sets of weights and biases are then either added or multiplied and the results are used to compute layer output.

3.6 Hypernetwork weight initialization

Hypernetworks are very sensitive to the learning rate and require careful weight initialization to ensure stable learning. Xavier [11] or Kaiming [13] weight initializations typically used for deep neural networks are insufficient for hypernetworks, as we want to achieve the corresponding initial distribution of generated weights, but we can affect it only through the weights of the hypernetwork. To solve this issue, [4] propose hypernetwork weight initialization scheme which leads to the initial target network weights following the distribution given by either Xavier or Kaiming initialization. We adopt this process for our FMLP network, but find it insufficient for SIREN networks, as they require specific initialization described by [30]. Therefore, to simulate SIREN initialization in the initial stage of training, we set the starting weights for the last layer of the hypernetwork according to uniform distribution with boundaries close to zero. and set the distribution of biases to the one proposed in the SIREN paper. This makes the frequencies in the subsequent target network layers grow slowly in the initial stages of network training, which is crucial to stable learning with periodic activation functions.

4 Experiments

4.1 Setting

For most experiments, unless explicitly stated otherwise, we use the VCTK dataset downsampled to $f = 22\,050$ Hz, with recording length set to 32 768 samples, which is around 1.5 seconds. We obtain inputs by cropping the original audio to the desired length. We use the recordings of the last 10 speakers in the VCTK dataset as a validation set. We set batch size to for all the models to 16, and train each model for 2500 epochs with 10 000 training samples per epoch. We use a OneCycleLR [32] optimizer, varying learning rate depending on the target network model type, as we observe that SIREN networks require smaller learning rates to maintain stability during training.

During the evaluation, we focus on the quality of reconstructions obtained with target networks. Example of these reconstructions can be seen in Figure 3. Since there is no consensus on a single approach for the quantitative evaluation




Fig. 3. Example of VCTK validation sample reconstructed with HyperSound.

of audio quality, we assess the reconstruction results with multiple metrics. We use quantitative metrics such as MSE, Log-Spectral Distance (LSD) [18] and SI-SNR [19]. However, since these metrics do not fully capture the perceptual quality of the audio signal, we also include metrics such as PESQ [29], STOI [35] and CDPAM [20] to approximately measure how well the reconstructed signal sounds to the human ear.

Unless explicitly mentioned otherwise, all the network hyperparameters are set as described in Section 3. For training all of the hypernetworks, we use the same data augmentations as in RAVE[2], i.e. random crop, phase mangle, and dequantization.

4.2 Reconstruction quality with FMLP target network

We measure the quality of reconstructions obtained on the VCTK dataset with HyperSound using different variants of the target network. We compare the quality of reconstructions with the state-of-the-art RAVE [2] autoencoder retrained with sampling rate of 22 050 Hz. We use the same set of data augmentations for training both RAVE and HyperSound, and train RAVE for 3×10^6 steps.

We consider multiple variants of target networks that match the desired compression ratio (measured as a ratio between the number of samples in the input recording and the number of target network parameters). We set the hypernetwork learning rate to 10^{-4} and use the AdamW optimizer. All target networks use five hidden dense layers with biases, and we set the width of each hidden layer to an equal value, so that we obtain the desired compression ratio. We explore more realistic scenarios with compression ratios above $1\times$, but also include the results for target networks with higher parameter counts as an insight into the representational capacities of over-parameterized networks.

We show the results for RAVE in Table 2a, and for our method with FMLP target network Table 2b. The results show that, while over-parameterization of the target network is beneficial for the reconstruction quality, the experiments

Table 2. HyperSound reconstruction error on the VCTK validation set compared with other methods. a) We compare our method with the state-of-the-art reconstruction model RAVE. b) We test different variants of the FMLP target network with different sizes as measured by the compression ratio (CR). c) We compare the reconstruction error obtained with SIREN target networks of size equal to FMLP with CR 1 and various values of ω_i and ω_0 . d) We also show the upper bound of INR reconstruction obtained by training a single model for each data point, using both FMLP and SIREN networks.

Model		MSE	LSD	SI-SNR	PESQ	STOI	CDPAM
<i>Ideal metric behavior:</i>		$\rightarrow 0$	$\rightarrow 0$	$\rightarrow 100$	$\rightarrow 4.5$	$\rightarrow 1$	$\rightarrow 0$
a) <i>RAVE</i> [2]		0.049	0.99	-23.50	1.75	0.87	0.35
b) <i>HyperSound FMLP</i>	CR						
	10 \times	0.013	1.80	-0.25	1.14	0.71	0.43
	4 \times	0.011	1.80	1.27	1.17	0.73	0.42
	2 \times	0.011	1.73	2.74	1.21	0.75	0.38
	1 \times	0.009	1.61	2.99	1.27	0.76	0.36
	0.5 \times	0.009	1.53	4.13	1.34	0.78	0.30
	0.125 \times	0.007	1.50	4.48	1.42	0.80	0.30
c) <i>HyperSound SIREN</i>	ω_0	ω_i					
	30	30	0.020	4.14	-4.82	1.18	0.66
	100	30	0.018	3.56	-3.66	1.19	0.65
	500	30	0.033	2.52	-17.59	1.18	0.65
	1000	30	0.036	2.50	-22.63	1.19	0.66
	2000	30	0.037	2.59	-23.93	1.19	0.67
	2000	100	0.036	1.63	-22.11	1.18	0.67
d) <i>Individual INRs</i> (<i>upper bound</i>)	Model						
	FMLP	0.001	1.39	17.19	2.31	0.95	0.23
	SIREN	<0.001	1.40	27.53	3.05	0.97	0.28

with smaller target networks still obtain passable scores across all metrics. We also notice that our method outperforms RAVE in pure reconstruction metrics across all selected target network sizes, but obtains worse scores on metrics that measure spectral distance and perceptual quality.

4.3 Training SIREN target networks

Despite [1] showing a close relationship between SIREN and FMLP INRs, we observe that they behave very differently during training when compared to FMLP networks. First, we notice that without the hand-crafted hypernetwork weight initialization, as described in Section 3.3, SIREN networks fail to learn any meaningful signal. Additionally, we notice that SIRENs require significantly slower learning rates to converge, and we use a learning rate of 10^{-6} ($1/100$ the learning rate we use in FMLP networks). We would like to use high values of

Table 3. Comparison of different methods of generating INRs with hypernetworks described in Section 3.5. We notice that sharing initial layers of the target network is beneficial to reconstruction quality while making the effective size of both hypernetwork and target network smaller.

model type	MSE	LSD	SI-SNR
<i>standard</i>	0.009	1.61	2.99
<i>multiplicative</i>	0.009	1.61	2.77
<i>additive</i>	0.008	1.52	2.98
<i>first layer shared</i>	0.008	1.58	3.67
<i>first two layers shared</i>	0.008	1.60	3.44

ω_0 and ω_i , which improve the network’s capability to represent high frequencies, but increasing those values makes training the network very unstable and often prevents convergence to any meaningful solution altogether.

Table 2c shows the results obtained for SIREN networks with different values of ω_0 and ω_i . We use network architecture corresponding to the FMLP model with a compression ratio of $1\times$. Increasing the frequencies ω_0 and ω_i slightly improves the perceptual quality of the reconstructions, but makes the training process less stable. In further experiments, we focus on FMLP networks, as they produce better results and are easier to train.

4.4 Upper bound of the INR reconstruction quality

We compute our metrics for SIREN and FMLP INRs trained in vanilla fashion to get an upper bound for reconstruction error. We select a subset of the validation data, train a model for each example minimizing the mean squared reconstruction error and report the mean metric values as an upper bound for HyperSound. We only use a subset of the validation dataset, as the mean of the metrics quickly stabilises after several recordings. We use basic variants of the FMLP and SIREN networks ($\omega_0 = 2000$, $\omega_i = 30$), with parameters other than layer weights and biases set at fixed values as described in Section 3.1 and the basic model size as described in previous Sections.

The comparison of our results is described in Table 2d. Unsurprisingly, vanilla INRs achieve significantly better quality compared to our method and RAVE. Contrary to results for HyperSound, SIREN INRs seem to work better than FMLP models when optimized directly for a single signal. We hypothesise that this can be explained by the increased complexity of hypernetwork optimization and ReLU nonlinearity being more robust to weight perturbations.

4.5 Exploring different ways to use weights generated by the hypernetwork

We explore possible extensions of the hypernetwork framework for INR generation, comparing three approaches described in Section 3.5 with the standard variant

Table 4. Impact of the frequency weighting. A higher value of p indicates a stronger bias towards high frequencies in loss function, while $p = 0$ equals no weighting. Our results indicate that focusing on high-frequency content in signals leads to better reconstructions.

p	MSE	LSD	SI-SNR
<i>0.0 [baseline]</i>	0.009	1.61	2.99
<i>0.2</i>	0.008	1.50	3.38
<i>0.5</i>	0.008	1.57	3.09
<i>1.0</i>	0.007	1.48	3.60

Table 5. Comparison of HyperSound reconstruction quality on different datasets. Our method still performs reasonably well on datasets other than VCTK, without any additional hyperparameter tuning.

Dataset	MSE	LSD	SI-SNR
<i>VCTK</i>	0.009	1.61	2.99
<i>LJ Speech</i>	0.014	1.65	-1.08
<i>LibriSpeech</i>	0.011	2.15	0.89

of the target network. We conduct experiments using the network architecture corresponding to the FMLP with compression ratio $1\times$ from Section 4.2.

The results of our experiments are described in Table 4.3. We notice that learning a shared set of weights for the first layer in the target network leads to better results than in the case of the standard model, but sharing more than one layer is not beneficial. Sharing weights of the first layer can be thought of as learning a shared feature extractor for smaller hypernetwork and target network models, which means that effectively we reduce the computational cost of the method and provide more compressed representations. Other variants such as multiplicative and additive do not improve the results noticeably while inducing additional computational cost due to the added complexity.

4.6 Higher-frequency weighting in the loss function

We notice that our solution tends to produce representations that focus on lower frequencies, which lowers the perceptual quality of the reconstructions. We aim to improve this quality by assigning higher weights to the higher frequencies in the STFT loss, as described in Section 3.4. As shown in Table 4, this simple modification improves the reconstruction quality of the model, which indicates that audio processing can greatly benefit from domain-specific loss functions.

4.7 Results on other datasets

In our experiments, we focus on the VCTK dataset, as most publicly available audio datasets are either small, include data that does not match our sampling rate or input length requirements, or only include data from a single source or speaker. However, training hypernetworks requires large and diverse datasets and benefits from data augmentations, which are harder to apply in a reasonable way to music or environmental sound datasets. Therefore, we evaluate HyperSound reconstruction quality on additional speech datasets, i.e., LJ Speech and LibriSpeech *train-clean-360* datasets. We present the results of this evaluation in Table 5, showing that HyperSound still achieves reasonable results despite



Fig. 4. t-SNE visualisation of the INR weights obtained on VCTK validation set. Different colors of the dots indicate different speakers. Speakers are generally clustered, and we notice a clear separation of the male and female recordings.

no additional dataset-specific hyperparameter tuning apart from learning rate reduction for LibriSpeech.

4.8 Representational capability of INRs

Finally, we conduct a simple experiment visualizing the weights of the INRs obtained on the validation set of VCTK with t-SNE [38] and show the results in Figure 4. We group the INRs by speaker identity, as its the best-defined criterion to distinguish between recordings in the dataset. We notice that different speakers occupy distinct regions in the t-SNE subspace, with a clear separation between recordings of male and female speech. This proves that representations produced by HyperSound capture the underlying structure of the data.

5 Conclusions

We propose the first method of obtaining generalizable Implicit Neural Representations for audio signals, named HyperSound. We adapt the hypernetwork framework to audio processing through domain-specific modifications and demonstrate that our model performs reasonably well on the VCTK dataset, producing reconstructions quantitatively comparable to the state-of-the-art RAVE autoencoder. We investigate the use of SIREN and Fourier-mapped perceptron networks as INRs generated by the hypernetwork and try alternative ways of applying hypernetworks in the INR generation, which improve upon the standard method.

INRs produced by our method can represent the signal in compressed way, using less weights than the number of samples in the original recording. Additionally, we demonstrate that our method performs reasonably across different datasets with minimal hyperparameter tuning and that the representations produced by the method align with the underlying structure of the data.

Limitations and future work. As theoretical guarantees of hypernetwork training are hard to define, we find our method sensitive to the hyperparameter choice and its training requires a lot of trial and error. Moreover, the misalignment between the human perception of the sound and metrics optimized in deep neural networks means that the perceptual quality of reconstructions generated with our method leaves room for improvement. We hope that further work in those directions can improve the stability of our method and better calibrate the results to human perception.

Ethical Statement

This paper belongs to the line of fundamental research and all the experiments in this paper were performed using publicly available data. However, we acknowledge that there are potential ethical implications of our work that need to be considered.

As with all machine learning models, biases from the training data can be encoded into the model, leading to inaccurate or discriminatory behavior towards underrepresented groups.

Furthermore, the data reconstructed with Implicit Neural Representations always contains some degree of error, and models trained with different hyperparameters can produce varying representations. These properties make INRs a potential tool to evade copyright detection, and current detection algorithms are not equipped to work on data stored as INR weights, further compounding the issue.

Acknowledgements

This work was supported by Foundation for Polish Science with Grant No POIR.04.04.00-00-14DE/18-00 carried out within the Team-Net program co-financed by the European Union under the European Regional Development Fund, and by the National Centre of Science (Poland) Grant No. 2020/39/B/ST6/01511. Filip Szatkowski and Tomasz Trzcinski are supported by National Centre of Science (Poland) Grant No. 2022/45/B/ST6/02817. Przemysław Spurek is supported by the National Centre of Science (Poland) Grant No. 2021/43/B/ST6/01456. Jacek Tabor research has been supported by a grant from the Priority Research Area DigiWorld under the Strategic Programme Excellence Initiative at Jagiellonian University.

References

- [1] Nuri Benbarka et al. “Seeing Implicit Neural Representations As Fourier Series”. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. Jan. 2022, pp. 2041–2050.
- [2] Antoine Caillon and Philippe Esling. “RAVE: A variational autoencoder for fast and high-quality neural audio synthesis”. In: *arXiv preprint 2111.05011* (2021).
- [3] Eric R Chan et al. “PI-GAN: Periodic implicit generative adversarial networks for 3d-aware image synthesis”. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2021, pp. 5799–5809.
- [4] Oscar Chang, Lampros Flokas, and Hod Lipson. “Principled weight initialization for hypernetworks”. In: *International Conference on Learning Representations*. 2019.
- [5] Yinbo Chen and Xiaolong Wang. “Transformers as Meta-learners for Implicit Neural Representations”. In: *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVII*. Springer. 2022, pp. 170–187.
- [6] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. “Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)”. In: *arXiv preprint arXiv:1511.07289* (2015).
- [7] Alexandre Défossez et al. “SING: Symbol-to-Instrument Neural Generator”. In: *Advances in neural information processing systems* 31 (2018).
- [8] Emilien Dupont et al. “From data to functa: Your data point is a function and you should treat it like one”. In: *arXiv preprint arXiv:2201.12204* (2022).
- [9] Jesse Engel et al. “Neural Audio Synthesis of Musical Notes with WaveNet Autoencoders”. In: *International Conference on Machine Learning*. PMLR. 2017, pp. 1068–1077.
- [10] Elizabeth Fons et al. “HyperTime: Implicit Neural Representation for Time Series”. In: *arXiv preprint arXiv:2208.05836* (2022).
- [11] Xavier Glorot and Yoshua Bengio. “Understanding the difficulty of training deep feedforward neural networks”. In: *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings. 2010, pp. 249–256.
- [12] David Ha, Andrew Dai, and Quoc V Le. “HyperNetworks”. In: *arXiv preprint arXiv:1609.09106* (2016).
- [13] Kaiming He et al. “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification”. In: *Proceedings of the IEEE international conference on computer vision*. 2015, pp. 1026–1034.
- [14] Adam Kania et al. “Hypernerfgan: Hypernetwork approach to 3d nerf gan”. In: *arXiv preprint arXiv:2301.11631* (2023).
- [15] Tero Karras et al. “Progressive Growing of GANs for Improved Quality, Stability, and Variation”. In: *International Conference on Learning Representations*. 2018.

- [16] Sylwester Klocek et al. “Hypernetwork functional image representation”. In: *International Conference on Artificial Neural Networks*. Springer. 2019, pp. 496–510.
- [17] Kundan Kumar et al. “MelGAN: Generative Adversarial Networks for Conditional Waveform Synthesis”. In: *Advances in neural information processing systems* 32 (2019).
- [18] Haohe Liu et al. “Neural Vocoder is All You Need for Speech Super-resolution”. In: *arXiv preprint arXiv:2203.14941* (2022).
- [19] Yi Luo and Nima Mesgarani. “TasNet: time-domain audio separation network for real-time, single-channel speech separation”. In: *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. 2018, pp. 696–700.
- [20] Pranay Manocha et al. “CDPAM: Contrastive learning for perceptual audio similarity”. In: *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. 2021, pp. 196–200.
- [21] Soroush Mehri et al. “SampleRNN: An Unconditional End-to-End Neural Audio Generation Model”. In: *arXiv preprint arXiv:1612.07837* (2016).
- [22] Ishit Mehta et al. “Modulated Periodic Activations for Generalizable Local Functional Representations”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, pp. 14214–14223.
- [23] Ben Mildenhall et al. “NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis”. In: *Communications of the ACM* 65.1 (2021), pp. 99–106.
- [24] Norman Müller et al. “DiffRF: Rendering-Guided 3D Radiance Field Diffusion”. In: *arXiv preprint arXiv:2212.01206* (2022).
- [25] Aaron Oord et al. “Parallel WaveNet: Fast High-Fidelity Speech Synthesis”. In: *International conference on machine learning*. PMLR. 2018, pp. 3918–3926.
- [26] Aaron van den Oord et al. “WaveNet: A Generative Model for Raw Audio”. In: *arXiv preprint arXiv:1609.03499* (2016).
- [27] Hendrik Purwins et al. “Deep learning for audio signal processing”. In: *IEEE Journal of Selected Topics in Signal Processing* 13.2 (2019), pp. 206–219.
- [28] Nasim Rahaman et al. “On the spectral bias of neural networks”. In: *International Conference on Machine Learning*. PMLR. 2019, pp. 5301–5310.
- [29] Antony W Rix et al. “Perceptual evaluation of speech quality (PESQ)—a new method for speech quality assessment of telephone networks and codecs”. In: *2001 IEEE international conference on acoustics, speech, and signal processing. Proceedings (Cat. No. 01CH37221)*. Vol. 2. IEEE. 2001, pp. 749–752.
- [30] Vincent Sitzmann et al. “Implicit Neural Representations with Periodic Activation Functions”. In: *Advances in Neural Information Processing Systems* 33 (2020), pp. 7462–7473.

- [31] Vincent Sitzmann et al. “Metasdf: Meta-learning signed distance functions”. In: *Advances in Neural Information Processing Systems* 33 (2020), pp. 10136–10147.
- [32] Leslie N Smith and Nicholay Topin. “Super-convergence: Very fast training of neural networks using large learning rates”. In: *Artificial intelligence and machine learning for multi-domain operations applications*. Vol. 11006. SPIE. 2019, pp. 369–386.
- [33] Przemysław Spurek et al. “Hypernetwork approach to generating point clouds”. In: *Proceedings of Machine Learning Research* 119 (2020).
- [34] Yannick Strümpfer et al. “Implicit neural representations for image compression”. In: *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVI*. Springer. 2022, pp. 74–91.
- [35] Cees H Taal et al. “A short-time objective intelligibility measure for time-frequency weighted noisy speech”. In: *2010 IEEE international conference on acoustics, speech and signal processing*. IEEE. 2010, pp. 4214–4217.
- [36] J Tabor, T Trzciński, et al. “Hyperpocket: Generative point cloud completion”. In: *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2022, pp. 6848–6853.
- [37] Matthew Tancik et al. “Fourier features let networks learn high frequency functions in low dimensional domains”. In: *Advances in Neural Information Processing Systems* 33 (2020), pp. 7537–7547.
- [38] Laurens Van der Maaten and Geoffrey Hinton. “Visualizing data using t-SNE.” In: *Journal of machine learning research* 9.11 (2008).
- [39] Johannes Von Oswald et al. “Continual learning with hypernetworks”. In: *arXiv preprint arXiv:1906.00695* (2019).
- [40] Ryuichi Yamamoto, Eunwoo Song, and Jae-Min Kim. “Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram”. In: *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. 2020, pp. 6199–6203.
- [41] Neil Zeghidour et al. “SoundStream: An End-to-End Neural Audio Codec”. In: *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 30 (2021), pp. 495–507.
- [42] Dominic Zhao et al. “Meta-Learning via Hypernetworks”. In: (2020).
- [43] D Zimny, T Trzciński, and P Spurek. “Points2NeRF: Generating Neural Radiance Fields from 3D point cloud”. In: *arXiv preprint arXiv:2206.01290* (2022).
- [44] Jan Zuiderveld, Marco Federici, and Erik J Bekkers. “Towards Lightweight Controllable Audio Synthesis with Conditional Implicit Neural Representations”. In: *arXiv preprint arXiv:2111.08462* (2021).