

---

# Implicit Neural Representations with Periodic Activation Functions

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1        Implicitly defined, continuous, differentiable signal representations parameterized  
2        by neural networks have emerged as a powerful paradigm, offering many possible  
3        benefits over conventional representations. However, current network architec-  
4        tures for such implicit neural representations are incapable of modeling signals  
5        with fine detail, and fail to represent a signal’s spatial and temporal derivatives,  
6        despite the fact that these are essential to many physical signals defined implicitly  
7        as the solution to partial differential equations. We propose to leverage periodic  
8        activation functions for implicit neural representations and demonstrate that these  
9        networks, dubbed sinusoidal representation networks or SIRENS, are ideally suited  
10        for representing complex natural signals and their derivatives. We analyze SIREN  
11        activation statistics to propose a principled initialization scheme and demonstrate  
12        the representation of images, wavefields, video, sound, and their derivatives. Fur-  
13        ther, we show how SIRENS can be leveraged to solve challenging boundary value  
14        problems, such as particular Eikonal equations (yielding signed distance functions),  
15        the Poisson equation, and the Helmholtz and wave equations. Lastly, we combine  
16        SIRENS with hypernetworks to learn priors over the space of SIREN functions.

## 17 1 Introduction

18 We are interested in a class of functions  $\Phi$  that satisfy equations of the form

$$F(\mathbf{x}, \Phi, \nabla_{\mathbf{x}}\Phi, \nabla_{\mathbf{x}}^2\Phi, \dots) = 0, \quad \Phi : \mathbf{x} \mapsto \Phi(\mathbf{x}). \quad (1)$$

19 This implicit problem formulation takes as input the spatial or spatio-temporal coordinates  $\mathbf{x} \in \mathbb{R}^m$   
20 and, optionally, derivatives of  $\Phi$  with respect to these coordinates. Our goal is then to learn a neural  
21 network that parameterizes  $\Phi$  to map  $\mathbf{x}$  to some quantity of interest while satisfying the constraint  
22 presented in Equation (1). Thus,  $\Phi$  is implicitly defined by the relation defined by  $F$  and we refer to  
23 neural networks that parameterize such implicitly defined functions as *implicit neural representations*.  
24 As we show in this paper, a surprisingly wide variety of problems across scientific fields fall into this  
25 form, such as modeling many different types of discrete signals in image, video, and audio processing  
26 using a continuous and differentiable representation, learning 3D shape representations via signed  
27 distance functions [1–4], and, more generally, solving boundary value problems, such as the Poisson,  
28 Helmholtz, or wave equations.

29 A continuous parameterization offers several benefits over alternatives, such as discrete grid-based  
30 representations. For example, due to the fact that  $\Phi$  is defined on the continuous domain of  $\mathbf{x}$ , it can be  
31 significantly more memory efficient than a discrete representation, allowing it to model fine detail that  
32 is not limited by the grid resolution but by the capacity of the underlying network architecture. Being  
33 differentiable implies that gradients and higher-order derivatives can be computed analytically, for  
34 example using automatic differentiation, which again makes these models independent of conventional

35 grid resolutions. Finally, with well-behaved derivatives, implicit neural representations may offer a  
36 new toolbox for solving inverse problems, such as differential equations.

37 For these reasons, implicit neural representations have seen significant research interest over the  
38 last year (Sec. 2). Most of these recent representations build on ReLU-based multilayer perceptrons  
39 (MLPs). While promising, these architectures lack the capacity to represent fine details in the  
40 underlying signals, and they typically do not represent the derivatives of a target signal well. This  
41 is partly due to the fact that ReLU networks are piecewise linear, their second derivative is zero  
42 everywhere, and they are thus incapable of modeling information contained in higher-order derivatives  
43 of natural signals. While alternative activations, such as tanh or softplus, are capable of representing  
44 higher-order derivatives, we demonstrate that their derivatives are often not well behaved and also  
45 fail to represent fine details.

46 To address these limitations, we leverage MLPs with periodic activation functions for implicit neural  
47 representations. We demonstrate that this approach is not only capable of representing details in the  
48 signals better than ReLU-MLPs, or positional encoding strategies proposed in concurrent work [5],  
49 but that these properties also uniquely apply to the derivatives, which is critical for many applications  
50 we explore in this paper.

51 To summarize, the contributions of our work include:

- 52 • A continuous implicit neural representation using periodic activation functions that fits  
53 complicated signals, such as natural images and 3D shapes, and their derivatives robustly.
- 54 • An initialization scheme for training these representations and validation that distributions  
55 of these representations can be learned using hypernetworks.
- 56 • Demonstration of applications in: image, video, and audio representation; 3D shape re-  
57 construction; solving first-order differential equations that aim at estimating a signal by  
58 supervising only with its gradients; and solving second-order differential equations.

## 59 2 Related Work

60 **Implicit neural representations.** Recent work has demonstrated the potential of fully connected  
61 networks as continuous, memory-efficient implicit representations for shape parts [6, 7], objects [1, 4,  
62 8, 9], or scenes [10–12]. These representations are typically trained from some form of 3D data as  
63 either signed distance functions [1, 4, 8–12] or occupancy networks [2, 13]. In addition to representing  
64 shape, some of these models have been extended to also encode object appearance [3, 5, 10, 14, 15],  
65 which can be trained using (multiview) 2D image data using neural rendering [16]. Temporally aware  
66 extensions [17] and variants that add part-level semantic segmentation [18] have also been proposed.

67 **Periodic nonlinearities.** Periodic nonlinearities have been investigated repeatedly over the past  
68 decades, but have so far failed to robustly outperform alternative activation functions. Early work  
69 includes Fourier neural networks, engineered to mimic the Fourier transform via single-hidden-  
70 layer networks [19, 20]. Other work explores neural networks with periodic activations for simple  
71 classification tasks [21–23] and recurrent neural networks [24–28]. It has been shown that such  
72 models have universal function approximation properties [29–31]. Compositional pattern producing  
73 networks [32, 33] also leverage periodic nonlinearities, but rely on a combination of different  
74 nonlinearities via evolution in a genetic algorithm framework. Motivated by the discrete cosine  
75 transform, Klocek et al. [34] leverage cosine activation functions for image representation but they  
76 do not study the derivatives of these representations or other applications explored in our work.  
77 Inspired by these and other seminal works, we explore MLPs with periodic activation functions for  
78 applications involving implicit neural representations and their derivatives, and we propose principled  
79 initialization and generalization schemes.

80 **Neural DE Solvers.** Neural networks have long been investigated in the context of solving differ-  
81 ential equations (DEs) [35], and have previously been introduced as implicit representations for this  
82 task [36]. Early work on this topic involved simple neural network models, consisting of MLPs or  
83 radial basis function networks with few hidden layers and hyperbolic tangent or sigmoid nonlinear-  
84 ities [36–38]. The limited capacity of these shallow networks typically constrained results to 1D  
85 solutions or simple 2D surfaces. Modern approaches to these techniques leverage recent optimization




Figure 1: Comparison of different implicit network architectures fitting a ground truth image (top left). The representation is only supervised on the target image but we also show first- and second-order derivatives of the function fit in rows 2 and 3, respectively.

86 frameworks and auto-differentiation, but use similar architectures based on MLPs. Still, solving more  
 87 sophisticated equations with higher dimensionality, more constraints, or more complex geometries  
 88 is feasible [39–41]. However, we show that the commonly used MLPs with smooth, non-periodic  
 89 activation functions fail to accurately model high-frequency information and higher-order derivatives  
 90 even with dense supervision.

91 Neural ODEs [42] are related to this topic, but are very different in nature. Whereas implicit neural  
 92 representations can be used to directly solve ODEs or PDEs from supervision on the system dynamics,  
 93 neural ODEs allow for continuous function modeling by pairing a conventional ODE solver (e.g.,  
 94 implicit Adams or Runge-Kutta) with a network that parameterizes the dynamics of a function. The  
 95 proposed architecture may be complementary to this line of work.

### 96 3 Formulation

97 Our goal is to solve problems of the form presented in Equation (1). We cast this as  
 98 a feasibility problem, where a function  $\Phi$  is sought that satisfies a set of  $M$  constraints  
 99  $\{\mathcal{C}_m(\mathbf{a}(\mathbf{x}), \Phi(\mathbf{x}), \nabla\Phi(\mathbf{x}), \dots)\}_{m=1}^M$ , each of which relate the function  $\Phi$  and/or its derivatives to  
 100 quantities  $\mathbf{a}(\mathbf{x})$ :

$$\text{find } \Phi(\mathbf{x}) \quad \text{subject to } \mathcal{C}_m(\mathbf{a}(\mathbf{x}), \Phi(\mathbf{x}), \nabla\Phi(\mathbf{x}), \dots) = 0, \quad \forall \mathbf{x} \in \Omega_m, \quad m = 1, \dots, M \quad (2)$$

101 This problem can be cast in a loss function that penalizes deviations from each of the constraints on  
 102 their domain  $\Omega_m$ :

$$\mathcal{L} = \int_{\Omega} \sum_{m=1}^M \mathbf{1}_{\Omega_m}(\mathbf{x}) \|\mathcal{C}_m(\mathbf{a}(\mathbf{x}), \Phi(\mathbf{x}), \nabla\Phi(\mathbf{x}), \dots)\| d\mathbf{x}, \quad (3)$$

103 with the indicator function  $\mathbf{1}_{\Omega_m}(\mathbf{x}) = 1$  when  $\mathbf{x} \in \Omega_m$  and 0 when  $\mathbf{x} \notin \Omega_m$ . In practice, the  
 104 loss function is enforced by sampling  $\Omega$ . A dataset  $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{a}_i(\mathbf{x}_i))\}_i$  is a set of tuples of  
 105 coordinates  $\mathbf{x}_i \in \Omega$  along with samples from the quantities  $\mathbf{a}(\mathbf{x}_i)$  that appear in the constraints.  
 106 Thus, the loss in Equation (3) is enforced on coordinates  $\mathbf{x}_i$  sampled from the dataset, yielding  
 107 the loss  $\tilde{\mathcal{L}} = \sum_{i \in \mathcal{D}} \sum_{m=1}^M \|\mathcal{C}_m(\mathbf{a}(\mathbf{x}_i), \Phi(\mathbf{x}_i), \nabla\Phi(\mathbf{x}_i), \dots)\|$ . In practice, the dataset  $\mathcal{D}$  is sampled  
 108 dynamically at training time, approximating  $\mathcal{L}$  better as the number of samples grows, as in Monte  
 109 Carlo integration.

110 We parameterize functions  $\Phi_{\theta}$  as fully connected neural networks with parameters  $\theta$ , and solve the  
 111 resulting optimization problem using gradient descent.

#### 112 3.1 Periodic Activations for Implicit Neural Representations

113 We propose SIREN, a simple neural network architecture for implicit neural representations that uses  
 114 the sine as a periodic activation function:

$$\Phi(\mathbf{x}) = \mathbf{W}_n (\phi_{n-1} \circ \phi_{n-2} \circ \dots \circ \phi_0)(\mathbf{x}) + \mathbf{b}_n, \quad \mathbf{x}_i \mapsto \phi_i(\mathbf{x}_i) = \sin(\mathbf{W}_i \mathbf{x}_i + \mathbf{b}_i). \quad (4)$$

115 Here,  $\phi_i : \mathbb{R}^{M_i} \mapsto \mathbb{R}^{N_i}$  is the  $i^{\text{th}}$  layer of the network. It consists of the affine transform defined by  
 116 the weight matrix  $\mathbf{W}_i \in \mathbb{R}^{N_i \times M_i}$  and the biases  $\mathbf{b}_i \in \mathbb{R}^{N_i}$  applied on the input  $\mathbf{x}_i \in \mathbb{R}^{M_i}$ , followed  
 117 by the sine nonlinearity applied to each component of the resulting vector.

118 Interestingly, any derivative of a SIREN *is itself* a SIREN, as the derivative of the sine is a cosine, i.e.,  
 119 a phase-shifted sine (see supplemental). Therefore, the derivatives of a SIREN inherit the properties  
 120 of SIRENS, enabling us to supervise any derivative of SIREN with “complicated” signals. In our  
 121 experiments, we demonstrate that when a SIREN is supervised using a constraint  $\mathcal{C}_m$  involving the  
 122 derivatives of  $\phi$ , the function  $\phi$  remains well behaved, which is crucial in solving many problems,  
 123 including boundary value problems (BVPs).

124 We will show that SIRENS can be initialized with some control over the distribution of activations,  
 125 allowing us to create deep architectures. Furthermore, SIRENS converge significantly faster than  
 126 baseline architectures, fitting, for instance, a single image in a few hundred iterations, taking a few  
 127 seconds on a modern GPU, while featuring higher image fidelity (Fig. 1).

128 **A simple example: fitting an im-**  
 129 **age.** Consider the case of finding

130 the function  $\Phi : \mathbb{R}^2 \mapsto \mathbb{R}^3, \mathbf{x} \rightarrow$   
 131  $\Phi(\mathbf{x})$  that parameterizes a given discrete  
 132 image  $f$  in a continuous fashion. The image defines a dataset  
 133  $\mathcal{D} = \{(\mathbf{x}_i, f(\mathbf{x}_i))\}_i$  of pixel coordi-  
 134 nates  $\mathbf{x}_i = (x_i, y_i)$  associated with  
 135 their RGB colors  $f(\mathbf{x}_i)$ . The only constraint  
 136  $\mathcal{C}$  enforces is that  $\Phi$  shall output  
 137 image colors at pixel coordinates,  
 138 solely depending on  $\Phi$  (none of its  
 139 derivatives) and  $f(\mathbf{x}_i)$ , with the form  
 140  $\mathcal{C}(f(\mathbf{x}_i), \Phi(\mathbf{x})) = \Phi(\mathbf{x}_i) - f(\mathbf{x}_i)$   
 141 which can be translated into the loss  
 142  $\tilde{\mathcal{L}} = \sum_i \|\Phi(\mathbf{x}_i) - f(\mathbf{x}_i)\|^2$ . In Fig. 1,

143 we fit  $\Phi_\theta$  using comparable network architectures with different activation functions to a natural  
 144 image. We supervise this experiment only on the image values, but also visualize the gradients  $\nabla f$   
 145 and Laplacians  $\Delta f$ . While only two approaches, a ReLU network with positional encoding (P.E.) [5]  
 146 and our SIREN, accurately represent the ground truth image  $f(\mathbf{x})$ , SIREN is the only network capable  
 147 of also representing the derivatives of the signal. Additionally, we run a simple experiment where  
 148 we fit a short video with 300 frames and with a resolution of  $512 \times 512$  pixels using both ReLU and  
 149 SIREN MLPs. As seen in Figure 2, our approach is successful in representing this video with an  
 150 average peak signal-to-noise ratio close to 30 dB, outperforming the ReLU baseline by about 5 dB.  
 151 We also show the flexibility of SIRENS by representing audio signals in the supplement.  
 152



Figure 2: Example frames from fitting a video with SIREN and ReLU-MLPs. Our approach faithfully reconstructs fine details like the whiskers. Mean (and standard deviation) of the PSNR over all frames is reported.

### 153 3.2 Distribution of activations, frequencies, and a principled initialization scheme

154 We present a principled initialization scheme necessary for the effective training of SIRENS. While  
 155 presented informally here, we discuss further details, proofs and empirical validation in the supple-  
 156 mental material. The key idea in our initialization scheme is to preserve the distribution of activations  
 157 through the network so that the final output at initialization does not depend on the number of layers.  
 158 Note that building SIRENS with not carefully chosen uniformly distributed weights yielded poor  
 159 performance both in accuracy and in convergence speed.

160 To this end, let us first consider the output distribution of a single sine neuron with the uniformly  
 161 distributed input  $x \sim \mathcal{U}(-1, 1)$ . The neuron’s output is  $y = \sin(ax + b)$  with  $a, b \in \mathbb{R}$ . It  
 162 can be shown that for any  $a > \frac{\pi}{2}$ , i.e. spanning at least half a period, the output of the sine is  
 163  $y \sim \text{arcsine}(-1, 1)$ , a special case of a U-shaped Beta distribution and independent of the choice of  
 164  $b$ . We can now reason about the output distribution of a neuron. Taking the linear combination of  $n$   
 165 inputs  $\mathbf{x} \in \mathbb{R}^n$  weighted by  $\mathbf{w} \in \mathbb{R}^n$ , its output is  $y = \sin(\mathbf{w}^T \mathbf{x} + b)$ . Assuming this neuron is in  
 166 the second layer, each of its inputs is arcsine distributed. When each component of  $\mathbf{w}$  is uniformly  
 167 distributed such as  $w_i \sim \mathcal{U}(-c/\sqrt{n}, c/\sqrt{n})$ ,  $c \in \mathbb{R}$ , we show (see supplemental) that the dot product  
 168 converges to the normal distribution  $\mathbf{w}^T \mathbf{x} \sim \mathcal{N}(0, c^2/6)$  as  $n$  grows. Finally, feeding this normally

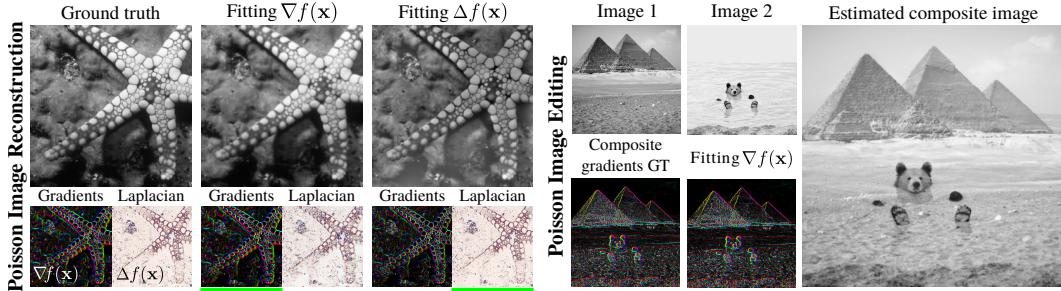


Figure 3: **Poisson image reconstruction:** An image (left) is reconstructed by fitting a SIREN, supervised either by its gradients or Laplacians (underlined in green). The results, shown in the center and right, respectively, match both the image and its derivatives well. **Poisson image editing:** The gradients of two images (top) are fused (bottom left). SIREN allows for the composite (right) to be reconstructed using supervision on the gradients (bottom right).

169 distributed dot product through another sine is also arcsine distributed for any  $c > \sqrt{6}$ . Note that the  
 170 weights of a SIREN can be interpreted as angular frequencies while the biases are phase offsets. Thus,  
 171 larger frequencies appear in the networks for weights with larger magnitudes. For  $|\mathbf{w}^T \mathbf{x}| < \pi/4$ ,  
 172 the sine layer will leave the frequencies unchanged, as the sine is approximately linear. In fact, we  
 173 empirically find that a sine layer keeps spatial frequencies approximately constant for amplitudes  
 174 such as  $|\mathbf{w}^T \mathbf{x}| < \pi$ , and increases spatial frequencies for amplitudes above this value<sup>1</sup>.

175 Hence, we propose to draw weights with  $c = 6$  so that  $w_i \sim \mathcal{U}(-\sqrt{6/n}, \sqrt{6/n})$ . This ensures that  
 176 the input to each sine activation is normal distributed with a standard deviation of 1. Since only a  
 177 few weights have a magnitude larger than  $\pi$ , the frequency throughout the sine network grows only  
 178 slowly. Finally, we propose to initialize the first layer of the sine network with weights so that the  
 179 sine function  $\sin(\omega_0 \cdot \mathbf{W}\mathbf{x} + \mathbf{b})$  spans multiple periods over  $[-1, 1]$ . We found  $\omega_0 = 30$  to work  
 180 well for all the applications in this work. The proposed initialization scheme yielded fast and robust  
 181 convergence using the ADAM optimizer for all experiments in this work.

## 182 4 Experiments

183 In this section, we leverage SIRENS to solve challenging boundary value problems using different types  
 184 of supervision of the derivatives of  $\Phi$ . We first solve the Poisson equation via direct supervision of its  
 185 derivatives. We then solve a particular form of the Eikonal equation, placing a unit-norm constraint  
 186 on gradients, parameterizing the class of signed distance functions (SDFs). SIREN significantly  
 187 outperforms ReLU-based SDFs, capturing large scenes at a high level of detail. We then solve  
 188 the second-order Helmholtz partial differential equation, and the challenging inverse problem of  
 189 full-waveform inversion. Finally, we combine SIRENS with hypernetworks, learning a prior over the  
 190 space of parameterized functions. All code and data will be made publicly available.

### 191 4.1 Solving the Poisson Equation

192 We demonstrate that the proposed representation is not only able to accurately represent a function  
 193 and its derivatives, but that it can also be supervised solely by its derivatives, i.e., the model is never  
 194 presented with the actual function values, but only values of its first or higher-order derivatives.

195 An intuitive example representing this class of problems is the Poisson equation. The Poisson  
 196 equation is perhaps the simplest elliptic partial differential equation (PDE) which is crucial in physics  
 197 and engineering, for example to model potentials arising from distributions of charges or masses.  
 198 In this problem, an unknown ground truth signal  $f$  is estimated from discrete samples of either its  
 199 gradients  $\nabla f$  or Laplacian  $\Delta f = \nabla \cdot \nabla f$  as

$$\mathcal{L}_{\text{grad.}} = \int_{\Omega} \|\nabla_{\mathbf{x}} \Phi(\mathbf{x}) - \nabla_{\mathbf{x}} f(\mathbf{x})\| d\mathbf{x}, \quad \text{or} \quad \mathcal{L}_{\text{lapl.}} = \int_{\Omega} \|\Delta \Phi(\mathbf{x}) - \Delta f(\mathbf{x})\| d\mathbf{x}. \quad (5)$$

<sup>1</sup>Formalizing the distribution of output frequencies throughout SIRENS proves to be a hard task and is out of the scope of this work.




Figure 4: Shape representation. We fit signed distance functions parameterized by implicit neural representations directly on point clouds. Compared to ReLU implicit representations, our periodic activations significantly improve detail of objects (left) and complexity of entire scenes (right).

200 **Poisson image reconstruction.** Solving the Poisson equation enables the reconstruction of images  
 201 from their derivatives. We show results of this approach using SIREN in Fig. 3. Supervising the implicit  
 202 representation with either ground truth gradients via  $\mathcal{L}_{\text{grad}}$ , or Laplacians via  $\mathcal{L}_{\text{lapl}}$ , successfully  
 203 reconstructs the image. Remaining intensity variations are due to the ill-posedness of the problem.

204 **Poisson image editing.** Images can be seamlessly fused in the gradient domain [43]. For this  
 205 purpose,  $\Phi$  is supervised using  $\mathcal{L}_{\text{grad}}$ , of Eq. (5), where  $\nabla_{\mathbf{x}}f(\mathbf{x})$  is a composite function of the  
 206 gradients of two images  $f_{1,2}$ :  $\nabla_{\mathbf{x}}f(\mathbf{x}) = \alpha \cdot \nabla f_1(x) + (1 - \alpha) \cdot \nabla f_2(x)$ ,  $\alpha \in [0, 1]$ . Fig. 3 shows  
 207 two images seamlessly fused with this approach.

## 208 4.2 Representing Shapes with Signed Distance Functions

209 Inspired by recent work on shape representation with differentiable signed distance functions  
 210 (SDFs) [1, 4, 9], we fit SDFs directly on oriented point clouds using both ReLU-based implicit  
 211 neural representations and SIRENS. This amounts to solving a particular Eikonal boundary value  
 212 problem that constrains the norm of spatial gradients  $|\nabla_{\mathbf{x}}\Phi|$  to be 1 almost everywhere. Note that  
 213 ReLU networks are seemingly ideal for representing SDFs, as their gradients are locally constant and  
 214 their second derivatives are 0. Adequate training procedures for working directly with point clouds  
 215 were described in prior work [4, 9]. We fit a SIREN to an oriented point cloud using a loss of the form

$$\mathcal{L}_{\text{sdf}} = \int_{\Omega} \left\| |\nabla_{\mathbf{x}}\Phi(\mathbf{x})| - 1 \right\| d\mathbf{x} + \int_{\Omega_0} \|\Phi(\mathbf{x})\| + (1 - \langle \nabla_{\mathbf{x}}\Phi(\mathbf{x}), \mathbf{n}(\mathbf{x}) \rangle) d\mathbf{x} + \int_{\Omega \setminus \Omega_0} \psi(\Phi(\mathbf{x})) d\mathbf{x}, \quad (6)$$

216 Here,  $\psi(\mathbf{x}) = \exp(-\alpha \cdot |\Phi(\mathbf{x})|)$ ,  $\alpha \gg 1$  penalizes off-surface points for creating SDF values close  
 217 to 0.  $\Omega$  is the whole domain and we denote the zero-level set of the SDF as  $\Omega_0$ . The model  $\Phi(x)$  is  
 218 supervised using oriented points sampled on a mesh, where we require the SIREN to respect  $\Phi(\mathbf{x}) = 0$   
 219 and its normals  $\mathbf{n}(\mathbf{x}) = \nabla f(\mathbf{x})$ . During training, each minibatch contains an equal number of points  
 220 on and off the mesh, each one randomly sampled over  $\Omega$ . As seen in Fig. 4, the proposed periodic  
 221 activations significantly increase the details of objects and the complexity of scenes that can be  
 222 represented by these neural SDFs, parameterizing a full room with only five fully connected layers.

## 223 4.3 Solving the Helmholtz and Wave Equations

224 The Helmholtz and wave equations are second-order partial differential equations related to the  
 225 physical modeling of diffusion and waves. They are closely related through a Fourier-transform  
 226 relationship, with the Helmholtz equation given as

$$H(m)\Phi(\mathbf{x}) = -f(\mathbf{x}), \text{ with } H(m) = (\Delta + m(\mathbf{x})w^2). \quad (7)$$

227 Here,  $f(\mathbf{x})$  represents a known source function,  $\Phi(\mathbf{x})$  is the unknown wavefield, and the squared  
 228 slowness  $m(\mathbf{x}) = 1/c(\mathbf{x})^2$  is a function of the wave velocity  $c(\mathbf{x})$ . In general, the solutions to the  
 229 Helmholtz equation are complex-valued and require numerical solvers to compute. As the Helmholtz  
 230 and wave equations follow a similar form, we discuss the Helmholtz equation here, with additional  
 231 results and discussion for the wave equation in the supplement.

232 **Solving for the wavefield.** We solve for the wavefield by parameterizing  $\Phi(\mathbf{x})$  with a SIREN. To  
 233 accommodate a complex-valued solution, we configure the network to output two values, interpreted  
 234 as the real and imaginary parts. Training is performed on randomly sampled points  $\mathbf{x}$  within the




Figure 5: **Direct Inversion:** We solve the Helmholtz equation for a single point source placed at the center of a medium (green dot) with uniform wave propagation velocity (top left). The SIREN solution closely matches a principled grid solver [44] while other network architectures fail to find the correct solution. **Neural Full-Waveform Inversion (FWI):** A scene contains a source (green) and a circular wave velocity perturbation centered at the origin (top left). With the scene velocity known *a priori*, SIREN directly reconstructs a wavefield that closely matches a principled grid solver [44] (bottom left, middle left). For FWI, the velocity and wavefields are reconstructed with receiver measurements (blue dots) from sources triggered in sequence (green, red dots). The SIREN velocity model outperforms a principled FWI solver [45], accurately predicting wavefields. FWI MSE values are calculated across all wavefields and the visualized real wavefield corresponds to the green source.

235 domain  $\Omega = \{\mathbf{x} \in \mathbb{R}^2 \mid \|\mathbf{x}\|_\infty < 1\}$ . The network is supervised using a loss function based  
 236 on the Helmholtz equation,  $\mathcal{L}_{\text{Helmholtz}} = \int_{\Omega} \lambda(\mathbf{x}) \|H(m)\Phi(\mathbf{x}) + f(\mathbf{x})\|_1 d\mathbf{x}$ , with  $\lambda(\mathbf{x}) = k$ , a  
 237 hyperparameter, when  $f(\mathbf{x}) \neq 0$  (corresponding to the inhomogeneous contribution to the Helmholtz  
 238 equation) and  $\lambda(\mathbf{x}) = 1$  otherwise (for the homogenous part). Each minibatch contains samples from  
 239 both contributions and  $k$  is set so the losses are approximately equal at the beginning of training. In  
 240 practice, we use a slightly modified form of Equation (7) to include the perfectly matched boundary  
 241 conditions that are necessary to ensure a unique solution [44] (see supplement for details).

242 Results are shown in Fig. 5 for solving the Helmholtz equation in two dimensions with spatially  
 243 uniform wave velocity and a single point source (modeled as a Gaussian with  $\sigma^2 = 10^{-4}$ ). The  
 244 SIREN solution is compared with a principled solver [44] as well as other neural network solvers. All  
 245 evaluated network architectures use the same number of hidden layers as SIREN but with different  
 246 activation functions. In the case of the RBF network, we prepend an RBF layer with 1024 hidden  
 247 units and use a tanh activation. SIREN is the only representation capable of producing a high-fidelity  
 248 reconstruction of the wavefield. We also note that the tanh network has a similar architecture to recent  
 249 work on neural PDE solvers [40], except we increase the network size to match SIREN.

250 **Neural full-waveform inversion (FWI).** In many wave-based sensing modalities (radar, sonar,  
 251 seismic imaging, etc.), one attempts to probe and sense across an entire domain using sparsely placed  
 252 sources (i.e., transmitters) and receivers. FWI uses the known locations of sources and receivers to  
 253 jointly recover the entire wavefield and other physical properties, such as permittivity, density, or  
 254 wave velocity. Specifically, the FWI problem can be described as [46]

$$\arg \min_{m, \Phi} \sum_{i=1}^N \int_{\Omega} |\text{III}_r(\Phi_i(\mathbf{x}) - r_i(\mathbf{x}))|^2 d\mathbf{x} \text{ s.t. } H(m)\Phi_i(\mathbf{x}) = -f_i(x), 1 \leq i \leq N, \forall \mathbf{x} \in \Omega, \quad (8)$$

255 where there are  $N$  sources,  $\text{III}_r$  samples the wavefield at the receiver locations, and  $r_i(x)$  models  
 256 receiver data for the  $i$ th source.

257 We first use a SIREN to directly solve Eq. 7 for a known wave velocity perturbation, obtaining an  
 258 accurate wavefield that closely matches that of a principled solver [44] (see Fig. 5, right). Without  
 259 *a priori* knowledge of the velocity field, FWI is used to jointly recover the wavefields and velocity.  
 260 Here, we use 5 sources and place 30 receivers around the domain, as shown in Fig. 5. Using the  
 261 principled solver, we simulate the receiver measurements for the 5 wavefields (one for each source)  
 262 at a single frequency of 3.2 Hz, which is chosen to be relatively low for improved convergence. We  
 263 pre-train SIREN to output 5 complex wavefields and a squared slowness value for a uniform velocity.  
 264 Then, we optimize for the wavefields and squared slowness using a penalty method variation [46]  
 265 of Eq. 8 (see the supplement for additional details). In Fig. 5, we compare to an FWI solver based  
 266 on the alternating direction method of multipliers [45, 47]. With only a single frequency for the  
 267 inversion, the principled solver is prone to converge to a poor solution for the velocity. As shown in




Figure 6: Generalizing across implicit functions parameterized by SIRENS on the CelebA dataset [48]. Image inpainting results are shown for various numbers of context pixels in  $O_j$ .

268 Fig. 5, SIREN converges to a better velocity solution and accurate solutions for the wavefields. All  
 269 reconstructions are performed or shown at  $256 \times 256$  resolution to avoid noticeable stair-stepping  
 270 artifacts in the circular velocity perturbation.

#### 271 4.4 Learning a Space of Implicit Functions

272 A powerful concept that has emerged for implicit representations is to learn priors over the space  
 273 of functions that define them [1, 2, 10]. Here we demonstrate that the function space parameterized  
 274 by SIRENS also admits the learning of powerful priors. Each of these SIRENS  $\Phi_j$  are fully defined  
 275 by their parameters  $\theta_j \in \mathbb{R}^l$ . Assuming that all parameters  $\theta_j$  of a class exist in a  $k$ -dimensional  
 276 subspace of  $\mathbb{R}^l$ ,  $k < l$ , then these parameters can be well modeled by latent code vectors in  $\mathbf{z} \in \mathbb{R}^k$ .  
 277 Like in neural processes [49–51], we condition these latent code vectors on partial observations of  
 278 the signal  $O \in \mathbb{R}^m$  through an encoder

$$C : \mathbb{R}^m \rightarrow \mathbb{R}^k, \quad O_j \mapsto C(O_j) = \mathbf{z}_j, \quad (9)$$

279 and use a ReLU hypernetwork [52], to map the latent code to the weights of a SIREN, as in [10]:

$$\Psi : \mathbb{R}^k \rightarrow \mathbb{R}^l, \quad \mathbf{z}_j \mapsto \Psi(\mathbf{z}_j) = \theta_j. \quad (10)$$

280 We replicated the experiment from [49] on the CelebA dataset [48] using a set encoder. Additionally,  
 281 we show results using a convolutional neural network encoder which operates on sparse images.  
 282 Interestingly, this improves the quantitative and qualitative performance on the inpainting task.  
 283

284 At test time, this enables reconstruction from sparse pixel observations, and, thereby, inpainting. Fig. 6 shows test-time  
 285 reconstructions from a varying number of pixel observations. Note that these inpainting  
 286 results were all generated using the same model, with the same parameter values. Tab. 1 reports a quantitative comparison  
 287 to [49], demonstrating that generalization over SIREN representations is at least equally as powerful as generalization over images.

Table 1: Quantitative comparison to Conditional Neural Processes [49] (CNPs) on the  $32 \times 32$  CelebA test set. Metrics are reported in pixel-wise mean squared error.

Number of Context Pixels	10	100	1000
CNP [49]	0.039	0.016	0.009
Set Encoder + Hypernet.	0.035	0.013	0.009
CNN Encoder + Hypernet.	<b>0.033</b>	<b>0.009</b>	<b>0.008</b>

## 294 5 Discussion and Conclusion

295 The question of how to represent a signal is at the core of many problems across science and  
 296 engineering. Implicit neural representations may provide a new tool for many of these by offering  
 297 a number of potential benefits over conventional continuous and discrete representations. We  
 298 demonstrate that periodic activation functions are ideally suited for representing complex natural  
 299 signals and their derivatives using implicit neural representations. We also prototype several boundary  
 300 value problems that our framework is capable of solving robustly. There are several exciting avenues  
 301 for future work, including the exploration of other types of inverse problems and applications in areas  
 302 beyond implicit neural representations, for example neural ODEs [42].

303 With this work, we make important contributions to the emerging field of implicit neural representation  
 304 learning and its applications.



## 305 **Broader Impact**

306 The proposed SIREN representation enables accurate representations of natural signals, such as  
307 images, audio, and video in a deep learning framework. This may be an enabler for downstream  
308 tasks involving such signals, such as classification for images or speech-to-text systems for audio.  
309 Such applications may be leveraged for both positive and negative ends. SIREN may in the future  
310 further enable novel approaches to the generation of such signals. This has potential for misuse in  
311 impersonating actors without their consent. For an in-depth discussion of such so-called DeepFakes,  
312 we refer the reader to a recent review article on neural rendering [16].

## 313 **References**

- 314 [1] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF:  
315 Learning continuous signed distance functions for shape representation. *Proc. CVPR*, 2019.
- 316 [2] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy  
317 networks: Learning 3d reconstruction in function space. In *Proc. CVPR*, 2019.
- 318 [3] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu:  
319 Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proc. ICCV*, pages  
320 2304–2314, 2019.
- 321 [4] Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *Proc. CVPR*,  
322 2020.
- 323 [5] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng.  
324 Nerf: Representing scenes as neural radiance fields for view synthesis. *arXiv preprint arXiv:2003.08934*,  
325 2020.
- 326 [6] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser.  
327 Learning shape templates with structured implicit functions. In *Proc. ICCV*, pages 7154–7164, 2019.
- 328 [7] Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. Deep structured  
329 implicit functions. *arXiv preprint arXiv:1912.06126*, 2019.
- 330 [8] Mateusz Michalkiewicz, Jhony K Pontes, Dominic Jack, Mahsa Baktashmotlagh, and Anders Eriksson.  
331 Implicit surface representations as layers in neural networks. In *Proc. ICCV*, pages 4743–4752, 2019.
- 332 [9] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization  
333 for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020.
- 334 [10] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous  
335 3d-structure-aware neural scene representations. In *Proc. NeurIPS*, 2019.
- 336 [11] Chiyu Max Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, and Thomas  
337 Funkhouser. Local implicit grid representations for 3d scenes. *arXiv preprint arXiv:2003.08981*, 2020.
- 338 [12] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional  
339 occupancy networks. *arXiv preprint arXiv:2003.04618*, 2020.
- 340 [13] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proc. CVPR*,  
341 pages 5939–5948, 2019.
- 342 [14] Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields:  
343 Learning texture representations in function space. In *Proc. ICCV*, 2019.
- 344 [15] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric  
345 rendering: Learning implicit 3d representations without 3d supervision. In *Proc. CVPR*, 2020.
- 346 [16] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo  
347 Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, et al. State of the art on neural rendering.  
348 *Proc. Eurographics*, 2020.
- 349 [17] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d recon-  
350 struction by learning particle dynamics. In *Proc. ICCV*, 2019.
- 351 [18] Amit Kohli, Vincent Sitzmann, and Gordon Wetzstein. Inferring semantic information with 3d neural  
352 scene representations. *arXiv preprint arXiv:2003.12673*, 2020.
- 353 [19] R. Gallant and H. White. There exists a neural network that does not make avoidable mistakes. In *IEEE*  
354 *Int. Conference on Neural Networks*, pages 657–664, 1988.
- 355 [20] Abylay Zhumekenov, Malika Uteuliyeva, Olzhas Kabdолоv, Rustem Takhanov, Zhenisbek Assylbekov,  
356 and Alejandro J Castro. Fourier neural networks: A comparative study. *arXiv preprint arXiv:1902.03011*,  
357 2019.

- 358 [21] Josep M Sopena, Enrique Romero, and Rene Alquezar. Neural networks with periodic and monotonic  
359 activation functions: a comparative study in classification problems. In *Proc. ICANN*, 1999.
- 360 [22] Kwok-wo Wong, Chi-sing Leung, and Sheng-jiang Chang. Handwritten digit recognition using multilayer  
361 feedforward neural networks with periodic and monotonic activation functions. In *Object recognition  
362 supported by user interaction for service robots*, volume 3, pages 106–109. IEEE, 2002.
- 363 [23] Giambattista Parascandolo, Heikki Huttunen, and Tuomas Virtanen. Taming the waves: sine as activation  
364 function in deep neural networks. 2016.
- 365 [24] Peng Liu, Zhigang Zeng, and Jun Wang. Multistability of recurrent neural networks with nonmonotonic  
366 activation functions and mixed time delays. *IEEE Trans. on Systems, Man, and Cybernetics: Systems*, 46  
367 (4):512–523, 2015.
- 368 [25] Renée Koplun and Eduardo D Sontag. Using fourier-neural recurrent networks to fit sequential input/output  
369 data. *Neurocomputing*, 15(3-4):225–248, 1997.
- 370 [26] M Hisham Choueiki, Clark A Mount-Campbell, and Stanley C Ahalt. Implementing a weighted least  
371 squares procedure in training a neural network to solve the short-term load forecasting problem. *IEEE  
372 Trans. on Power systems*, 12(4):1689–1694, 1997.
- 373 [27] René Alquézar Mancho. *Symbolic and connectionist learning techniques for grammatical inference*.  
374 Universitat Politècnica de Catalunya, 1997.
- 375 [28] JM Sopena and R Alquezar. Improvement of learning in recurrent networks by substituting the sigmoid  
376 activation function. In *Proc. ICANN*, pages 417–420. Springer, 1994.
- 377 [29] Emmanuel J Candès. Harmonic analysis of neural networks. *Applied and Computational Harmonic  
378 Analysis*, 6(2):197–218, 1999.
- 379 [30] Shaobo Lin, Xiaofei Guo, Feilong Cao, and Zongben Xu. Approximation by neural networks with scattered  
380 data. *Applied Mathematics and Computation*, 224:29–35, 2013.
- 381 [31] Sho Sonoda and Noboru Murata. Neural network with unbounded activation functions is universal  
382 approximator. *Applied and Computational Harmonic Analysis*, 43(2):233–268, 2017.
- 383 [32] Kenneth O Stanley. Compositional pattern producing networks: A novel abstraction of development.  
384 *Genetic programming and evolvable machines*, 8(2):131–162, 2007.
- 385 [33] Alexander Mordvintsev, Nicola Pezzotti, Ludwig Schubert, and Chris Olah. Differentiable image parame-  
386 terizations. *Distill*, 3(7):e12, 2018.
- 387 [34] Sylwester Klocek, Łukasz Maziarka, Maciej Wołczyk, Jacek Tabor, Jakub Nowak, and Marek Śmieja.  
388 Hypernetwork functional image representation. In *Proc. ICANN*, pages 496–510. Springer, 2019.
- 389 [35] Hyuk Lee and In Seok Kang. Neural algorithm for solving differential equations. *Journal of Computational  
390 Physics*, 91(1):110–131, 1990.
- 391 [36] Isaac E Lagaris, Aristidis Likas, and Dimitrios I Fotiadis. Artificial neural networks for solving ordinary  
392 and partial differential equations. *IEEE Trans. on neural networks*, 9(5):987–1000, 1998.
- 393 [37] Shouling He, Konrad Reif, and Rolf Unbehauen. Multilayer neural networks for solving a class of partial  
394 differential equations. *Neural networks*, 13(3):385–396, 2000.
- 395 [38] Nam Mai-Duy and Thanh Tran-Cong. Approximation of function and its derivatives using radial basis  
396 function networks. *Applied Mathematical Modelling*, 27(3):197–220, 2003.
- 397 [39] Justin Sirignano and Konstantinos Spiliopoulos. Dgm: A deep learning algorithm for solving partial  
398 differential equations. *Journal of Computational Physics*, 375:1339–1364, 2018.
- 399 [40] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep  
400 learning framework for solving forward and inverse problems involving nonlinear partial differential  
401 equations. *Journal of Computational Physics*, 378:686–707, 2019.
- 402 [41] Jens Berg and Kaj Nyström. A unified deep artificial neural network approach to partial differential  
403 equations in complex geometries. *Neurocomputing*, 317:28–41, 2018.
- 404 [42] Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential  
405 equations. In *Proc. NIPS*, pages 6571–6583, 2018.
- 406 [43] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. *ACM Trans. on Graphics*, 22  
407 (3):313–318, 2003.
- 408 [44] Zhongying Chen, Dongsheng Cheng, Wei Feng, and Tingting Wu. An optimal 9-point finite difference  
409 scheme for the helmholtz equation with pml. *International Journal of Numerical Analysis & Modeling*, 10  
410 (2), 2013.
- 411 [45] Hossein S Aghamiry, Ali Gholami, and Stéphane Operto. Improving full-waveform inversion by wavefield  
412 reconstruction with the alternating direction method of multipliers. *Geophysics*, 84(1):R139–R162, 2019.

- 413 [46] Tristan Van Leeuwen and Felix J Herrmann. Mitigating local minima in full-waveform inversion by  
414 expanding the search space. *Geophysical Journal International*, 195(1):661–667, 2013.
- 415 [47] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization  
416 and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in*  
417 *Machine learning*, 3(1):1–122, 2011.
- 418 [48] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proc.*  
419 *ICCV*, December 2015.
- 420 [49] Marta Garnelo, Dan Rosenbaum, Chris J Maddison, Tiago Ramalho, David Saxton, Murray Shana-  
421 han, Yee Whye Teh, Danilo J Rezende, and SM Eslami. Conditional neural processes. *arXiv preprint*  
422 *arXiv:1807.01613*, 2018.
- 423 [50] SM Ali Eslami, Danilo Jimenez Rezende, Frederic Besse, Fabio Viola, Ari S Morcos, Marta Garnelo,  
424 Avraham Ruderman, Andrei A Rusu, Ivo Danihelka, Karol Gregor, et al. Neural scene representation and  
425 rendering. *Science*, 360(6394):1204–1210, 2018.
- 426 [51] Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol Vinyals,  
427 and Yee Whye Teh. Attentive neural processes. *Proc. ICLR*, 2019.
- 428 [52] David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. In *Proc. ICLR*, 2017.