

Modulated Periodic Activations for Generalizable Local Functional Representations

Ishit Mehta¹ Michaël Gharbi² Connelly Barnes²
 Eli Shechtman² Ravi Ramamoorthi¹ Manmohan Chandraker¹
¹UC San Diego ²Adobe Research

Abstract

Multi-Layer Perceptrons (MLPs) make powerful functional representations for sampling and reconstruction problems involving low-dimensional signals like images, shapes and light fields. Recent works have significantly improved their ability to represent high-frequency content by using periodic activations or positional encodings. This often came at the expense of generalization: modern methods are typically optimized for a single signal. We present a new representation that generalizes to multiple instances and achieves state-of-the-art fidelity. We use a dual-MLP architecture to encode the signals. A synthesis network creates a functional mapping from a low-dimensional input (e.g. pixel-position) to the output domain (e.g. RGB color). A modulation network maps a latent code corresponding to the target signal to parameters that modulate the periodic activations of the synthesis network. We also propose a local-functional representation which enables generalization. The signal’s domain is partitioned into a regular grid, with each tile represented by a latent code. At test time, the signal is encoded with high-fidelity by inferring (or directly optimizing) the latent code-book. Our approach produces generalizable functional representations of images, videos and shapes, and achieves higher reconstruction quality than prior works that are optimized for a single signal.

1. Introduction

Functional neural representations using Multi-Layer Perceptrons (MLPs) have garnered renewed interest for their conceptual simplicity and ability to approximate complex signals like images, videos, audio recordings [38], light-fields [26] and implicitly-defined 3D shapes [8, 32, 2]. They have shown to be more compact and efficient than their discrete counterparts [21, 39]. While recent contributions have focused on improving the accuracy of these representations, in particular to model complex signals with high-frequency details [38, 44, 26], it is still challenging to generalize them to unseen signals. Recent approaches typically require train-

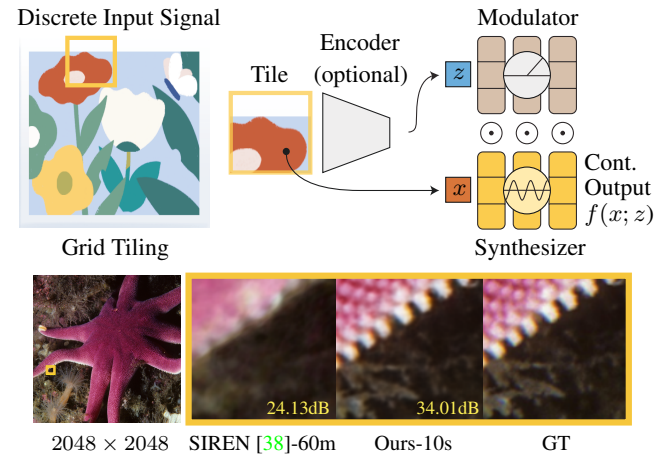


Figure 1: (Top) We propose a new method to encode discrete signals as neural-functional representations. Each signal is represented as a structured grid, with each tile defined by a latent code. The tiles are encoded as continuous signals by a synthesis network using periodic activations. The latent codes corresponding to the tiles are used to modulate the activations using a modulation network. (Bottom) Our generalizable method is faster and enables higher quality reconstructions than previous methods.

ing a separate MLP for *each* signal [26, 9]. Previous efforts sought to improve generalization by imposing priors on the functional space spanned by the MLP parameterization [32, 35], using hypernetworks [14, 38], or via meta-learning [37]. But multi-instance generalization still causes significant degradations in quality.

We introduce a neural functional representation that *simultaneously* achieves high-reconstruction quality and generalizes to multiple instances. Our approach can encode functional representations for multiple discrete signals using a *single* model. Unlike previous works, which train a model for each signal, it can do so in a single feed-forward pass (Fig. 1). We represent each signal using a low-dimensional latent code. These codes serve as conditioning variables in a functional mapping that uses two MLPs: a modulator and a synthesis network. The *synthesis* network implements a map-




Figure 2: Obtaining a neural representation of an image using SIREN [38] requires overfitting a new MLP to every new image instance. This works well in practice for low-resolution images, but with the same model, we observe a sharp drop in reconstruction accuracy for high-resolution images. Our method can be used to encode *multiple* images at high-resolution with only feed-forward operations without gradient computation.

ping from coordinates (*e.g.* spatial position) to signal values (*e.g.* RGB colors). It uses the sine function as activation, which enables accurate reconstructions of high-frequency content [38], but also makes naive conditioning strategies ineffective (§ 3.2). The *modulator* is the key to generalization. It consumes latent code and outputs, at each layer, parameters that modulate the amplitude, phase and frequency of periodic activations in the synthesis network. The modulator uses ReLU activations. Our model can either be used as an autoencoder, where the latent codes are produced by a third network (the encoder); or as an auto-decoder, where the latent codes are optimized jointly with all the network parameters.

As we show in Figure 2, the quality of functional representations that fit images as-a-whole, degrades as we increase the target resolution. High-resolution images typically have a broad power spectrum, thereby requiring more expensive models to represent them functionally. But images are usually much simpler locally: simple edges and textures re-occur commonly across images that are otherwise quite distinct at the global level. This motivates our strategy to exploit locality. We partition the signal domain into a regular tiling, and assign each tile a latent code (Fig. 1). By “zooming in” on the local structure, computing functional approximations that generalize becomes more tractable [23], because simple parts exhibit fewer variations than complete objects [13]. Some recent work has explored locality, but they focus on relatively simple, low-frequency signals like signed distance fields [4, 13], which can locally be well approximated using a single linear decision boundary—well in the purview of ReLU-based MLPs. For more complex signals like images and videos, even local patterns contain high-frequency components that a standard ReLU-MLP fails

to reconstruct (Figure 3). We show that locality, together with our model architecture, makes it possible to obtain functional representations of large, complex signals. Compared to previous methods, ours produces qualitatively and quantitatively better functional representations, with improved generalization capabilities. Concretely, our contributions are as follows:

- A local neural-functional representation that enables generalization and achieves high fidelity. We use a set of local functions defined on a tiling of the input domain that combine to reconstruct the target signal.
- A new network architecture, which uses modulation and synthesis sub-networks for high-fidelity functional neural representations of images, shapes and videos.

2. Related Work

Continuous Representations of Visual Signals Our work builds upon the extensive use of Multi Layer Perceptrons (MLPs) to encode images [34, 42], videos [38], shapes [32, 8] and 3D scenes [40, 21, 3, 26]. Once trained, these models yield continuous representations that can be queried at arbitrary locations in the signal’s input domain. They have had significant impact in view-synthesis [21, 26] and other interpolation problems [27, 45]. Similar approaches have been used for end-to-end differentiable texture mapping [28] and volumetric rendering [29].

Periodic Activations Lapedes and Farber [16] show the earliest use of periodic activations in neural networks. They observe that networks with more than one layer with periodic activations are difficult to train, and often converge to undesired local minima. This problem is formalized further in [31]. For small datasets, Sopena *et al.* [41] show compelling results using sine activations in the first layer and monotonic functions in the others. This is similar to preconditioning the input using a Fourier basis, which was shown to be useful in feature visualization [30] and image synthesis [28]. More recently, Tancik *et al.* [44] proposed Fourier Feature Networks (FFN), where they encode the MLP’s input into a high-dimensional space using a random sampling of Fourier basis functions. Concurrently, Sitzmann *et al.* [38] showed that, with careful network initializations, sine activations can be used in all layers. They demonstrate regressions of small, single images and videos, as well as more complex shapes. However, as we show in Section 4, these networks struggle with larger datasets, or individual instances when the complexity is increased. In Figure 2, we show how quality degrades with these networks as we regress a progressively higher-resolution image, or a longer video. Our work lifts these limitations by exploiting locality, and introduces an effective modulation mechanism to enable generalization.

Instance-conditioned Implicit Functions A major limitation of current implicit representations is that they need to be optimized for each test signal individually, unlike more established models that only require a forward pass at test time, having been trained on large datasets. Building implicit models with similar generalization properties is typically done via a conditioning mechanism, using latent variables. Conditioning by concatenating the latent code with an MLP’s spatial input coordinates has been successful in signed-distance fields regression tasks [8, 32]. Schwarz *et al.* [35] use the same strategy for MLPs encoding radiance fields, although with limited resolution. In Section 3.2, we show why this conditioning-by-concatenation approach is inadequate for MLPs with sine activations. Conditional hypernetworks [11] achieves similar goals. A hypernetwork estimates all the parameters of a hyponetwork [38] from the latent code. Hypernetworks are prohibitively expensive, in both compute and memory, thereby limiting the resolution of the reconstructed signal. We propose a new approach to modulate the implicit function based on a conditioning variable, inspired from attention mechanisms [46]. In a concurrent work, Chan *et al.* [5] estimate all modulation parameters using a hypernetwork to enable generalization for implicit functions.

Local Implicit Functions Our work bears resemblance with other methods which use local representations for shapes [4, 25] and radiance fields [18]. Chen *et al.* show that a similar representation can also be used for image superresolution [7]. More recently, Takikawa *et al.* [43] and Martel *et al.* [24] showed that shapes and images can be encoded at high resolution using hierarchical-implicit representations with varying levels of detail.

3. Method

We introduce a novel parameterization of neural functions approximating signals defined on a Euclidean input domain \mathbb{R}^n . Our pipeline is illustrated in Figure 1. Our model can simultaneously encode a large number of functions, each of which is summarized into a latent code. The latent codes are processed by a modulation network, which conditionally modulates the activations of a synthesis network, that acts as a template for the functional mapping (§ 3.1). We show that this new architecture is crucial to generate functional representation of multiple signals with a single model, a task where our approach significantly outperforms previous work that use a single MLP (§ 3.2). This ability to generalize lets use compute functional representations for signals with much higher-resolution than previously possible, and with much higher fidelity. For this, we decompose the input signals into local tiles, each represented with a latent code (§ 3.3). The latent codes can either be estimated from the discrete input by a convolutional encoder (§ 3.4.1), or they can be optimized simultaneously with the model parameters




Figure 3: **Sine vs. ReLU synthesizer.** MLPs using ReLU activations (*left*) fail to reconstruct high-frequency components of the target signal; here a 6×6 grid of Perlin [33] texture patches, where the horizontal and vertical frequencies increase from the top-left to bottom-right corner (*right*). Sine-MLPs (*middle*) are able to reconstruct elements from a wider frequency spectrum.

(§ 3.4.2), as in [4].

3.1. Modulated Periodic Activations

Concretely, we define our functional representation as a continuous conditional mapping.

$$f_\theta : \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}^m. \quad (1)$$

f_θ is a neural network, with parameters θ , $d = 256$ is the dimension of the latent space. In the case of images, we use $n = 2$ for pixel coordinates, $m = 3$ for the RGB color values. We find the optimal model parameters θ by minimizing a domain-specific reconstruction loss $\mathcal{L}(f_\theta(x_i; z), y_i)$ on a dataset of signal values $y_i \in \mathbb{R}^m$, sampled at coordinate $x_i \in \mathbb{R}^n$. The signal is encoded into a latent code z . Our network architecture has two components—a synthesis network (§ 3.1.1), and a modulation network (§ 3.1.2).

3.1.1 Synthesis Network

The synthesis network defines a continuous function from the spatial coordinates of a discrete signal like an image to its output domain (*e.g.* color). It is a composition of K hidden layers, with hidden features h_1, \dots, h_K . Each layer uses a periodic-nonlinear activation function and is defined recursively as:

$$h_i = \alpha_i \odot \sin(w_i h_{i-1} + b_i), \quad (2)$$

with $w_i \in \mathbb{R}^{d_i \times d_{i-1}}$ and $b_i \in \mathbb{R}^{d_i}$, the learnable weights and biases for layer i , and $\alpha_i \in \mathbb{R}^{d_i}$ a modulation variable discussed in Section 3.2. We set $h_0 := x \in \mathbb{R}^n$ to be the input coordinates. The sine function is applied pointwise and \odot denotes element-wise multiplication. Sine activations have proven to be beneficial in modeling high-frequency signals [38]. We confirm this by comparing our synthesis network to an alternative that uses ReLU activations in Figure 3.

3.1.2 Modulation Network

We modulate the activations of the synthesis network with a second MLP using ReLU activations, which acts on the latent code z corresponding to the target signal. It is defined recursively as:

$$h'_0 = \text{ReLU}(w'_0 z + b'_0), \quad (3)$$

$$\alpha_{i+1} = h'_{i+1} = \text{ReLU}(w'_{i+1} [h'_i \ z]^T + b'_{i+1}), \quad (4)$$

where $\text{ReLU}(\cdot) = \max(0, \cdot)$ and w'_i, b'_i are the weights and biases of the MLP. We establish an explicit relationship between w' and z , by feeding in z at every layer in the modulation network in the form of a skip connection. As can be seen from Equations 4 and 2, the latent codes z can modulate the amplitude of the sine activations of each hidden layer in the synthesis network, through the modulation parameters α_i . Furthermore, expanding Equation 2, we get

$$h_i = \alpha_i \odot \sin(w_i \cdot (\alpha_{i-1} \odot \sin(\dots)) + b_i), \quad (5)$$

which shows the latent codes also indirectly control the *frequency* and *phase shift* of the sinusoids in subsequent layers. Figure 4 illustrates the expressivity of our modulation mechanism visually.

3.2. Expressivity of Modulation

A simpler alternative to using a separate modulation network would be to concatenate the latent codes and the input coordinates, and use the resultant vector as a *single* input to the synthesis network. This strategy has shown to be fruitful for ReLU-based synthesis networks for encoding signed distance fields [32]. However, we find it consistently fails with sine activations (see Section 4.1 for details).

In this alternative conditioning mechanism, the network takes the concatenation $[x, z]$ as input, so the first layer can be rewritten as:

$$h_1 = \sin(w_{1,x}x + w_{1,z}z + b_1), \quad (6)$$

where $w_{1,x}$ and $w_{1,z}$ are submatrices of w_1 , corresponding to x and z respectively. The latent codes z can therefore only act as a *phase shift*, $w_{1,z}z$, on the first layer. This severely limits expressivity, in contrast to our model, where the latent code z modulates the amplitude, frequency *and* phase-shift of the functional representation at *all* layers, via the α_i . Figure 4 illustrates the difference in expressivity between our model and a concatenation-based MLP.

3.3. Local Functional Representations

The ability to generalize to more than one signal gives us an additional opportunity. Rather than computing a single neural function for the entire signal, we decompose the domain into a regular grid, and calculate a local continuous representation for each tile (illustrated in Figure 5). Concretely, we assign each tile a latent code z_i , so that the entire




Figure 4: **Our modulator improves expressivity.** Conditioning a sine-synthesizer by concatenation does not yield much control over the MLP’s internal feature maps, thereby limiting the expressivity of the latent space. Here, we show feature maps obtained at the second layer of a randomly initialized synthesis network conditionally modulated on four different latent vectors z_i . Concatenating z_i with the input (left) only changes the phase of the signals at the second layer. Using our modulator sub-network (right) provides much more control over the internal feature maps, with variations in phase, amplitude and frequency.

signal is represented by a codebook $\{z_i\}$, and the corresponding neural functions $\{f_\theta(\cdot, z_i)\}$, whose first argument is the normalized local coordinates in the tile $x \in [0, 1]^n$.

Continuity at tile boundaries. In practice, to eliminate visual discontinuities at the tile boundaries, the images are split into a set of overlapping tiles. When evaluating the continuous representation, the contribution of overlapping tiles is weighted n -linearly according to the distance between the point and the tile centers (Fig. 5).

3.4. Training procedure

We present two modes of training of our model. In the auto-encoder setting, (§ 3.4.1), the latent codes are estimated using a discrete encoder. In the auto-decoder configuration (§ 3.4.2), the latent codes are randomly initialized and optimized with the network parameters as in [32].




Figure 5: **Continuity at tile boundaries.** We partition the target signals into uniform grids of overlapping tiles. Here, an image (left) is partitioned into 4 tiles (middle) with overlapping areas. During inference the tiles are blended using an n -linear (bilinear here) weighting scheme (right). The colors represent the resultant contribution of each tile at a given position. (Best viewed in color)

3.4.1 Auto-encoder

Unless otherwise specified, we use our model in an auto-encoder configuration. Auto-encoding lets us to build a continuous representation, from discrete input signals, using an auxiliary encoder network (shown in Figure 1). This could be useful in spatial super-resolution (images, videos), frame interpolation (videos), or reconstruction problems from sparse samples (lightfields, compression).

3.4.2 Auto-decoder configuration

In the auto-decoder configuration, we jointly optimize the network parameters θ , and the latent codes for all the training signals. That is, we do not use the optional encoder of Figure 1. We use this configuration in our shape reconstruction experiments, as proposed by [32]. After training, we obtain a functional representation for new, unseen test signals by sampling a new latent code z for the unseen signal, and optimizing it with the same objective used during training, but this time keeping the network parameters θ constant. We initialize all latent codes $z \sim \mathcal{N}(0, s^2)$ as Gaussian random vectors with $s = 10^{-2}$.

4. Experiments

We demonstrate two classes of experiments, on three domains (images, videos, 3D shapes). First, we demonstrate the generalization capabilities of the proposed model (§ 4.1) in a *global* setting. That is, we compute functional representations for many discrete signals, each of which represented (as a whole) by a latent code (i.e., without the tiling procedure described in Section 3.3). Second, we show how our model can be used to learn *local* functional representations of discrete signals (§ 4.2), with high reconstruction quality. In this set of experiments, each signal is defined using a latent codebook (one code per tile of the input signal). We also show our model can be applied to other multi-domain tasks, such as image relighting (§ 4.3), where the function’s input is a 2D pixel coordinate and a 3D lighting direction.

We compare to state-of-the-art MLP-based functional baselines, illustrated in Figure 6, together with our model. These are:

ReLU/FFN A standard MLP with two inputs—latent code and sample coordinates. In case of FFN [44], the sample coordinates are transformed using a random Fourier matrix with scale σ .

SIREN+ A single MLP with sine activations adapted from [38], with an additional input for the latent code z , concatenated with the coordinates x .

HyperNet-SIREN A SIREN whose weights are conditionally predicted using a hypernetwork [11], as described




Figure 6: **Comparison Baselines.** In all our experiments, we compare with state-of-the-art baselines that use MLPs to encode visual signals. For consistency, we keep the size of the MLPs constant across all models.

Method	PSNR \uparrow	PSNR- $2\times \uparrow$	Params.
ReLU	28.31	24.91	1.1M
SIREN+ [38]	20.15	19.19	1.1M
FFN ($\sigma = 10$) [44]	26.48	4.22	1.2M
FFN ($\sigma = 1$) [44]	27.87	24.37	1.2M
HyperNet-SIREN [11, 38]	26.70	24.54	72M
Ours	29.42	25.49	1.3M
Ours - Large Conv.	29.64	26.05	17M

Table 1: **Comparison of generalization ability on CelebA.** We train conditional functional on 200K 32×32 images from the CelebA [20] dataset, with a conditional latent code per image. We show PSNR on the test set at $1\times$ resolution to show reconstruction accuracy, and at $2\times$ resolution to show the interpolation behaviors of the functions. We mark the best result in yellow. Without careful tuning of σ , FFN produces discontinuous functions (red). HyperNetwork requires $\sim 60\times$ more parameters (orange).

in [38]. Instead of a modulator sub-network, a hyper-network takes the latent code z as input, and predicts all the parameters of the synthesis network (the w_i, b_i).

4.1. Global Functional Representation

Images We run our image experiments on the CelebA [20] and CIFAR-10 [15] datasets separately. We use our model in an auto-encoder setting (§ 3.4.1). A convolutional encoder estimates a latent code for each image. From the latent codes, we decode a functional representation for each image. All the images are resampled to 32×32 resolution for training. The parameters of the modulator, synthesizer and encoder are trained simultaneously to minimize the sum of a reconstruction loss,

$$\mathcal{L}(f_\theta(x; z_i), y) = \|f_\theta(x; z_i) - y\|_2^2, \quad (7)$$

We train all the models for 1000 epochs. Our train/test splits contain 167K/33K and 60K/10K images respectively. We use 128×128 center-crops for CelebA and entire image for CIFAR-10 as ground truth. The images are resampled to 64×64 and 32×32 using bicubic sampling.

Method	PSNR \uparrow	PSNR- $2\times$ \uparrow
ReLU	23.65	25.18
SIREN+ [38]	<i>n/a</i>	<i>n/a</i>
FFN ($\sigma = 10$) [44]	22.75	5.75
FFN ($\sigma = 1$) [44]	24.97	25.51
HyperNet-SIREN [11, 38]	18.07	19.16
Ours	25.73	27.05

Table 2: **Comparison of generalization ability on CIFAR-10.** We also show results on CIFAR-10 [15] dataset, which contains more diverse images than the CelebA dataset. The numbers are reported for the test set with 10K images. PSNR values for $2\times$ resolution are computed against ground-truth images at $1\times$ resolution upsampled using bilinear interpolation. SIREN+ fails to converge on this dataset.

We evaluate generalization by sampling f_θ at pixel centers, at the input image resolution 32×32 ($1\times$) and computing the PSNR. Additionally, we evaluate continuity by sampling f_θ more finely, at 64×64 ($2\times$) resolution, and comparing to ground-truth resampled to the same resolution; we do *not* train the models with these higher resolution targets. Table 1, summarizes our result on the CelebA dataset, and Table 2 on CIFAR-10 [15]. SIREN+ struggles with generalization, and in the case of CIFAR-10, does not even converge. We hypothesize this is due to the higher image variability in CIFAR-10, in comparison to CelebA where faces are aligned. We observe a similar behavior with HyperNet-SIREN. As shown in Tables 1 and 2, the scaling parameter σ of FFN [44] is critical for continuity: PSNR $2\times$ drops with the recommended $\sigma = 10$ value. Since in the case of HyperNet-SIREN, the last layer predicts *all* the parameters of the hyponetwork, it makes the last layer highly over-parameterized. This leads to *slow* training, *unstable* convergence and *inefficient* memory usage. We show reconstructions on CelebA test images in Figure 7.

Shapes Generative modeling of 3D shapes has recently been driven by implicit neural representations [32, 8] trained to regress a shape’s signed distance field (SDF), by sampling discrete locations in the 3D space. The shape can be reconstructed from the learned SDF using sphere tracing [19] or marching cubes. We show that our model is a powerful replacement for the conditional ReLU-MLPs typically used for this application; it can encode SDFs more accurately. For this experiment, we sample 500K points for each shape in the *cars* category of ShapeNet [6]. Half these points are sampled close to the surface, the remaining are randomly sampled inside the unit sphere encompassing the shapes [12]. We use a similar training objective as in case of images (Eq. 7), but we replace the L_2 loss with an L_1 penalty in the fidelity term. Following [32], all conditional models are trained as auto-decoders for this experiment




Figure 7: **Functional representation of CelebA images.** We show qualitative comparisons with other neural representations which use periodic activations. SIREN [38] with the latent codes concatenated with the input pixel coordinates fails to recover sharp details. The HyperNet [11] (HypN.) conditioned SIREN [38] has $\sim 60\times$ more parameters.




Figure 8: **Implicit shape encoding.** Our model used as an auto-decoder with a latent code z_i assigned to each shape from the *cars* category in the ShapeNet [6] dataset. Compared to DeepSDF [32] which uses a single MLP with ReLU activations, our model is able to recover finer high-frequency details. Note the highlighted regions.

(§ 3.4.2). Table 3 shows quantitative comparisons in terms of bi-directional Chamfer distance, computed between the ground-truth shapes and the reconstructions. We show renderings in Figure 8. Compared to DeepSDF [32], we produce higher-quality reconstructions, with finer details. As for images, we found SIREN+ does not converge. In this comparison, we do not include recent improvements that are orthogonal to our contribution, e.g., improvements to the spatial sampling [9], training procedure [37] or loss functions [10]. These improvements would benefit our method as well as the baselines.

Videos For videos, we train our model on ~ 90 K videos from the Vimeo-90k *septuplet* dataset [47]. Each video is 7 frames long and has a spatial resolution of 448×256 . During training, we randomly crop $32 \times 32 \times 7$ tiles from the videos. We use a 3D convolutional encoder to predict the latent codes. At test time, the videos are structured in a grid and each tile is reconstructed with the estimated latent

Method	Chamfer Distance		
	Median ↓	Mean ↓	Std ↓
DeepSDF [32]	0.00284	0.00363	0.00559
DeepSDF + SIREN [32, 38]	<i>n/a</i>	<i>n/a</i>	<i>n/a</i>
DeepSDF + FFN [32, 44]	0.00399	0.00519	0.00757
Ours	0.00230	0.00273	0.00285

Table 3: **Comparison of generalization ability on ShapeNet.** Our model can be used in an auto-decoder setting, where the latent codes for each shape are optimized simultaneously with the network parameters. We implicitly represent 3514 car shapes from the ShapeNet [6] dataset and report the bi-directional Chamfer distance to the ground truth shapes.

Method	PSNR ↑
ReLU	23.06
SIREN+ [38]	<i>n/a</i>
FFN [44]	19.38
HyperNet-SIREN [11, 38]	23.61
Ours	25.28

Table 4: **Comparison of generalization ability on Vimeo-90k.** For videos, we observe trends similar to the image experiment. We report reconstruction PSNR on 7K videos from the Vimeo-90k dataset.

Method	Local	PSNR ↑	Overfit	Time ↓
ReLU		18.94	✓	~60m
SIREN [38]		22.88	✓	~60m
FFN [44]		28.48	✓	~60m
Ours-ReLU	✓	34.73		13s
Ours	✓	38.03		15s

Table 5: **Encode high-res images faster.** We use our model as a functional auto-encoder to encode images from Div2K [1] dataset. The global methods are *overfit* on each image separately and do not generalize. Our local methods are pre-trained and numbers are reported on a test set.

code. The train-test split is used as provided in the dataset. We show quantitative comparisons in Table 4.

4.2. Local Functional Representations

Images Our dual-MLP model can also be used to generalize to high-resolution implicit functions. We train our model on 100 images from Div2K [1]. Each image has a long-side resolution of 2K and split into 32×32 overlapping tiles. We found 32×32 tile size to provide good reconstruction accuracy as well as good interpolation properties. The total number of tiles in the training set is $\sim 2M$. Each tile is encoded using our method as an auto-encoder. At test time we sample 50 unseen images, and encode them using the trained model. Since other baselines do not generalize, we train a separate MLP for each of the images individually for

Method	Local	Chamfer Distance ($\times 10^{-5}$)	
		Scene A (9 shapes)	Scene B (9 Shapes)
ReLU [32]		2.16	3.61
SIREN [32, 38]		2.00	7.36
FFN [32, 44]		1.93	3.54
Ours-ReLU	✓	1.37	8.37
Ours	✓	1.32	2.40

Table 6: **Implicit encoding of large 3D scenes.** Our method can be used to approximate a signed-distance-field for large scenes. The global methods are overfit on each scene separately. Our local models are trained on Scene A and then generalized to Scene B by freezing the weights of the model and optimizing the latent codes.

global methods (*i.e.* ReLU, SIREN, FFN). Reconstruction PSNR at $1 \times$ resolution is reported in Table 5. Additionally, we perform an ablation on our model by using a standard ReLU MLP with our local parameterization.



Figure 9: **SIREN struggles with high-frequencies in temporal domain.** Our model faithfully reconstructs high-frequencies in both the spatial and temporal dimensions. We show the last frame of a video segment before a jump-cut to a different shot. This is a sharp transition in time, which SIREN struggles to model accurately.

Shape For this experiment, we collect 18 high-resolution ($\sim 2M$ triangles) shapes from the the ThreedScans project [17]. These shapes are split into two Scenes A and B, each of which have 9 shapes. We compute a ground truth signed-distance-field (SDF) as in the global shapes experiment (§ 4.1) for supervision. We use our model in an auto-decoder configuration. We train it on $48 \times 48 \times 48$ voxels extracted from Scene A, and we evaluate reconstruction accuracy on Scene B, where we only optimize the latent codes. For our global baselines, we overfit the models individually for each scene. We extract meshes from the learned neural SDFs using marching cubes [22], and report the chamfer-distance from the ground truth to the reconstructions in Table 6. In Figure 10 we show Scene A renders using all the baselines and our method.

Video Similar to experiments shown in [38], we encode high-resolution videos using our method. We use 6 videos (pexels.com) with 1920×1080 resolution and downsample them to 640×256 . Each video is split into a grid of $32 \times 32 \times 7$ tiles for the local model. The reconstruction PSNR is reported in Table 7. Our model is pre-trained on




Figure 10: **Encoding large 3D scenes.** Our method can be used to encode signed-distance-fields (SDFs) for large scenes (example shown on the left). The local representation enables the MLP f_θ to distribute its capacity sparsely in the scene. As a result, the estimated SDF is not overly *smooth* like we observe in the case of global methods.

Method	Local	PSNR \uparrow	Overfit	Time \downarrow
ReLU		19.28	✓	~15hr
FFN [44]		20.87	✓	~15hr
SIREN [38]		25.19	✓	~15hr
Ours-Generalized	✓	25.21		~1m

Table 7: **Encoding high-res videos faster.** We use our pre-trained model on Vimeo-90k [47] to encode high-res videos faster. To achieve a similar reconstruction accuracy, overfitting other models take upto 15 hours.

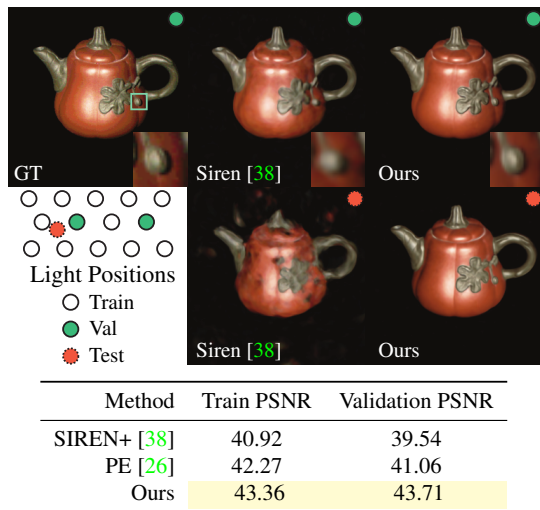


Figure 11: **Image-based relighting using MLPs.** Using light direction and image coordinates concatenated as input to a SIREN fails to reconstruct cast-shadows and specularities faithfully (insets). Our method of using conditional modulation reconstructs at higher fidelity with smoother interpolation in the light domain.

Vimeo-90k [47] (§ 4.1) and tested on the collected videos. It is able to achieve similar reconstruction accuracy as to the one obtained by previous methods while being $\sim 1000\times$ faster. We found that SIRENs [38] struggle to reconstruct high-frequency content for complex and varied video signals, in both the spatial and time dimensions as shown in Figure 9.

4.3. Image-based Relighting

We perform image-based relighting, where the input is a pixel coordinate $x \in \mathbb{R}^2$ and a lighting direction is taken as the latent code $z \in \mathbb{R}^3$. The output is radiance $y \in \mathbb{R}^3$. We use 90 cropped and aligned 300×300 images of a *real* scene, captured in an OLAT (One-Light-At-a-Time) setup [36]. The 90 images are divided into a (80, 10) train-val split. We train MLPs to learn the function $y_{ij} = f_\theta(x_{ij}; z_i)$, where y_{ij} is the radiance at x_{ij} with z_i as the light direction. A standard L_2 loss is used to train the parameters θ . At test time, we pick arbitrary z_i 's within reasonable bounds and use f_θ to synthesize images with the corresponding light direction. For comparison, we implement SIREN+ [38] and a ReLU MLP with positional encodings (PE) [26] (we found them to work better than FFN) with naïve concatenation of domains, *i.e.* $[x_{ij}, z_i]$ as the input. Figure 11 shows a scene relit using light-directions not present in the training set. Our method reconstructs light-dependent effects like shadows and specular highlights with higher fidelity. We find SIREN+ to struggle with interpolation in this experiment.

5. Conclusion

We propose a novel method for representing signals using multi-layer perceptrons (MLPs). We show that partitioning the signal domain into tiles simplifies the signal locally. This leads to representing images, videos and shapes using MLPs with high-quality reconstructions. MLPs with ReLU activations fail to reconstruct high-frequency components of the signals. Instead, we use sine activations which we show to work with a wider frequency spectrum. Using local models requires MLPs to be conditioned on latent codes. We show that concatenating latent codes with the input hinders expressivity. Our method uses a dual-MLP architecture instead. The proposed model also generalizes to multiple instances of these signals. Our local parameterization is general enough to be applied in other applications that use implicit neural functions [26].

6. Acknowledgements

This work was supported in part by ONR grants N000142012529, N000141712687, Adobe, the Ronald L. Graham Chair, NSF grant 1730158, NSF CAREER 1751365, Amazon Research Award and the UC San Diego Center for Visual Computing.



References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017. 7
- [2] Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2565–2574, 2020. 1
- [3] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Deep reflectance volumes: Relightable reconstructions from multi-view photometric images. *arXiv preprint arXiv:2007.09892*, 2020. 2
- [4] Rohan Chhabra, Jan Eric Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. *arXiv preprint arXiv:2003.10983*, 2020. 2, 3
- [5] Eric Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *arXiv*, 2020. 3
- [6] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 6, 7
- [7] Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. *arXiv preprint arXiv:2012.09161*, 2020. 3
- [8] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019. 1, 2, 3, 6
- [9] Thomas Davies, Derek Nowrouzezahrai, and Alec Jacobson. Overfit neural networks as a compact shape representation. *arXiv preprint arXiv:2009.09808*, 2020. 1, 6
- [10] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020. 6
- [11] David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016. 3, 5, 6, 7
- [12] Zekun Hao, Hadar Averbuch-Elor, Noah Snaveley, and Serge Belongie. Dualsdf: Semantic shape manipulation using a two-level representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7631–7641, 2020. 6
- [13] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, and Thomas Funkhouser. Local implicit grid representations for 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6001–6010, 2020. 2
- [14] Sylwester Kloczek, Łukasz Maziarka, Maciej Wołczyk, Jacek Tabor, Jakub Nowak, and Marek Śmieja. Hypernetwork functional image representation. In *International Conference on Artificial Neural Networks*, pages 496–510. Springer, 2019. 1
- [15] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 5, 6
- [16] Alan Lapedes and Robert Farber. Nonlinear signal processing using neural networks: Prediction and system modelling. Technical report, 1987. 2
- [17] Oliver Laric. Three d scans. 7
- [18] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33, 2020. 3
- [19] Shaohui Liu, Yinda Zhang, Songyou Peng, Boxin Shi, Marc Pollefeys, and Zhaopeng Cui. Dist: Rendering deep implicit signed distance function with differentiable sphere tracing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2019–2028, 2020. 6
- [20] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 5
- [21] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751*, 2019. 1, 2
- [22] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics*, 21(4):163–169, 1987. 7
- [23] Dhruv Mahajan, Ira Kemelmacher Shlizerman, Ravi Ramamoorthi, and Peter Belhumeur. A theory of locally low dimensional light transport. In *ACM SIGGRAPH 2007 papers*, pages 62–es. 2007. 2
- [24] Julien N.P. Martel, David B. Lindell, Connor Z. Lin, Eric R. Chan, Marco Monteiro, and Gordon Wetzstein. Acorn: Adaptive coordinate networks for neural representation. *ACM Trans. Graph. (SIGGRAPH)*, 2021. 3
- [25] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 3
- [26] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *arXiv preprint arXiv:2003.08934*, 2020. 1, 2, 8
- [27] Hans-Peter Seidel, Mojtaba Berman, Karol Myszkowski, and Tobias Ritschel. X-fields: Implicit neural view-, light- and time-image interpolation. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia 2020)*, 39(6), 2020. 2
- [28] Alexander Mordvintsev, Nicola Pezzotti, Ludwig Schubert, and Chris Olah. Differentiable image parameterizations. *Distill*, 2018. <https://distill.pub/2018/differentiable-parameterizations>. 2
- [29] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3504–3515, 2020. 2
- [30] Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. Feature visualization. *Distill*, 2017. <https://distill.pub/2017/feature-visualization>. 2

- [31] Giambattista Parascandolo, Heikki Huttunen, and Tuomas Virtanen. Taming the waves: sine as activation function in deep neural networks. 2016. [2](#)
- [32] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [33] Ken Perlin. Improving noise. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 681–682, 2002. [3](#)
- [34] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. [2](#)
- [35] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *arXiv preprint arXiv:2007.02442*, 2020. [1](#), [3](#)
- [36] Boxin Shi, Zhe Wu, Zhipeng Mo, Dinglong Duan, Sai-Kit Yeung, and Ping Tan. A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3707–3716, 2016. [8](#)
- [37] Vincent Sitzmann, Eric Chan, Richard Tucker, Noah Snavely, and Gordon Wetzstein. Metasdf: Meta-learning signed distance functions. *Advances in Neural Information Processing Systems*, 33, 2020. [1](#), [6](#)
- [38] Vincent Sitzmann, Julien NP Martel, Alexander W Bergman, David B Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *arXiv preprint arXiv:2006.09661*, 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [39] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2019. [1](#)
- [40] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems*, pages 1121–1132, 2019. [2](#)
- [41] Josep M Sopena, Enrique Romero, and Rene Alquezar. Neural networks with periodic and monotonic activation functions: a comparative study in classification problems. 1999. [2](#)
- [42] Kenneth O Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8(2):131–162, 2007. [2](#)
- [43] Towaki Takikawa, Joey Litalien, Kangxue Yin, Karsten Kreis, Charles Loop, Derek Nowrouzezahrai, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Neural geometric level of detail: Real-time rendering with implicit 3D shapes. 2021. [3](#)
- [44] Matthew Tancik, Pratul P Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *arXiv preprint arXiv:2006.10739*, 2020. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [45] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019. [2](#)
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [3](#)
- [47] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision (IJCV)*, 127(8):1106–1125, 2019. [6](#), [8](#)