

Learning Spatially Collaged Fourier Bases for Implicit Neural Representation

Jason Chun Lok Li*, Chang Liu*, Binxiao Huang, Ngai Wong

Department of Electrical and Electronic Engineering, The University of Hong Kong
 {jasonlcl, lcon7, huangbx7}@connect.hku.hk, nwong@eee.hku.hk

Abstract

Existing approaches to Implicit Neural Representation (INR) can be interpreted as a global scene representation via a linear combination of Fourier bases of different frequencies. However, such universal basis functions can limit the representation capability in local regions where a specific component is unnecessary, resulting in unpleasant artifacts. To this end, we introduce a learnable spatial mask that effectively dispatches distinct Fourier bases into respective regions. This translates into collaging Fourier patches, thus enabling an accurate representation of complex signals. Comprehensive experiments demonstrate the superior reconstruction quality of the proposed approach over existing baselines across various INR tasks, including image fitting, video representation, and 3D shape representation. Our method outperforms all other baselines, improving the image fitting PSNR by over 3dB and 3D reconstruction to 98.81 IoU and 0.0011 Chamfer Distance.

Introduction

Implicit Neural Representation (INR) has attracted increased attention lately as an effective framework for encoding complex signals in diverse fields, encompassing image fitting, 3D shape representation, and novel view synthesis (Sitzmann et al. 2020; Saragadam et al. 2023; Mildenhall et al. 2021). Using INR, discontinuous signals such as images, videos, and 3D point clouds can be represented by a lightweight neural network, which outputs corresponding data values given the query coordinates. Hence, the neural network transforms the original discontinuous data into a continuous latent mapping function, which sheds new light on several downstream tasks, for example, data compression, 3D shape reconstruction, and super-resolution (Sitzmann et al. 2020; Saragadam et al. 2023).

Albeit the success of INR, conventional methods usually model the target signal through a linear combination of universal basis functions spanning various frequencies (Sitzmann et al. 2020; Tancik et al. 2020; Fathy et al. 2020), which are usually waveform Fourier features. Consequently, these methods inevitably introduce the same strength of each basis feature to the whole space, even when the component

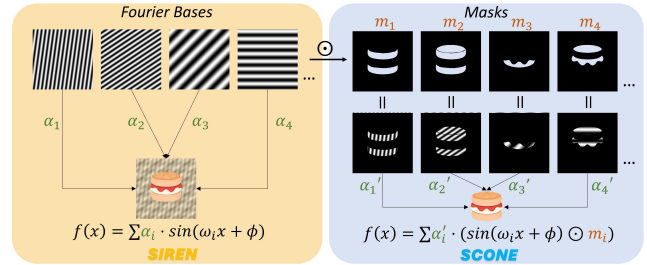


Figure 1: (Left) traditional INR such as SIREN can be viewed as a linear combination of Fourier features of different frequencies, where the globally added Fourier features are universally applied to the whole frame even for some regions where these frequencies are unnecessary. (Right) The proposed SCONE utilizes spatial masks for each Fourier feature and collages the frequency patches to form a more precise reconstruction.

is *not* necessary for certain regions. Even though such unnecessary components can be counteracted by combining a large number of different frequency bases, this strains the network and limits its representability. As also pointed out by (Dou et al. 2023), the global combination manner results in larger model sizes and/or sub-optimal performance, especially for local regions where existing methods can introduce undesirable artifacts and fuzzy details.

In view of such limitation, this paper presents a simple yet effective architecture, dubbed **S**patially **C**ollaged **c**oordinate **N**etworks (SCONE). In contrast to existing approaches employing a combination of global Fourier features, SCONE incorporates learnable spatial masks to allocate specific Fourier bases to their corresponding regions. As illustrated in Fig. 1, SCONE learns a spatial mask for each Fourier feature to avoid adding the wave component in regions where it is unnecessary. Consequently, SCONE distinguishes between the distinct tasks of learning geometry-dependent masks and Fourier features, permitting the Fourier features to be spatially and adaptively decoupled. By adopting a different paradigm of spatial collaging instead of global overlaying, the proposed technique facilitates a more precise and granular representation of intricate signals through the efficient assembly of mask-cropped Fourier

*These authors contributed equally.

patches. Additionally, it has been observed that SCONE converges faster than baseline methods, suggesting the gradual assembly of Fourier patches can provide better learning dynamics. In short, SCONE effectively overcomes the limitations of existing global basis functions and offers superior reconstruction quality across a wide spectrum of tasks.

To demonstrate the efficacy of SCONE, we extensively evaluate SCONE on various INR tasks, including image fitting, video representation, and 3D shape representation, and compare its performance against existing baselines. The experimental results *consistently* show the superiority of SCONE in terms of reconstruction quality, overtaking its competing counterparts. For image fitting, SCONE outperforms all baselines by a significant margin over 3dB PSNR. It also achieves a new state-of-the-art (SOTA) performance of 98.81% IoU and 0.0011 Chamfer Distance on 3D shape representation. Our contributions are threefold:

- We identify the limitation in existing methods of globally combining Fourier bases and propose SCONE to decouple the bases spatially.
- SCONE is the first-ever coordinate network that explicitly collages patches of different frequencies.
- Comprehensive experiments are conducted to evaluate SCONE on various tasks. Our method consistently sets new SOTA on all tasks.

Related Work

Implicit Neural Representations (INRs), also known as coordinate networks, can encode continuous signals in a compact form using constant memory, depending only on model capacity, regardless of the resolution of their discrete representation (e.g., pixels or voxel grids). They take low-dimensional coordinates as inputs and are trained to learn the input-output mapping in an end-to-end fashion via gradient-based optimization. INRs have been applied to represent a wide variety of signals, including, but not limited to, images (Saragadam et al. 2023; Lindell et al. 2022; Shekarforoush et al. 2022), occupancy networks (Chen and Zhang 2019; Mescheder et al. 2019), signed distance functions (Sitzmann, Zollhöfer, and Wetzstein 2019; Park et al. 2019), as well as solving differential equations (Sitzmann et al. 2020; Fathony et al. 2020).

INRs first gained attention in the field of novel view synthesis (Mildenhall et al. 2021), which uses a multilayer perceptron (MLP) to represent a 5D continuous radiance field (viz. neural radiance field), which maps radiance directions and spatial coordinates to view-dependent color and density. The model is trained using 2D images from various views, and a differential volume rendering equation is employed to render the MLP output for novel views. Positional encoding is utilized (Vaswani et al. 2017) to represent high-frequency scene content effectively.

Following that, Ref. (Tancik et al. 2020) systematically shows that the positional encoding adopted is one type of Fourier-feature mapping that enables a coordinate-based MLP to overcome its inductive spectral bias of failing in learning high-frequency content. With the use of Neural Tangent Kernel (NTK) (Jacot, Gabriel, and Hongler 2018)

theory, the authors show that Fourier feature mapping can successfully transform the effective NTK of an MLP into a stationary kernel with a tunable bandwidth, thus improving its capability of fitting natural signals such as images. Apart from positional encoding, Gaussian Random Fourier Features (RFFs) are also proposed, which leads to improved reconstruction quality.

However, even when the coordinates are first transformed by positional encoding via Fourier features, the use of ReLU MLPs, whose second derivative is zero, limits the model’s ability to represent the derivatives of a target signal. This constraint has motivated the development of sinusoidal representation networks (SIREN) (Sitzmann et al. 2020). Specifically, SIREN replaces the ReLU activation with periodic sinusoidal nonlinearities, which can accurately represent natural signals and their derivatives. In addition to sine activations, Wavelet Implicit Neural Representation (WIRE) (Saragadam et al. 2023) attempts to utilize a continuous complex Gabor wavelet activation function that is optimally concentrated in space-frequency. This provides appropriate inductive biases for representing images.

On the other hand, instead of stacking layers in a composite manner, multiplicative filter networks (MFNs) (Fathony et al. 2020), a family of INRs with better interpretability, repeatedly apply linear functions of sinusoidal or Gabor filters to the input, which are then element-wise multiplied together. The entire expression can be rewritten as a linear combination of an exponential number of sinusoidal or Gabor wavelet bases. Two works have been developed to extend the theoretical understanding and practicality of MFNs. BACON (Lindell et al. 2022) is a variant of MFNs that has band-limited properties, resulting in an analytical Fourier spectrum. Whereas RMFN (Shekarforoush et al. 2022) introduces residual connections and a novel initialization scheme, allowing more control of the frequency spectrum at each stage of optimization and enabling coarse-to-fine estimation.

By taking a step back, a recent work (Yüce et al. 2022) attempts to provide an overarching explanation for the underlying theory behind INRs, including SIREN and MFNs. Results from harmonic analysis and NTK theory are utilized to demonstrate that most INR families are analogous to structured signal dictionaries. This structure enables INRs to express signals with an exponentially increasing frequency support using a number of parameters that grow linearly with depth.

Two approaches closely related to our work are GaborNet instantiation of MFN and WIRE. They enhance the locality of INRs by employing spatially compact wavelet filters to improve their ability to represent local regions. But rather differently, our proposed SCONE accurately represents local structures *without* relying on the computationally intensive wavelet transform (which may involve complex number computations). We achieve this by collaging Fourier patches with the help of positional-dependent masks. A multi-channeled mask is applied to Fourier bases of varying frequencies in each stage, resulting in Fourier patches. These patches are then mixed using a linear layer. The resulting output is subsequently used as an input to generate

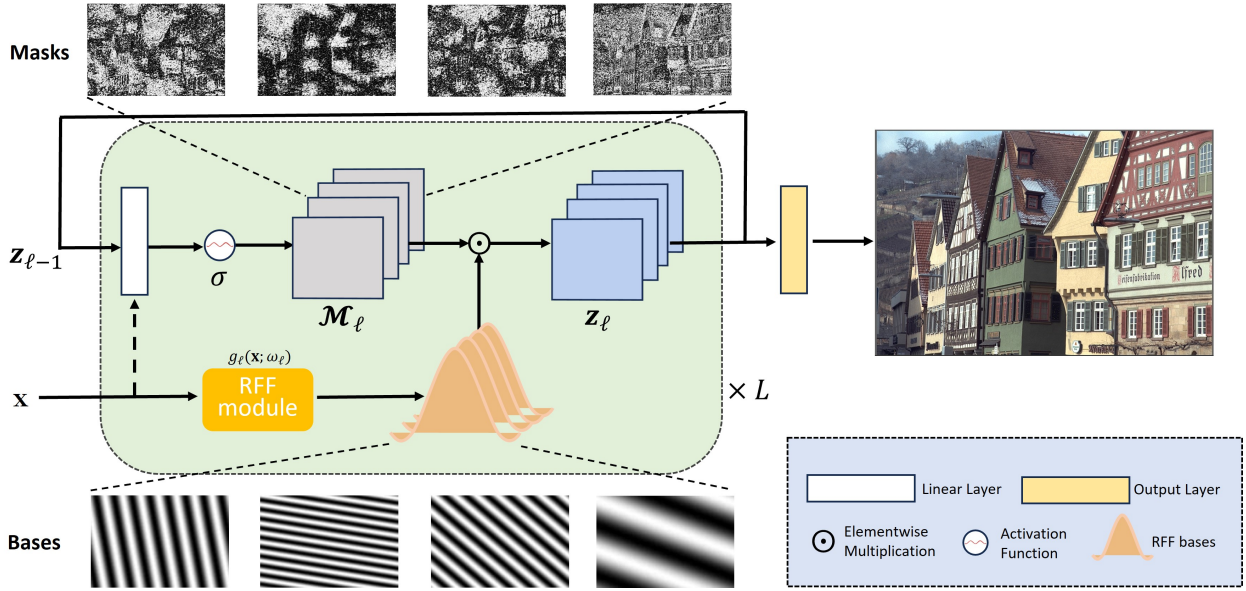


Figure 2: An overview of SCONE architecture. The dashed arrow denotes that the initial hidden features z_0 are from the coordinates x . σ is the activation function with output range $[0, 1]$. Independent Fourier bases in each level are masked to form Fourier patches that are then spatially collaged by the next linear layer. This enables the assignment of global bases to their proper spatial positions, which are then reused in more than one region across the signal.

the mask for the next stage. The whole process can be interpreted as making a collage by assembling different textures, as illustrated in Fig. 1. Unlike Gabor filters or wavelet activations, where each wavelet waveform can only appear in specific spatial locations, the adaptive mask in SCONE can effectively dispatch global bases to multiple locations in the signal.

Preliminary

An INR or a coordinate network is a method that uses a neural network to encode continuous signals. The network is trained to learn the complex relationships between the coordinates and their corresponding features, in order to accurately encode the signal continuously. For example, an image can be described by a function that maps each pixel’s coordinates $\mathbf{x} = (x, y)$ in \mathbb{R}^2 to its corresponding RGB value $\mathbf{y} = (r, g, b)$ in \mathbb{R}^3 . With a single image datapoint as a collection of coordinate and feature pairs $\mathbf{d} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$, the INR fits a neural network $f_\theta : \mathbb{R}^m \rightarrow \mathbb{R}^c$ with parameters θ to the datapoint by minimizing the loss function:

$$\mathcal{L}(\theta, \mathbf{d}) = \sum_{i=1}^n \|f_\theta(\mathbf{x}_i) - \mathbf{y}_i\|_2 \quad (1)$$

The conventional parameterization of an L -layer INR exhibit a general form:

$$\begin{aligned} \mathbf{z}_0 &= \gamma(\mathbf{x}) \\ \mathbf{z}_\ell &= \delta(\mathbf{W}_\ell \mathbf{z}_{\ell-1} + \mathbf{b}_\ell), \quad \ell = 1, \dots, L-1 \\ \mathbf{z}_L &= \mathbf{W}_L \mathbf{z}_{L-1} + \mathbf{b}_L \end{aligned} \quad (2)$$

where $\mathbf{z}_\ell \in \mathbb{R}^{D_\ell}$ represents the post-activation latent feature, $\mathbf{W}_\ell \in \mathbb{R}^{D_\ell \times D_{\ell-1}}$ represents the weight matrix, and

$\mathbf{b}_\ell \in \mathbb{R}^{D_\ell}$ represents the bias vector. $\gamma(\mathbf{x})$ denotes the input mapping, also known as positional encoding, and $\delta(\cdot)$ represents the nonlinear activation function.

In particular, both Fourier feature networks (FFNs) (Tancik et al. 2020) and SIREN (Sitzmann et al. 2020) models can be expressed in such unified form. FFNs utilize Fourier feature mappings to lift the input coordinates into frequency space, followed by MLPs with ReLU nonlinearities. According to (Tancik et al. 2020), using Gaussian RFFs $\gamma(\mathbf{x}) = [\cos(2\pi\mathbf{B}\mathbf{x}), \sin(2\pi\mathbf{B}\mathbf{x})]^T$, where $\mathbf{B} \in \mathbb{R}^{m \times d}$ is a random matrix sampled from $\mathcal{N}(0, \sigma^2)$, yields the best results. On the other hand, SIREN can also be interpreted as an MLP with sine nonlinearity and the first layer being a feature mapping with $\gamma(\mathbf{x}) = \sin(\omega_0(\mathbf{W}_0\mathbf{x} + \mathbf{b}_0))$. Both σ^2 in FFNs and ω_0 in SIREN are global hyperparameters controlling the signal frequency of the input mapping.

Spatially Collaged Coordinate Networks

This section elucidates SCONE. As described earlier, SCONE can be interpreted as a spatial-aware combination of Fourier bases with different frequencies. As shown in Fig. 2, we iteratively generate spatial masks from previous-stage latent features, which will be applied for collaging Fourier Features with different frequencies. Finally, the SCONE output is derived as a linear combination of the collaged RFFs.

Random Fourier Features Given the input coordinates \mathbf{x} , SCONE first normalizes the inputs into a range of $[-1, 1]^2$ and transforms them into position encoding using RFFs:

$$g_\ell(\mathbf{x}; \omega_\ell) = \sin(\omega_\ell(\mathbf{B}_\ell \mathbf{x})) \quad (3)$$

where $\mathbf{B}_\ell \sim U(-\sqrt{1/2}, \sqrt{1/2})$ is a random matrix. This results in L independent Fourier bases, one for each layer

controlled by the hyperparameter ω_ℓ as depicted in the bottom part of Fig. 2.

Spatial Mask During the mask generation process, the input coordinates first pass through a linear mapping with a scaling factor as described as $\mathbf{z}_0 = \omega_0(\mathbf{W}_0\mathbf{x})$. Similar to SIREN, ω_0 can be seen as a global parameter to control the frequency of the mask-generating process.

Iteratively, SCONE generates spatial masks for collaging RFF bases, thereby enabling various RFF combinations at different regions of the image. Specifically, given last layer latent feature $\mathbf{z}_{\ell-1}$, the mask generating process takes place in following form:

$$\begin{aligned} \mathcal{M}_\ell &= \sigma(\mathbf{W}_{\ell-1}\mathbf{z}_{\ell-1} + \mathbf{b}_{\ell-1}) \\ \mathbf{z}_\ell &= \mathcal{M}_\ell \circ g_\ell(\mathbf{x}; \omega_\ell), \quad \ell = 1, \dots, L-1 \end{aligned} \quad (4)$$

where σ is an activation function and \circ denotes element-wise multiplication and the weight matrix $\mathbf{W}_\ell \in \mathbb{R}^{D_\ell \times D_{\ell-1}}$ is initialized using the same initialization scheme as SIREN.

By applying a corresponding spatial mask on each RFF, the frequency will only be added to certain region of the image instead of globally as in SIREN (Sitzmann et al. 2020). Moreover, different from existing methods such as WIRE (Saragadam et al. 2023), instead of manipulating the activation function to shrink the RFFs into a closed region, our method allows one RFF to be used at multiple regions with irregular or complex shapes. Without extra supervision signal, SCONE learns to generate spatial masks aligning with the target image distribution in a highly interpretable manner. We visualize the generated masks in the upper part of Fig. 2. We also provide a visualization of hidden layers’ activations in SCONE, as well as those of other models in the later Section for comparison.

In practice, the proposed model progressively learns simple to complex spatial masks. Taking advantage of the increasingly elaborate spatial masks, albeit using the same number of RFFs, it is observed that outputs from each iteration appear to contain recognizable shapes in deep layers. In the end, a fully connected layer generates the final output image as a linear combination of the last set of collaged Fourier bases.

Experiments

In this section, we evaluate our SCONE for various tasks of image fitting, video representation, and 3D shape representation. Our method is proven to surpass all the other baseline methods by a large margin. We also visualize the layer activations for interpretability.

Implementation Details

All models are trained for 10k iterations on Nvidia RTX 3090 GPUs, each equipped with a 24GB memory buffer. The training process utilizes the Adam optimizer (Kingma and Ba 2014), with the parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$, and employs the Mean Squared Error (MSE) loss without weight decay. Additionally, a cosine learning rate scheduler is applied, with a minimum learning rate of $1e-6$. We evaluate SCONE against several baselines, including ReLU

MLP with Gaussian RFF encoding (Tancik et al. 2020), SIREN (Sitzmann et al. 2020), GaborNet instantiation of MFN (Fathony et al. 2020), and WIRE (Saragadam et al. 2023). The implementation of all codes is carried out using the PyTorch (Paszke et al. 2019) framework, and the baselines are based on official codes released by authors of the respective models. For the below experiments, we choose \sin^2 as the activation function σ due to its desirable output range of $[0, 1]$ for generating soft masks and periodic property. In a later section, we conduct a thorough analysis of the impact of different activation functions, which demonstrates that SCONE is robust to the choice of activation function, not limited to \sin^2 . The excellent performance of SCONE can be attributed to the design of collaging Fourier features. Please refer to the ablation study for more details.

Image Fitting

The image representation task is performed on selected images from the Kodak dataset (Eastman Kodak Company 1999), in which each image has a resolution of 512×768 or 768×512 . Two widely adopted metrics, namely Peak Signal-to-Noise Ratio (PSNR) and structural similarity index (SSIM) (Wang et al. 2004), are used as evaluation metrics. To align the model sizes, all models have 256 hidden units in each layer, except for SCONE and WIRE, which has 250 and 360 hidden units, respectively. We find that naively increasing the number of layers in WIRE leads to inferior performance. To fix the number of parameters to around 200k, both SIREN and SCONE have 5 layers, while the others have 4. We conduct hyperparameter sweeps for each model to select the best hyperparameters and initial learning rate. Fig. 3 demonstrates that SCONE outperforms all baselines by a significant margin ($> 3\text{dB}$). The ReLU MLP with Gaussian RFF exhibits high-frequency noise on relatively flat surfaces. MFN displays severe ringing artifacts around the edges. WIRE struggles with learning color, resulting in poor saturation (as seen in the grass). Both SIREN and SCONE achieve good reconstruction, with SCONE producing sharper and finer textures (as seen in the house next to the tower). We also compare the convergence speed of SCONE with other baselines in Fig. 4. It can be seen that MFN has a fast initial convergence speed but quickly becomes flat, failing to learn high-quality representations. In our setting with 200k parameters, SIREN performs fairly well and outperforms all other baselines, including WIRE. This has also been pointed out by WIRE (Saragadam et al. 2023), which states that SIREN catches up in terms of reconstruction quality when scaling up, making it a strong baseline to beat. Our approach, SCONE, not only converges the fastest but also achieves the highest final PSNR.

Video Representation

We further extend our approach from image fitting to video representation tasks. In this case, the input coordinates are now in $[-1, 1]^3$, including an additional temporal domain. All models have the same number of layers and hidden units as in the image-fitting task described above. Following (Sitzmann et al. 2020), the video¹ is first downsampled to 512×512 with 300 RGB frames. In each iteration, a batch

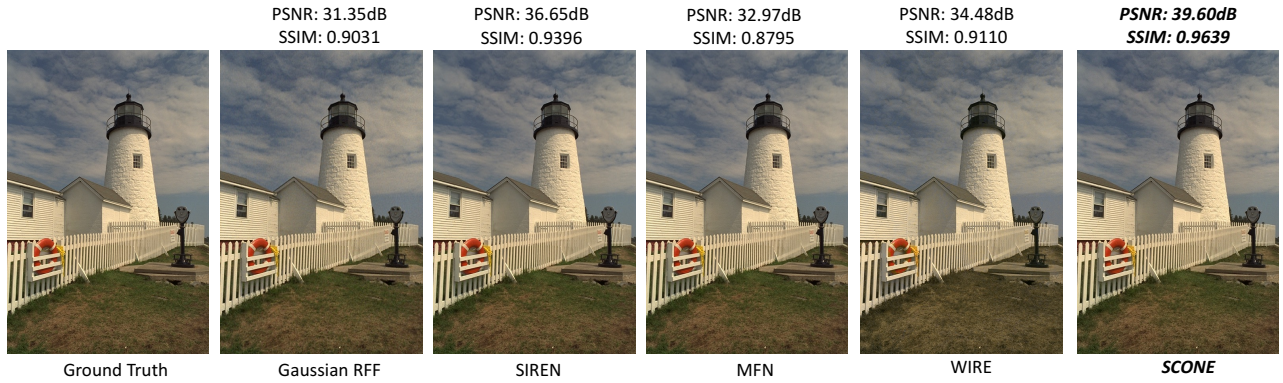


Figure 3: The result on image representation task. Kodak19 is selected for comparison.

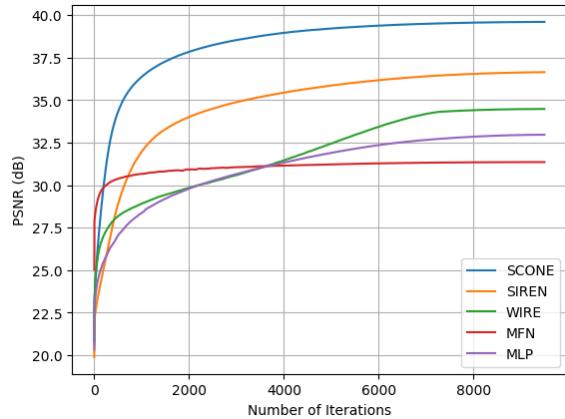


Figure 4: PSNR vs. Training Iterations: The plot depicts the PSNR against the training iterations for the image fitting task (Kodak09). With the aid of geometry-dependent spatial masks, SCONE converges at a much faster rate and achieves the highest reconstruction accuracy.

of 262144 random data points is sampled. The mean and standard deviation of PSNR and SSIM across all frames are reported. Table 1 clearly demonstrates that SCONE can be extended to more complex signals, such as video, where data lies in the spatial-temporal domain, achieving the most accurate reconstruction in terms of PSNR and SSIM.

3D Shape Representation

For the 3D shape representation task, we use the Stanford 3D scan dataset². The dataset contains point clouds with normals of different 3D objects. Following (Lindell et al. 2022), for each step during training, we extract a batch of points from the original point cloud and add a random noise, which consists of two different levels of perturbation, following Laplace distribution with scales of $1e-1$ and $1e-3$, respectively. The normals for these perturbed points are ap-

¹<https://www.pexels.com/video/the-full-facial-features-of-a-pet-cat-3040808>

²<http://graphics.stanford.edu/data/3Dscanrep/>

Method	#Params (K)	PSNR(dB)	SSIM
Gaussian RFF	198.1	21.74 ± 0.71	0.35 ± 0.02
SIREN	199.1	26.92 ± 0.87	0.70 ± 0.02
MFN	206.3	24.02 ± 0.73	0.49 ± 0.02
WIRE	196.1	26.90 ± 0.78	0.67 ± 0.02
SCONE (ours)	194.0	27.41 ± 0.83	0.70 ± 0.01

Table 1: Results on video representation task. The mean \pm standard deviation across all frames is reported. SCONE achieves the best performance in terms of PSNR. It ties with SIREN in SSIM, but with smaller variance across frames.

proximated by taking the average of normals of the nearest three points in the original point cloud. After that, we use the approximated normals to calculate the signed distance to the surface as SDF (signed distance function) labels. MSE loss is used as supervision criteria. To ensure a fair comparison, we adjust SCONE and all the other baseline methods to have similar model sizes of around 400K parameters. Specifically, we use 8 layers for SCONE and empirically initialize the RFF with ω being $[70, 70, 60, 50, 40, 40, 30, 30]$. We train all the methods with the learning rate following their original implementations, and for 10K steps. Except for Gaussian RFF (Tancik et al. 2020) and WIRE (Saragadam et al. 2023), which are slower to converge, we train them for 20K steps. For evaluation, we uniformly sample $512 \times 512 \times 512$ points in the cube $[-0.5, 0.5]^3$. The networks then predict the signed distance for these points, outputting a tensor of shape $512 \times 512 \times 512$, each element denotes the signed distance to the 3D shape plane.

IoU and Chamfer Distance (CD) metrics are leveraged to quantify the reconstruction results. To compute the IoU scores, we transform the output tensor into an occupancy grid by setting the voxel value to 0 if the corresponding SDF value > 0 (i.e., outside the 3D shape), and 1 if $SDF \leq 0$ (i.e., inside the 3D shape). IoU is then computed between the ground truth and the predicted occupancy grids. For CD, we use the Trimesh library to convert the SDF tensor into a 3D mesh and take all the vertices from a mesh as a point cloud for calculating CD.

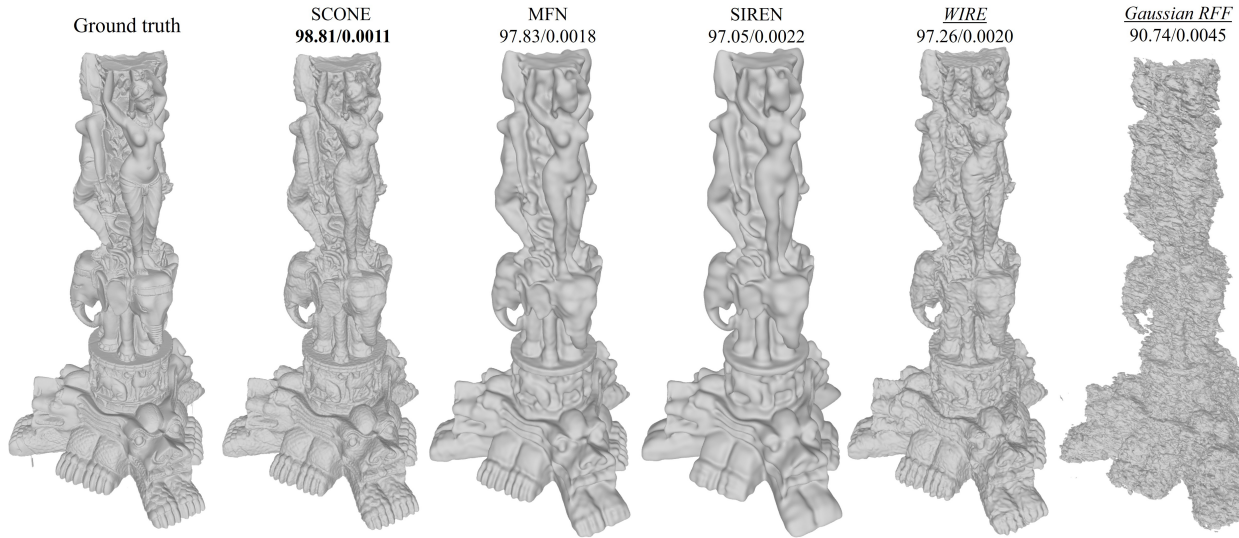


Figure 5: 3D shape reconstruction results. All the methods listed are adjusted to have similar model sizes of 400K parameters. All models are trained for 10K steps, except for Gaussian RFF and WIRE cannot converge until 20K steps. SCONE captures more details than other baseline methods, benefiting from collaged Fourier features’ design. Best viewed with zoom-in.

Qualitative results are shown in Fig. 5. As can be observed, our method captures details much better than the other baseline methods, yielding a 98.81% IoU and 0.0011 CD, 0.98% and 63% higher than the second best. This observation validates our motivation to use collaged Fourier features to improve the model’s expressiveness and ability to represent local details. Nonetheless, it is worth stressing the baselines’ results are slightly different from the values reported in previous studies because we aligned all the methods under the same experiment settings of (1) model sizes of 400K and training step of 10K for fairness, (2) random sampling points for training, (3) IoU and CD calculating method. The only two exceptions are WIRE (Saragadam et al. 2023), and Gaussian RFF (Tancik et al. 2020), who cannot converge and learn any shape within 10K steps. Therefore, for these two methods, we train them for a long period of 20K steps. However, even being trained with double iterations, the IoU and CD results are still relatively low compared to other methods.

SCONE Layer Visualization

Following WIRE (Saragadam et al. 2023), we visualize the layer activations of INRs being trained on the Siemens star test image for 10k iterations, which contain all spatial frequencies and orientations. We compare SCONE with WIRE and GaborNet MFN to highlight their differences in achieving space-frequency locality. Here, each model has 4 hidden layers with a dimension size of 64. The PSNR of SCONE, WIRE, and MFN are 38.47dB, 34.10dB, and 28.26dB, respectively. Fig. 6 visualizes the outputs of each hidden layer of the respective models. It is apparent that WIRE and MFN,

both utilizing Gabor wavelets, produce sparse activations that are compact in space. In contrast, each base of SCONE is global in nature, spatially spanning the entire channel. The spatial mask then dynamically assigns Fourier bases to the appropriate spatial locations based on the geometric structure of the image. This leads to potentially higher utilization of the channels.

Ablation Study

In this section, we conduct ablation studies with the image fitting task using the Kodak dataset to address two questions: (1) Does the performance gain truly come from spatially collaging Fourier bases? (2) Is SCONE robust to different choices of activation functions?

Comparison with Random Binary Mask

To justify the effectiveness of our design choice, we compare our method to the baseline that uses randomly generated binary masks instead of learnable soft masks in SCONE. The baseline method can be formulated as:

$$\mathbf{z}_\ell = \mathcal{M}_\ell \circ \mathbf{W}_{\ell-1} \mathbf{z}_{\ell-1} \circ g_\ell(\mathbf{x}; \omega_\ell)$$

where \circ denotes element-wise multiplication, $\mathcal{M}_\ell \in \{0, 1\}^{D_\ell \times H \times W}$ is the random binary mask (RBM) for the ℓ -th layer, D is the channel dimension, and H and W are the spatial sizes of the image. RBMs are randomly initialized for each layer and fixed during training.

As shown in Fig. 7, the RBM method has much lower PSNR scores compared to SCONE and SIREN, which justifies that SCONE’s improvement in performance indeed originates from collaging Fourier bases with learnable spatial masks.

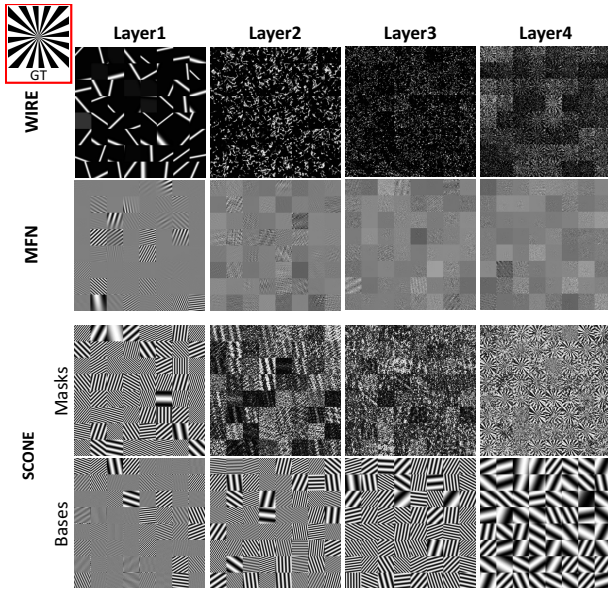


Figure 6: Visualization of activations from hidden layers. The Siemens star test image is used to highlight the mechanisms that differentiate SCONE from other INRs with wavelet filters or nonlinearities. Each layer consists of 64 units arranged in an 8×8 grid, with output of each unit having the same size as the input image (256×256). For WIRE, only the real parts of the activation outputs are shown.

Effect of the Choice of Activation Functions

In this section, we explore the effects of different activation functions on SCONE for generating spatial masks. As mentioned above, a soft mask with values ranging from 0 to 1 is expected. Several activation functions, such as *sigmoid*, $\exp(-x^2)$ and \sin^2 satisfies this property. For other activation functions such as \sin , \cos , \tanh , whose outputs are in finite range, we normalize the output to $[0, 1]$ by min-max normalization (e.g., $\frac{1}{2} \tanh(\cdot) + 0.5$).

Fig. 7 compares the performances of different activation function choices on the Kodak dataset. The prefix ‘N-’ denotes min-max normalization, and the ‘-L’ postfix indicates that the input is multiplied by a trainable scaler before applying the activation function, allowing for an adjustable input scale. The x -axis in Fig. 7 represents the data samples sorted by their average PSNR score over all methods, from easy to difficult, for better visualization.

It is evident in Fig. 7 that SCONE is not limited to specific activation functions, but is robust to different activation choices. SCONE with all activation choices yields significantly higher PSNR scores than the baseline SIREN. Additionally, the differences between various activation choices of SCONE are also small. In contrast, SIREN and RBM exhibit a lot of fluctuation from sample to sample, while the SCONE family shows a more consistent performance across simple and difficult images. Furthermore, the learnable activations (indicated by the ‘-L’ postfix) can bring small improvements to some activation functions (e.g., *sigmoid*, *tanh*), although these improvements are not substantial.

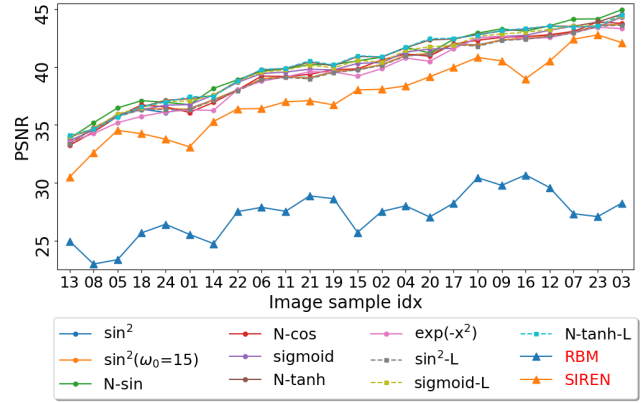


Figure 7: Ablation studies for activation function σ on the Kodak dataset. Image samples are sorted by difficulty (average PSNR) for better visualization. All variants of SCONE yield similarly high PSNR and outperform the baselines.

In the above experiment section, we select \sin^2 due to its favorable features: (i) output range being $[0, 1]$ acting as a soft mask; (ii) periodic function with smooth and continuous gradient, which is intuitively connected with Fourier features. However, it is important to note that our SCONE is not limited to \sin^2 and is highly robust to different choices of activation functions. In fact, for SCONE, all activation functions, including \sin^2 , yield similarly good results and significantly outperform the baseline methods. These results validate our motivation for collaging Fourier bases and demonstrate that SCONE exhibits good generality and robustness across different activation choices.

Conclusion

This paper presents Spatially Collaged Coordinate Networks (SCONE) as a novel solution to overcome the shortcomings of existing INRs whose universal basis functions can lead to subpar performance and undesirable artifacts. In contrast, SCONE exploits learnable spatial masks to assign particular Fourier bases to their respective regions, thus enabling a more accurate representation of complicated signals. In-depth analyses of SCONE’s performance on various INR tasks, such as image fitting, video representation and 3D shape representation, demonstrate excellent reconstruction quality. Specifically, SCONE achieves over 3dB PSNR improvement in image fitting and establishes new SOTA performance in 3D shape representation for general INRs with 99% IoU and 0.0011 Chamfer Distance, surpassing all competing methods by a large margin.

Acknowledgements

This work was supported in part by the Theme-based Research Scheme (TRS) project T45-701/22-R and in part by the General Research Fund (GRF) project 17209721 of the Research Grants Council (RGC), Hong Kong SAR.

References

- Chen, Z.; and Zhang, H. 2019. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5939–5948.
- Dou, Y.; Zheng, Z.; Jin, Q.; and Ni, B. 2023. Multiplicative Fourier Level of Detail. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1808–1817.
- Eastman Kodak Company. 1999. Kodak lossless true color image suite. <http://r0k.us/graphics/kodak/>. [Accessed 14-08-2023].
- Fathony, R.; Sahu, A. K.; Willmott, D.; and Kolter, J. Z. 2020. Multiplicative filter networks. In *International Conference on Learning Representations*.
- Jacot, A.; Gabriel, F.; and Hongler, C. 2018. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lindell, D. B.; Van Veen, D.; Park, J. J.; and Wetzstein, G. 2022. Bacon: Band-limited coordinate networks for multi-scale scene representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16252–16262.
- Mescheder, L.; Oechsle, M.; Niemeyer, M.; Nowozin, S.; and Geiger, A. 2019. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4460–4470.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Park, J. J.; Florence, P.; Straub, J.; Newcombe, R.; and Lovegrove, S. 2019. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 165–174.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Kopf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*, 8024–8035. Curran Associates, Inc.
- Saragadam, V.; LeJeune, D.; Tan, J.; Balakrishnan, G.; Veer-araghavan, A.; and Baraniuk, R. G. 2023. Wire: Wavelet implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18507–18516.
- Shekarforoush, S.; Lindell, D.; Fleet, D. J.; and Brubaker, M. A. 2022. Residual multiplicative filter networks for multiscale reconstruction. *Advances in Neural Information Processing Systems*, 35: 8550–8563.
- Sitzmann, V.; Martel, J.; Bergman, A.; Lindell, D.; and Wetzstein, G. 2020. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33: 7462–7473.
- Sitzmann, V.; Zollhöfer, M.; and Wetzstein, G. 2019. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *Advances in Neural Information Processing Systems*, 32.
- Tancik, M.; Srinivasan, P.; Mildenhall, B.; Fridovich-Keil, S.; Raghavan, N.; Singhal, U.; Ramamoorthi, R.; Barron, J.; and Ng, R. 2020. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33: 7537–7547.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Yüce, G.; Ortiz-Jiménez, G.; Besbinar, B.; and Frossard, P. 2022. A structured dictionary perspective on implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19228–19238.