




Learning Local Implicit Fourier Representation for Image Warping

Jaewon Lee¹, Kwang Pyo Choi², and Kyong Hwan Jin¹*

¹ Daegu Gyeongbuk Institute of Science and Technology (DGIST), Korea

² Samsung Electronics, Korea

ljw3136@dgist.ac.kr, kp5.choi@samsung.com, kyong.jin@dgist.ac.kr

Abstract. Image warping aims to reshape images defined on rectangular grids into arbitrary shapes. Recently, implicit neural functions have shown remarkable performances in representing images in a continuous manner. However, a standalone multi-layer perceptron suffers from learning high-frequency Fourier coefficients. In this paper, we propose a local texture estimator for image warping (LTEW) followed by an implicit neural representation to deform images into continuous shapes. Local textures estimated from a deep super-resolution (SR) backbone are multiplied by locally-varying Jacobian matrices of a coordinate transformation to predict Fourier responses of a warped image. Our LTEW-based neural function outperforms existing warping methods for asymmetric-scale SR and homography transform. Furthermore, our algorithm well generalizes arbitrary coordinate transformations, such as homography transform with a large magnification factor and equirectangular projection (ERP) perspective transform, which are not provided in training. Our source code is available at <https://github.com/jaewon-lee-b/ltew>.

Keywords: Image warping, Implicit neural representation, Fourier features, Jacobian, Homography transform, Equirectangular projection (ERP)

1 Introduction

Our goal is to deform images defined on rectangular grids into continuous shapes, referred to as image warping. Image warping is widely used in various computer vision and graphic tasks, such as image editing [11, 12], optical flow [43], image alignment [7, 21, 23, 38], and omnidirectional vision [2, 3, 5, 14, 28]. A conventional approach [21, 24] applies an inverse coordinate transformation to interpolate the missing RGB value in the input space. However, interpolation-based methods cause jaggging and blurring artifacts in output images. Recently, SRWarp [42] paved the way to reshape images with high-frequency details by adopting a deep single image super-resolution (SISR) architecture as a backbone.

SISR is a particular case of image warping [42]. A goal of SISR is to reconstruct a high-resolution (HR) image from its degraded low-resolution (LR) counterpart. Recent lines of research in solving SISR are to extract deep feature maps using advanced architectures [8, 29, 30, 32, 48, 51, 52] and upscale them to HR images at the end of a network [10, 19, 27, 39, 50]. Even though deep SISR

* Corresponding author.




Fig. 1. Implicit neural representation for image warping.

methods reconstruct visually clear HR images, directly applying them to our problem is limited since each local region of a warped image is stretched with different scale factors [17]. By reconsidering the warping problem as a spatially-varying SR task, SRWarp [42] shed light on deforming images with sharp edges. However, interpolation-based SRWarp shows limited performance in generalizing to a large-scale representation which is out of training range.

Recently, implicit neural functions have attracted significant attention in representing signals, such as image [10, 27], video [9], signed distance [36], occupancy [33], shape [22], and view synthesis [34, 41], in a continuous manner. A multi-layer perceptron (MLP) parameterizes such an implicit neural representation [40, 44] and takes coordinates as an input. Inspired by the recent implicit function success, LIIF [10] well generalizes to a large-scale rectangular SR beyond a training distribution. However, one shortcoming of implicit neural functions [37, 44] is that a standalone MLP with ReLUs is biased towards learning low-frequency content. To alleviate this spectral bias problem, Local Texture Estimator (LTE) [27] estimates Fourier features for an HR image from its LR counterpart motivated by Fourier analysis. While LTE achieved arbitrary-scale rectangular SR with high-frequency details, LTE representation fails to evaluate a frequency response for image warping due to its spatially-varying nature.

Given an image $\mathbf{I}^{\text{IN}} : X \mapsto \mathbb{R}^3$ and a differentiable and invertible coordinate transformation $f : X \mapsto Y$, we propose a local texture estimator for image warping (LTEW) followed by an implicit neural function representing $\mathbf{I}^{\text{WARP}} : Y \mapsto \mathbb{R}^3$, as in Fig. 1. Our algorithm leverages both Fourier features estimated from an input image and the Jacobian of coordinate transformation. In geometry, the determinant of the Jacobian indicates a local magnification ratio. Hence, we multiply spatially-varying Jacobian matrices to Fourier features for each pixel before our MLP represents \mathbf{I}^{WARP} . Furthermore, we point out that a spatially-varying prior for pixel shape is essential in enhancing a representational power of neural functions. The pixel shape described by orientation and curvature is numerically computed by gradient terms of given coordinate transformation.

We demonstrate that our proposed LTEW with a deep SISR backbone [48, 51, 52] surpasses existing warping methods [24, 42, 46] for both upscaling and homography transform. While previous warping techniques [42, 46] employ convolution and polynomial interpolation as a resampling module, our LTEW-based implicit neural function takes continuous coordinates as an input. Therefore, our proposed algorithm effectively generalizes in continuously representing \mathbf{I}^{WARP} , especially for homography transforms with a substantial magnification factor

($\times 4 \times 8$), which is not provided during a training phase. We further pay attention to omnidirectional imaging (ODI) [2, 3, 5, 14, 28] to verify the generalization ability of our algorithm. With the rapid advancement in virtual reality (VR), ODI has become crucial for product development. Equirectangular projection (ERP) is widely used in imaging pipelines of a head-mounted display (HMD) [20]. As a result of projection from spherical grids to rectangular grids, pixels are sparsely located near high latitudes. Since the proposed LTEW learns spatially-varying properties, our method qualitatively outperforms other warping methods in projecting perspective without extra training.

To summarize, the contributions of our work include:

- We propose a continuous neural representation for image warping by taking advantage of both Fourier features and spatially-varying Jacobian matrices of coordinate transformations.
- We claim that a spatially-varying prior for pixel shape described by orientation and curvature is significant in improving the representational capacity of the neural function.
- We demonstrate that our LTEW-based implicit neural function outperforms the existing warping methods for upscaling and homography transform, and unseen coordinate transformations.

2 Related Works

Image warping Image warping is a popular technique for various computer vision and graphics tasks, such as image editing [11, 12], optical flow estimation [43], and image alignment [7, 21, 23, 38]. A general technique for image warping [21] is finding a spatial location in input space and applying an interpolation kernel to calculate missing RGB values. Even though an interpolation-based image warping is differentiable and an easy-to-implement framework, an output image suffers from jaggging and blurring artifacts [42]. Recently, SRWarp [42] proposed an arbitrary image transformation framework by interpreting an image warping task as a spatially-varying SR problem. SRWarp shows noticeable performance gain in arbitrary SR, including homography transform using an adaptive warping layer. However, the generalization ability of SRWarp is limited for unseen transformations, like a homography transform with a large magnification factor.

Implicit neural representation (INR) Motivated from the fact that neural network is a universal function approximator [18], INR is widely applied to represent continuous-domain signals. Conventionally, the memory requirement for data is quadratically (2D) or cubically (3D) proportional to signal resolution. In contrast, INR is a memory-efficient framework to store continuous signals since storage size is proportional to the number of model parameters rather than signal resolution. Recently, local INR [10, 22, 27] has been proposed to enhance the spatial resolution of input signals in an arbitrary manner. By using both feature maps from a deep neural encoder and relative coordinates (or *local grid* in [22, 28]), such approaches are capable of generalizing to unseen tasks, which

are not given during training. Inspired by previous works, our proposed LTEW utilizes Fourier features from a deep neural backbone and local grids to represent warped images under arbitrary coordinate transformations.

Spectral bias Early works [37, 44] have shown that INR parameterized by a standalone MLP with a ReLU activation fails to capture high-frequency details of signals. Dominant approaches for resolving this spectral bias problem are substituting ReLUs with a periodic function [40], projecting input coordinates into a high-dimensional Fourier [4, 27, 34, 44] or spline [47] feature space, and multiplying sinusoidal or Gabor filters [16]. Recently, LTE [27] achieved arbitrary-scale SR using INR by estimating Fourier information from an LR image. Unlike previous attempts [34, 44], Fourier feature space in LTE representation [27] is data-driven and characterizes texture maps in 2D space. However, considering a spatially-varying SR issue in image warping, directly applying LTE is limited to characterize the Fourier space of warped images.

Deep SISR After ESPCN [39] has proposed a memory-efficient upsampling layer based on sub-pixel convolution, advanced deep vision backbones, such as residual block [30], densely connected residual block [48, 52], channel attention [51], second-order attention [13], holistic attention [35], non-local network [32], are jointly employed to reconstruct high-quality images. Recently, by taking advantage of inductive bias in self-attention mechanism [15, 31], general-purpose image restoration backbones [8, 29] remarkably outperform convolution-based architectures using a large dataset. Despite their compelling representational power, we have to train and store several models for each scale factor. Current approaches [10, 19, 27, 46, 50] for arbitrary-scale SR are utilizing a dynamic filter network [19, 46], INR [10, 27], or transformer [50]. Unlike the previous arbitrary-scale SR methods, our LTEW represents images under arbitrary coordinate transformations, like homography transform, with only a single network.

Omnidirectional image (ODI) In a new era of VR, omnidirectional vision [2, 3, 5, 14, 28] becomes playing a crucial role in product development. While natural images are represented in the Euclidean space, ODIs are defined on the spherical coordinates. A common methods to project ODIs to the 2D plane is an ERP to be consistent with imaging pipelines in HMD [20]. One limitation of an ERP projected ODIs is that non-uniform spatial resolving power leads to severe spatial distortion near boundaries [28]. To handle this varying pixel densities across latitudes, Deng *et al.* [14] proposed a hierarchically adaptive network. Since our proposed LTEW utilizes Jacobian matrices of given coordinate transformation to learn spatially-varying property for image warping, the distortion caused by varying pixel densities can be safely projected without extra training.

3 Problem Formulation

Given an image $\mathbf{I}^{\text{IN}} : X \mapsto \mathbb{R}^3$, and a differentiable and invertible coordinate transformation $f : X \mapsto Y$, our goal is to formulate an implicit neural repre-

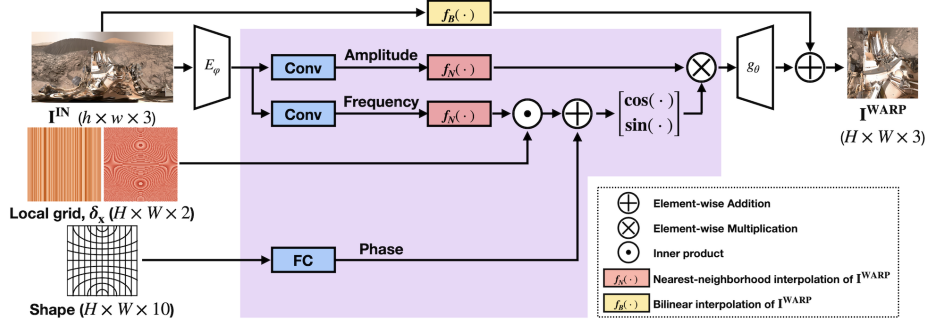


Fig. 2. A flowchart of our proposed Local Texture Estimator Warp (LTEW).

sensation of an $\mathbf{I}^{\text{WARP}} : Y \mapsto \mathbb{R}^3$ for image warping. The set $X \triangleq \{\mathbf{x} | \mathbf{x} \in \mathbb{R}^2\} : [0, h) \times [0, w)$ is an input coordinate space, and $Y \triangleq \{\mathbf{y} | \mathbf{y} = f(\mathbf{x}) \in \mathbb{R}^2\} : [0, H) \times [0, W)$ is an output coordinate space. In practice, warping changes an image resolution to preserve the density of pixels (In Fig. 2: $h \times w \rightarrow H \times W$). Note that neural representation parameterized by an MLP with ReLU activations fails to capture high-frequency details of signals [34, 37, 40, 44]. Recently, LTE [27] achieved arbitrary-scale SR for symmetric scale factors by estimating essential Fourier information. However, concerning the spatially-varying nature of warping problems [6], frequency responses of an input and a deformed image are inconsistent for each location. Therefore, we formalize a Local Texture Estimator Warp (LTEW), a frequency response estimator for image warping, by considering both Fourier features of an input image and spatially-varying property of coordinate transformations. We show that our LTEW is a generalized form of the LTE, allowing a neural representation to be biased towards learning high-frequency components while manipulating images under arbitrary coordinate transformations. In addition, we present shape-dependent phase estimation to enrich the information in output images.

3.1 Learning Fourier information for local neural representation

In local neural representation [10, 22], a neural representation g_θ is parameterized by an MLP with trainable parameters θ . A decoding function g_θ predicts RGB value for a query point $\mathbf{y} = f(\mathbf{x}) \in Y \subseteq \mathbb{R}^2$ as

$$\mathbf{I}^{\text{WARP}}[\mathbf{y}; \Theta] = \sum_{j \in \mathcal{J}} w_j g_\theta(\mathbf{z}_j, \mathbf{y} - f(\mathbf{x}_j)) \quad (1)$$

$$\text{where } \mathbf{z} = E_\phi(\mathbf{I}^{\text{IN}}), \quad (2)$$

$\Theta = [\theta, \varphi]$, \mathcal{J} is a set $\{j | j = [f^{-1}(\mathbf{y}) + [\frac{m}{w}, \frac{n}{h}], [m, n] \in [-1, 1]]\}$ [10, 21], w_j is a local ensemble coefficient [10, 22, 27], $\mathbf{z}_j \in \mathbb{R}^C$ indicates a latent variable for a index j , and $\mathbf{x}_j \in X \subseteq \mathbb{R}^2$ is a coordinate of \mathbf{z}_j .

Recent works [37, 44] have shown that a standard MLP structure suffers from learning high-frequency content. Lee *et al.* [27] modified the local neural

representation in Eq. (1) to overcome this spectral bias problem as

$$\mathbf{I}^{\text{WARP}}[\mathbf{y}; \Theta, \psi] = \sum_{j \in \mathcal{J}} w_j g_\theta(h_\psi(\mathbf{z}_j, \mathbf{y} - f(\mathbf{x}_j))) \quad (3)$$

where h_ψ is a Local Texture Estimator (LTE). LTE ($h_\psi(\cdot)$) contains two estimators; (1) an amplitude estimator ($h_a(\cdot) : \mathbb{R}^C \mapsto \mathbb{R}^{2D}$) (2) a frequency estimator ($h_f(\cdot) : \mathbb{R}^C \mapsto \mathbb{R}^{2 \times D}$). Specifically, an estimating function $h_\psi(\cdot, \cdot) : (\mathbb{R}^C, \mathbb{R}^2) \mapsto \mathbb{R}^{2D}$ is defined as

$$h_\psi(\mathbf{z}_j, \boldsymbol{\delta}_\mathbf{y}) = \mathbf{A}_j \odot \begin{bmatrix} \cos(\pi \langle \mathbf{F}_j, \boldsymbol{\delta}_\mathbf{y} \rangle) \\ \sin(\pi \langle \mathbf{F}_j, \boldsymbol{\delta}_\mathbf{y} \rangle) \end{bmatrix}, \quad (4)$$

$$\text{where } \mathbf{A}_j = h_a(\mathbf{z}_j), \mathbf{F}_j = h_f(\mathbf{z}_j), \quad (5)$$

$\boldsymbol{\delta}_\mathbf{y} = \mathbf{y} - f(\mathbf{x}_j)$ is a local grid, $\mathbf{A}_j \in \mathbb{R}^{2D}$ is an amplitude vector, $\mathbf{F}_j \in \mathbb{R}^{2 \times D}$ indicates a frequency matrix for an index j , $\langle \cdot, \cdot \rangle$ is an inner product, and \odot denotes element-wise multiplication. However, this formulation fails to represent warped images since \mathbf{F}_j is a frequency response of an input image \mathbf{I}^{IN} , which is different from that of a warped image \mathbf{I}^{WARP} [6]. In the following, we generalize LTE by considering a spatially-varying property of coordinate transformations.

3.2 Learning Fourier information with coordinate transformations

We linearize the given coordinate transformation f into affine transformations. A linear approximation of the local grid $\boldsymbol{\delta}_\mathbf{y}$ near a point \mathbf{x}_j is computed as

$$\begin{aligned} \boldsymbol{\delta}_\mathbf{y} &= \mathbf{y} - f(\mathbf{x}_j) = f(\mathbf{x}) - f(\mathbf{x}_j) \\ &= \{f(\mathbf{x}_j) + \mathbf{J}_f(\mathbf{x}_j)(\mathbf{x} - \mathbf{x}_j) + \mathcal{O}(\mathbf{x}^2)\} - f(\mathbf{x}_j) \\ &\simeq \mathbf{J}_f(\mathbf{x}_j)(\mathbf{x} - \mathbf{x}_j) = \mathbf{J}_f(\mathbf{x}_j)\boldsymbol{\delta}_\mathbf{x} \end{aligned} \quad (6)$$

where $\mathbf{J}_f(\mathbf{x}_j) \in \mathbb{R}^{2 \times 2}$ is the Jacobian matrix of coordinate transformation f at \mathbf{x}_j , $\mathcal{O}(\mathbf{x}^2)$ means terms of order \mathbf{x}^2 and higher, and $\boldsymbol{\delta}_\mathbf{x} = \mathbf{x} - \mathbf{x}_j$ is a local grid in input space X . By the affine theorem [6] and Eq. (6), a frequency response \mathbf{F}'_j of a warped image \mathbf{I}^{WARP} near a point $f(\mathbf{x}_j)$ is approximated as follows:

$$\mathbf{F}'_j \simeq \mathbf{J}_f^{-T}(\mathbf{x}_j)\mathbf{F}_j. \quad (7)$$

From Eq. (6) and Eq. (7), we generalize the estimating function in Eq. (4) as:

$$h_\psi(\mathbf{z}_j, \boldsymbol{\delta}_\mathbf{y}, f) = \mathbf{A}_j \odot \begin{bmatrix} \cos(\pi \langle \mathbf{F}'_j, \boldsymbol{\delta}_\mathbf{y} \rangle) \\ \sin(\pi \langle \mathbf{F}'_j, \boldsymbol{\delta}_\mathbf{y} \rangle) \end{bmatrix} \quad (8)$$

$$\simeq \mathbf{A}_j \odot \begin{bmatrix} \cos(\pi \langle \mathbf{J}_f^{-T}(\mathbf{x}_j)\mathbf{F}_j, \mathbf{J}_f(\mathbf{x}_j)\boldsymbol{\delta}_\mathbf{x} \rangle) \\ \sin(\pi \langle \mathbf{J}_f^{-T}(\mathbf{x}_j)\mathbf{F}_j, \mathbf{J}_f(\mathbf{x}_j)\boldsymbol{\delta}_\mathbf{x} \rangle) \end{bmatrix} \quad (9)$$

$$= \mathbf{A}_j \odot \begin{bmatrix} \cos(\pi \langle \mathbf{F}_j, \boldsymbol{\delta}_\mathbf{x} \rangle) \\ \sin(\pi \langle \mathbf{F}_j, \boldsymbol{\delta}_\mathbf{x} \rangle) \end{bmatrix} \quad (10)$$

$$= h_\psi(\mathbf{z}_j, \boldsymbol{\delta}_\mathbf{x}). \quad (11)$$




Fig. 3. Numerical Jacobian matrix and Hessian tensor for shape representation.

When comparing Eq. (4) and Eq. (11), we see that LTEW representation is capable of extracting Fourier information for warped images by utilizing the local grid in input coordinate space $\delta_{\mathbf{x}}$ instead of $\delta_{\mathbf{y}}$.

3.3 Shape-dependent phase estimation

For SISR tasks within symmetric scale factors, the pixel shape of upsampled images is square and spatially invariant. However, when it comes to image warping, pixels in resampled images are able to have arbitrary shapes and spatially-varying, as described in Fig. 3.(a). To address this issue, we represent the pixel shape $\mathbf{s}(\mathbf{y}) \in \mathbb{R}^{12}$ (\mathbb{R}^4 is for pixel orientation, and \mathbb{R}^8 is for pixel curvature) with a gradient of coordinate transformation for a point \mathbf{y} as

$$\mathbf{s}(\mathbf{y}) = [\tilde{\mathbf{J}}_{f^{-1}}(\mathbf{y}), \tilde{\mathbf{H}}_{f^{-1}}(\mathbf{y})] \quad (12)$$

where $[\cdot, \cdot]$ refers to the concatenation after flattening, $\tilde{\mathbf{J}}_{f^{-1}}(\mathbf{y}) \in \mathbb{R}^{2 \times 2}$ and $\tilde{\mathbf{H}}_{f^{-1}}(\mathbf{y}) \in \mathbb{R}^{2 \times 2 \times 2}$ denote the numerical Jacobian matrix, indicating an orientation of pixel, and the numerical Hessian tensor, specifying the degree of curvature, respectively. For shape representation, we apply an inverse coordinate transformation to a query point and its eight nearest points $(\mathbf{y} + [\frac{m}{W}, \frac{n}{H}])$ with $[m, n] \in [-1, 0, 1]$ and compute the difference to calculate numerical derivatives as described in Fig. 3.(b). Let us assume that given coordinate transformation f is in class C^2 , which means $f_{x_1 x_2} = f_{x_2 x_1}$. Hence; we use only six elements in $\tilde{\mathbf{H}}_{f^{-1}}(\mathbf{y})$ for shape representation.

Phase in Eq. (13) includes the information of edge locations or the shape of pixels [27]. Therefore, we redefine the estimating function in Eq. (11) as:

$$h_{\psi}(\mathbf{z}_j, \delta_{\mathbf{x}}, \mathbf{s}(\mathbf{y})) = \mathbf{A}_j \odot \begin{bmatrix} \cos\{\pi(\langle \mathbf{F}_j, \delta_{\mathbf{x}} \rangle + h_p(\mathbf{s}(\mathbf{y})))\} \\ \sin\{\pi(\langle \mathbf{F}_j, \delta_{\mathbf{x}} \rangle + h_p(\mathbf{s}(\mathbf{y})))\} \end{bmatrix} \quad (13)$$

where $h_p(\cdot) : \mathbb{R}^{10} \mapsto \mathbb{R}^D$ is a phase estimator.

Lastly, we add a bilinear interpolated image $\mathbf{I}_B[\mathbf{y}] = \mathbf{f}_B(\mathbf{I}^{\text{IN}}) \simeq \mathbf{I}^{\text{IN}}(f^{-1}(\mathbf{y}))$ to stabilize network convergence and aid LTEW in learning high-frequency details [25]. Thus, the local neural representation of a warped image \mathbf{I}^{WARP} with the proposed LTEW is formulated as follows:

$$\mathbf{I}^{\text{WARP}}[\mathbf{y}; \Theta, \psi] = \mathbf{I}_B[\mathbf{y}] + \sum_{j \in \mathcal{J}} w_j g_{\theta}(h_{\psi}(\mathbf{z}_j, f^{-1}(\mathbf{y}) - \mathbf{x}_j, \mathbf{s}(\mathbf{y}))) \quad (14)$$




Fig. 4. Training strategy for Local Texture Estimator Warp (LTEW).

4 Methods

4.1 Architecture details

Our LTEW-based image warping network consists of an encoder (E_φ), the LTEW (h_ψ , a purple shaded region in Fig. 2), and a decoder (g_θ). An encoder (E_φ) is designed with a deep SR network, such as EDSR [30], RCAN [51], RRDB [48], without upscaling modules. A decoder (g_θ) is a 4-layer MLP with ReLUs, and its hidden dimension is 256. Our LTEW (h_ψ) takes a local grid ($\delta_{\mathbf{x}}$), shape (\mathbf{s}), and feature map (\mathbf{z}) as input, and includes an amplitude estimator (h_a), a frequency estimator (h_f), and a phase estimator (h_p). An amplitude and a frequency estimator are implemented with 3×3 convolution layer having 256 channels, and a phase estimator is a single linear layer with 128 channels. We assume that a warped image has the same texture near point $f(\mathbf{x}_j)$. Hence, we find estimated Fourier information ($\mathbf{A}_j, \mathbf{F}_j$) at \mathbf{x}_j using the nearest-neighborhood interpolation. Then, estimated phase is added to an inner product between the local grid ($\delta_{\mathbf{x}}$) and estimated frequencies, as in Eq. (13). Before the decoder (g_θ) resamples images, we multiply amplitude and sinusoidal activation output.

4.2 Training strategy

We have two batches with a size \mathbf{B} : (1) Image batch $\{I_1, I_2, \dots, I_B\}$, $I_i[Y] \in \mathbb{R}^{H \times W \times 3}$. (2) Coordinate transformation batch $\{f_1, f_2, \dots, f_B\}$, where each $f_i : X \mapsto Y$ is differentiable and invertible. For an input image preparation, we first apply an inverse coordinate transformation as:

$$\{I_1[f_1^{-1}(Y)], I_2[f_2^{-1}(Y)], \dots, I_B[f_B^{-1}(Y)]\} \quad (15)$$

where $I_i[f_i^{-1}(Y)] \in \mathbb{R}^{h_i \times w_i \times 3}$. While avoiding void pixels, we crop input images $I_i^{crop}[f_i^{-1}(Y)] \in \mathbb{R}^{h \times w \times 3}$, where $h \leq \min(\{h_i\})$, $w \leq \min(\{w_i\})$. For a ground truth (GT) preparation, we randomly sample \mathbf{M} query points among valid coordinates for an i -th batch element as: $\{\mathbf{y}_1^i, \mathbf{y}_2^i, \dots, \mathbf{y}_M^i\}$, where $\mathbf{y}_j^i \in \mathbb{R}^2$. Then, we evaluate our LTEW for each query point with cropped input images to compute loss as in Fig. 4 by comparing with following GT batch:

$$\begin{pmatrix} I_1[\mathbf{y}_1^1] & \dots & I_B[\mathbf{y}_1^B] \\ \vdots & \ddots & \vdots \\ I_1[\mathbf{y}_M^1] & \dots & I_B[\mathbf{y}_M^B] \end{pmatrix} \quad (16)$$

where $I_i[\mathbf{y}_j^i]$ is an RGB value.

For asymmetric-scale SR, each scale factor s_x, s_y is randomly sampled from $\mathcal{U}(0.25, 1)$. For homography transform, we randomly sample inverse coordinate transformation from the following distribution, dubbed *in-scale*:

$$\begin{pmatrix} 1 & h_x & 0 \\ h_y & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \cos \theta & \sin \theta & 0 \\ -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ p_x & p_y & 1 \end{pmatrix} \quad (17)$$

where $h_x, h_y \sim \mathcal{U}(-0.25, 0.25)$ are for sheering, $\theta \sim \mathcal{N}(0, 0.15^\circ^2)$ is for rotation, $s_x, s_y \sim \mathcal{U}(0.35, 0.5)$ are for scaling, $t_x \sim \mathcal{U}(-0.75W, 0.125W)$, $t_y \sim \mathcal{U}(-0.75H, 0.125H)$, $p_x \sim \mathcal{U}(-0.6W, 0.6W)$, $p_y \sim \mathcal{U}(-0.6H, 0.6H)$ are for projection. We evaluate our LTEW for unseen transformations to verify the generalization ability. For asymmetric-scale SR and homography transform, untrained coordinate transformations are sampled from $s_x, s_y \sim \mathcal{U}(0.125, 0.25)$, dubbed *out-of-scale*, other parameter distributions for $h_x, h_y, \theta, t_x, t_y, p_x, p_y$ remain the same.

5 Experiments

5.1 Dataset and Training

We use a DIV2K dataset [1] of an NTIRE 2017 challenge [45] for training. For optimization, we use an L1 loss [30] and an Adam [26] with $\beta_1 = 0.9$, $\beta_2 = 0.999$. Networks are trained for 1000 epochs with batch size 16. The learning rate is initialized as $1e-4$ and reduced by 0.5 at [200, 400, 600, 800]. Due to the page limit, evaluation details are provided in the supplementary.

5.2 Evaluation

Asymmetric-scale SR We compare our LTEW for asymmetric-scale SR to RCAN [51], MetaSR [19], ArbSR [46] within *in-scale* in Table 1, Fig. 5 and *out-of-scale* in Table 2, Fig. 6. For RCAN [51], we first upsample LR images by a factor of 4 and resample using bicubic interpolation. For MetaSR [19], we first upsample input images by a factor of $\max(s_x, s_y)$ and downsample using a bicubic method as in [46]. Except for the case of Set14 ($\times 3.5, \times 2$) and B100 ($\times 1.5, \times 3$), LTEW outperforms existing methods within asymmetric scale factors in performance and visual quality for all scale factors and all datasets.

Homography transform We compare our LTEW for homography transform to RRDB [48] and SRWarp [42] in Table 3, Fig. 7, Fig. 8. For RRDB [48], we super-sample input images by a factor of 4 and transform using bicubic resampling as in [42]. We see that our LTEW surpasses existing homography transform methods in mPSNR and visual quality for both *in-scale* and *out-of-scale*

Table 1. Quantitative comparison with state-of-the-art methods for **asymmetric-scale SR** within **in-scale** on benchmark datasets (PSNR (dB)). **Red** and **blue** colors indicate the best and the second-best performance, respectively.

Method	Set5			Set14			B100			Urban100		
	$\times 1.5$ $\times 4$	$\times 1.5$ $\times 3.5$	$\times 1.6$ $\times 3.05$	$\times 4$ $\times 2$	$\times 3.5$ $\times 2$	$\times 3.5$ $\times 1.75$	$\times 4$ $\times 1.4$	$\times 1.5$ $\times 3$	$\times 3.5$ $\times 1.45$	$\times 1.6$ $\times 3$	$\times 1.6$ $\times 3.8$	$\times 3.55$ $\times 1.55$
Bicubic	30.01	30.83	31.40	27.25	27.88	27.27	27.45	28.86	27.94	25.93	24.92	25.19
RCAN [51]	34.14	35.05	35.67	30.35	31.02	31.21	29.35	31.30	29.98	30.72	28.81	29.34
MetaSR-RCAN [19]	34.20	35.17	35.81	30.40	31.05	31.33	29.43	31.26	30.09	30.73	29.03	29.67
Arb-RCAN [46]	34.37	35.40	36.05	30.55	31.27	31.54	29.54	31.40	30.22	31.13	29.36	30.04
LTEW-RCAN (ours)	34.45	35.46	36.12	30.57	31.21	31.55	29.62	31.40	30.24	31.25	29.57	30.21




Fig. 5. Qualitative comparison to other **asymmetric-scale SR** within **in-scale**. RCAN [51] is used as an encoder for all methods.

Table 2. Quantitative comparison with state-of-the-art methods for **asymmetric-scale SR** within **out-of-scale** on benchmark datasets (PSNR (dB)). **Red** and **blue** colors indicate the best and the second-best performance, respectively.

Method	Set5			Set14			B100			Urban100		
	$\times 3$ $\times 8$	$\times 3$ $\times 7$	$\times 3.2$ $\times 6.1$	$\times 8$ $\times 4$	$\times 7$ $\times 4$	$\times 7$ $\times 3.5$	$\times 8$ $\times 2.8$	$\times 3$ $\times 6$	$\times 7$ $\times 2.9$	$\times 3.2$ $\times 6$	$\times 3.2$ $\times 7.6$	$\times 7.1$ $\times 3.1$
Bicubic	25.69	26.35	26.84	24.27	24.62	24.79	24.67	25.58	24.98	22.55	21.92	22.15
RCAN [51]	29.00	30.01	30.46	26.48	26.94	27.11	26.06	27.19	26.47	25.52	24.50	24.84
MetaSR-RCAN [19]	28.75	29.74	30.38	26.32	26.85	27.03	26.07	27.15	26.45	25.50	24.47	24.84
Arb-RCAN [46]	28.37	29.35	30.08	26.06	26.63	26.84	25.91	27.14	26.40	25.36	24.12	24.61
LTEW-RCAN (ours)	29.26	30.16	30.64	26.60	27.06	27.25	26.25	27.28	26.62	25.85	24.79	25.18




Fig. 6. Qualitative comparison to other **asymmetric-scale SR** within **out-of-scale**. RCAN [51] is used as an encoder for all methods.

Table 3. Quantitative comparison with state-of-the-art methods for homography transform within in-scale (*isc*) and out-of-scale (*osc*) on DIV2KW and benchmarkW datasets (mPSNR (dB)). **Red** and **blue** colors indicate the best and the second-best performance, respectively.

Method	DIV2KW		Set5W		Set14W		B100W		Urban100W	
	<i>isc</i>	<i>osc</i>	<i>isc</i>	<i>osc</i>	<i>isc</i>	<i>osc</i>	<i>isc</i>	<i>osc</i>	<i>isc</i>	<i>osc</i>
Bicubic	27.85	25.03	35.00	28.75	28.79	24.57	28.67	25.02	24.84	21.89
RRDB [48]	30.76	26.84	37.40	30.34	31.56	25.95	30.29	26.32	28.83	23.94
SRWarp-RRDB [42]	31.04	26.75	37.93	29.90	32.11	25.35	30.48	26.10	29.45	24.04
LTEW-RRDB (ours)	31.10	26.92	38.20	31.07	32.15	26.02	30.56	26.41	29.50	24.25




Fig. 7. Qualitative comparison to other homography transform within in-scale. RRDB [48] is used as an encoder for all methods.




Fig. 8. Qualitative comparison to other homography transform within out-of-scale. RRDB [48] is used as an encoder for all methods.

ERP perspective projection We train our LTEW to perform homography transform and apply an unseen transformation: ERP perspective projection, to validate the generalization ability. Since we are not able to obtain high-quality GT for ERP perspective projection (due to JPEG compression artifact), we visually compare our method to RRDB [48] and SRWarp [42] in Fig. 9. For RRDB [48], we upsample input ERP images by a factor of 4 and interpolate them in a bicubic manner. Resolution of input ERP images is 1664×832 . As pointed out in [14], considering HMD’s limited hardware resources, storing and transmitting ERP images in full resolution is impractical. Therefore, we downsample HR ERP images by a factor of 4 and then project images to a size of 832×832 .




Fig. 9. Qualitative comparison to other **warping** for **ERP perspective projection**. RRDB [48] is used as an encoder for all methods.




Fig. 10. Visual demonstration of LTEW for **HR ERP perspective projection**. RRDB [48] and SRWarp [42] are not able to evaluate given ERP size (8190px×2529px) under CPU-computing with 512GB RAM.

with a field of view (FOV) 120° . From Fig. 9, we observe that RRDB [48] is limited in capturing high-frequency details, and SRWarp [42] shows artifacts near boundaries. In contrast, our proposed LTEW captures fine details without any artifacts near boundaries.

In Fig. 10, we project an HR ERP image with size 8190×2529 to 4095×4095 with a FOV 90° on the Intel Xeon Gold 6226R@2.90GHz and RAM 512GB. Note that RRDB [48] and SRWarp [42] are not able to evaluate given HR ERP size even under CPU-computing. Both RRDB and SRWarp need to upsample an ERP image by a factor of 4, consuming a massive amount of memory. In contrast, our LTEW is memory-efficient since we are able to query evaluation points sequentially. We notice that our LTEW is capable of restoring sharper and clearer edges compared to bicubic interpolation.

5.3 Ablation study

In Table. 4, we explore other arbitrary-scale SR methods [10, 50] for image warping and validate our design with extensive ablation studies. Rows 1-5 show that Fourier features provide performance gain compared to concatenation [10], transformer [50] or periodic activation [40]. Since [10, 50] are designed to perform rectangular SR, we retrain them after modifying *cell* in [10] (rows 2-3) and *token* in [50] (row 5) to our shape term as Eq. (12).

Table 4. Quantitative ablation study for **homography transform** within **in-scale** (*isc*) and **out-of-scale** (*osc*) on DIV2KW dataset (mPSNR (dB)). EDSR-baseline [30] is used as an encoder.

Method	Amp.	Long.	N_{F_j}	Act.	Jacob.	Hess.	<i>isc</i>	<i>osc</i>
LIIF ($-c$) [10]				ReLU	\times	\times	30.65(−0.28)	26.73(+0.00)
LIIF [10] + Eq. (12)				ReLU	\checkmark	\checkmark	30.74(−0.09)	26.66(−0.07)
LIIF [10] + Eq. (12) + [40]				ReLU	\checkmark	\checkmark	30.49(−0.34)	26.52(−0.21)
ITSRN ($-token$) [50]				ReLU	\times	\times	30.64(−0.19)	26.69(−0.04)
ITSRN [50] + Eq. (12)				ReLU	\checkmark	\checkmark	30.74(−0.09)	26.52(−0.21)
LTEW ($-A$)	\times	\checkmark	256	ReLU	\checkmark	\checkmark	30.77(−0.06)	26.66(−0.07)
LTEW ($-L$)	\checkmark	\times	256	ReLU	\checkmark	\checkmark	30.79(−0.04)	26.67(−0.06)
LTEW ($-F$)	\checkmark	\checkmark	128	ReLU	\checkmark	\checkmark	30.80(−0.03)	26.70(−0.03)
LTEW ($-P_J, -P_H, +c$ [27])	\checkmark	\checkmark	256	ReLU	$(2r_x, 2/r_y)$		25.23(−5.60)	25.91(−0.82)
LTEW ($-P_J, -P_H$)	\checkmark	\checkmark	256	ReLU	\times	\times	30.69(−0.14)	26.76 (+0.03)
LTEW ($-P_J$)	\checkmark	\checkmark	256	ReLU	\times	\checkmark	30.70(−0.13)	26.72(−0.01)
LTEW ($-P_H$)	\checkmark	\checkmark	256	ReLU	\checkmark	\times	30.80(−0.03)	26.73(+0.00)
LTEW + [40]	\checkmark	\checkmark	256	sin	\checkmark	\checkmark	30.80(−0.03)	26.71(−0.02)
LTEW	\checkmark	\checkmark	256	ReLU	\checkmark	\checkmark	30.83 (+0.00)	26.73(+0.00)

In rows 6-8, we remove an amplitude estimator (row 6), long skip connection (row 7), and reduce the number of estimated frequencies (row 8). We see that each component consistently enhances mPSNR of LTEW for both *in-scale* and *out-of-scale*. In row 9, we test LTEW with spatially-invariant cell as [27], specifically $r_x = 2w/W, r_y = 2h/H$. It causes a significant mPSNR drop for both *in-scale* and *out-of-scale*. From rows 10-12, we observe that both the Jacobian and Hessian shapes are significant in improving mPSNR only for *in-scale*. Inspired by [49], we hypothesize that INR performs superior interpolating unseen coordinates but relatively poorly extrapolating untrained shapes. Extrapolation for untrained shapes will be investigated in future work.

5.4 Fourier feature space




Fig. 11. Visualization of Fourier feature space from LTEW.

In Fig. 11, we visualize estimated Fourier feature space from LTEW and the discrete Fourier transform (DFT) of a warped image for validation. We first scatter estimated frequencies to 2D space and set the color according to magnitude. We observe that our LTEW (h_ψ) extracts Fourier information (\mathbf{F}_j) for an input image at \mathbf{x}_j ($\mathbf{I}^{\text{IN}}[\mathbf{x}_j]$). By observing pixels inside an encoder’s receptive field (RF), LTEW estimates dominant frequencies and corresponding Fourier coefficients for RF-sized local patches of an input image. Before our MLP (g_θ) represents $\mathbf{I}^{\text{WARP}}[f(\mathbf{x}_j)]$, the Fourier space (\mathbf{F}_j) is transformed by the Jacobian matrix $\mathbf{J}_f^{-T}(\mathbf{x}_j)$ and matched to a frequency response of an output image at $f(\mathbf{x}_j)$ ($\mathcal{F}(\mathbf{I}^{\text{WARP}}[f(\mathbf{x}_j)])$). We observe that LTEW utilizing the local grid in

output space ($\delta_{\mathbf{y}}$) instead of $\delta_{\mathbf{x}}$ diverges during training. We hypothesize that representing $\mathbf{I}^{\text{WARP}}[f(\mathbf{x}_j)]$ with frequencies (\mathbf{F}_j) of an input image makes the overall training procedure unstable. This indicates that a spatially-varying Jacobian matrix $\mathbf{J}_f^{-T}(\cdot)$ for the given coordinate transformation f is significant in predicting accurate frequency responses of warped images. By estimating Fourier response in latent space instead of directly applying DFT to input images, we avoid extracting undesirable frequencies due to aliasing, as discussed in [27].

5.5 Discussion

In Table 5, we compare model complexity and symmetric-scale SR performance of our LTEW for both *in-scale* and *out-of-scale* to other warping methods: ArbSR [46] and SRWarp [42]. Note that ArbSR [46] and SRWarp [42] are learned to perform asymmetric-scale SR and homography transform. Following [27], we use $\max(\mathbf{s}, \mathbf{s}_{tr})$ instead of \mathbf{s} in Eq. (13) for phase estimation. We see that LTEW significantly outperforms existing warping methods for *out-of-scale*, achieving competitive quality to [42, 46] for *in-scale*. A local ensemble [10, 22, 27] in LTEW, preventing blocky artifacts, makes the model more complex than ArbSR [46]. SRWarp [42] blends $\times 1$, $\times 2$, and $\times 4$ features, leading to increased model complexity than LTEW, which uses only an $\times 1$ feature map.

Table 5. Quantitative comparison with state-of-the-art warping methods for **symmetric-scale SR** on B100 dataset (PSNR (dB)). RCAN [51] and RRDB [48] are used as encoders. Computation time and memory consumption are measured with an 256×256 -sized input for an $\times 2$ task on NVIDIA RTX 3090 24GB.

Method	Training task	#Params.	Runtime	Memory	<i>in-scale</i>			<i>out-of-scale</i>	
					$\times 2$	$\times 3$	$\times 4$	$\times 6$	$\times 8$
Arb-RCAN [46]	Asymmetric -scale SR	16.6M	160ms	1.39GB	32.39	29.32	27.76	25.74	24.55
LTEW-RCAN (ours)		15.8M	283ms	1.77GB	32.36	29.30	27.78	26.01	24.95
SRWarp-RRDB [42]	Homography transform	18.3M	328ms	2.34GB	32.31	29.27	27.77	25.33	24.45
LTEW-RRDB (ours)		17.1M	285ms	1.79GB	32.35	29.29	27.76	25.98	24.95

6 Conclusions

In this paper, we proposed the continuous neural representation by learning Fourier characteristics of images warped by the given coordinate transformation. Particularly, we found that shape-dependent phase estimation and long skip connection enable MLP to predict signals more accurately. We demonstrated that the LTEW-based neural function outperforms existing warping techniques for asymmetric-scale SR and homography transform. Moreover, our method effectively generalizes untrained coordinate transformations, specifically *out-of-scale* and ERP perspective projection.

Acknowledgement This work was partly supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2021R1A4A1028652), the DGIST R&D Program of the Ministry of Science and ICT (No. 22-IJRP-01) and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. IITP-2021-0-02068).

References

1. Agustsson, E., Timofte, R.: NTIRE 2017 Challenge on Single Image Super-Resolution: Dataset and Study. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (July 2017)
2. Arican, Z., Frossard, P.: Joint Registration and Super-Resolution With Omnidirectional Images. *IEEE Transactions on Image Processing* **20**(11), 3151–3162 (2011)
3. Bagnato, L., Boursier, Y., Frossard, P., Vanderghenst, P.: Plenoptic based super-resolution for omnidirectional image sequences. In: 2010 IEEE International Conference on Image Processing. pp. 2829–2832 (2010)
4. Benbarka, N., Höfer, T., Riaz, H.u.M., Zell, A.: Seeing Implicit Neural Representations As Fourier Series. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 2041–2050 (January 2022)
5. Boomsma, W., Frellsen, J.: Spherical convolutions and their application in molecular modelling. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 30. Curran Associates, Inc. (2017)
6. Bracewell, R., Chang, K.Y., Jha, A., Wang, Y.H.: Affine theorem for two-dimensional Fourier transform. *Electronics Letters* **29**(3), 304–304 (1993)
7. Chan, K.C., Wang, X., Yu, K., Dong, C., Loy, C.C.: BasicVSR: The Search for Essential Components in Video Super-Resolution and Beyond. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4947–4956 (June 2021)
8. Chen, H., Wang, Y., Guo, T., Xu, C., Deng, Y., Liu, Z., Ma, S., Xu, C., Xu, C., Gao, W.: Pre-Trained Image Processing Transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12299–12310 (June 2021)
9. Chen, H., He, B., Wang, H., Ren, Y., Lim, S.N., Shrivastava, A.: NeRV: Neural Representations for Videos. In: Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems*. vol. 34, pp. 21557–21568. Curran Associates, Inc. (2021)
10. Chen, Y., Liu, S., Wang, X.: Learning Continuous Image Representation With Local Implicit Image Function. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8628–8638 (June 2021)
11. Chiang, M.C., Boulton, T.: Efficient super-resolution via image warping. *Image and Vision Computing* **18**(10), 761–771 (2000)
12. Chiang, M.C., Boulton, T.: Efficient image warping and super-resolution. In: Proceedings Third IEEE Workshop on Applications of Computer Vision. WACV’96. pp. 56–61 (1996)
13. Dai, T., Cai, J., Zhang, Y., Xia, S.T., Zhang, L.: Second-Order Attention Network for Single Image Super-Resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
14. Deng, X., Wang, H., Xu, M., Guo, Y., Song, Y., Yang, L.: LAU-Net: Latitude adaptive upscaling network for omnidirectional image super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9189–9198 (June 2021)
15. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3–7, 2021. OpenReview.net (2021)

16. Fathony, R., Sahu, A.K., Willmott, D., Kolter, J.Z.: Multiplicative Filter Networks. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net (2021)
17. Greene, N., Heckbert, P.S.: Creating Raster Omnimax Images from Multiple Perspective Views Using the Elliptical Weighted Average Filter. *IEEE Computer Graphics and Applications* **6**(6), 21–27 (1986)
18. Hornik, K., Stinchcombe, M., White, H.: Multilayer Feedforward Networks Are Universal Approximators. *Neural Netw.* **2**(5), 359–366 (Jul 1989)
19. Hu, X., Mu, H., Zhang, X., Wang, Z., Tan, T., Sun, J.: Meta-SR: A Magnification-Arbitrary Network for Super-Resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
20. Huang, M., Shen, Q., Ma, Z., Bovik, A.C., Gupta, P., Zhou, R., Cao, X.: Modeling the Perceptual Quality of Immersive Images Rendered on Head Mounted Displays: Resolution and Compression. *IEEE Transactions on Image Processing* **27**(12), 6039–6050 (2018)
21. Jaderberg, M., Simonyan, K., Zisserman, A., kavukcuoglu, k.: Spatial Transformer Networks. In: Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 28. Curran Associates, Inc. (2015)
22. Jiang, C.M., Sud, A., Makadia, A., Huang, J., Niessner, M., Funkhouser, T.: Local Implicit Grid Representations for 3D Scenes. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
23. Jiang, W., Trulls, E., Hosang, J., Tagliasacchi, A., Yi, K.M.: COTR: Correspondence Transformer for Matching Across Images. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 6207–6217 (October 2021)
24. Keys, R.: Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing* **29**(6), 1153–1160 (1981)
25. Kim, J., Lee, J.K., Lee, K.M.: Accurate Image Super-Resolution Using Very Deep Convolutional Networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016)
26. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. In: Bengio, Y., LeCun, Y. (eds.) *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings* (2015)
27. Lee, J., Jin, K.H.: Local Texture Estimator for Implicit Representation Function. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1929–1938 (June 2022)
28. Lee, Y., Jeong, J., Yun, J., Cho, W., Yoon, K.J.: SpherePHD: Applying CNNs on a Spherical PolyHeDron Representation of 360deg Images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
29. Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: SwinIR: Image Restoration Using Swin Transformer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops. pp. 1833–1844 (October 2021)
30. Lim, B., Son, S., Kim, H., Nah, S., Mu Lee, K.: Enhanced Deep Residual Networks for Single Image Super-Resolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (July 2017)

31. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin Transformer: Hierarchical Vision Transformer Using Shifted Windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10012–10022 (October 2021)
32. Mei, Y., Fan, Y., Zhou, Y.: Image Super-Resolution With Non-Local Sparse Attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3517–3526 (June 2021)
33. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy Networks: Learning 3D Reconstruction in Function Space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
34. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In: Proceedings of the European Conference on Computer Vision (ECCV) (August 2020)
35. Niu, B., Wen, W., Ren, W., Zhang, X., Yang, L., Wang, S., Zhang, K., Cao, X., Shen, H.: Single Image Super-Resolution via a Holistic Attention Network. In: Proceedings of the European Conference on Computer Vision (ECCV) (August 2020)
36. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
37. Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y., Courville, A.: On the Spectral Bias of Neural Networks. In: Chaudhuri, K., Salakhutdinov, R. (eds.) Proceedings of the 36th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 97, pp. 5301–5310. PMLR (09–15 Jun 2019)
38. Sajjadi, M.S.M., Vemulapalli, R., Brown, M.: Frame-Recurrent Video Super-Resolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
39. Shi, W., Caballero, J., Huszar, F., Totz, J., Aitken, A.P., Bishop, R., Rueckert, D., Wang, Z.: Real-Time Single Image and Video Super-Resolution Using an Efficient Sub-Pixel Convolutional Neural Network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016)
40. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit Neural Representations with Periodic Activation Functions. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) Advances in Neural Information Processing Systems. vol. 33, pp. 7462–7473. Curran Associates, Inc. (2020)
41. Sitzmann, V., Zollhoefer, M., Wetzstein, G.: Scene Representation Networks: Continuous 3D-Structure-Aware Neural Scene Representations. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 32. Curran Associates, Inc. (2019)
42. Son, S., Lee, K.M.: SRWarp: Generalized Image Super-Resolution under Arbitrary Transformation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7782–7791 (June 2021)
43. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: PWC-Net: CNNs for Optical Flow Using Pyramid, Warping, and Cost Volume. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)

44. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*. vol. 33, pp. 7537–7547. Curran Associates, Inc. (2020)
45. Timofte, R., Agustsson, E., Van Gool, L., Yang, M.H., Zhang, L.: NTIRE 2017 Challenge on Single Image Super-Resolution: Methods and Results. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops* (July 2017)
46. Wang, L., Wang, Y., Lin, Z., Yang, J., An, W., Guo, Y.: Learning a Single Network for Scale-Arbitrary Super-Resolution. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 4801–4810 (October 2021)
47. Wang, P.S., Liu, Y., Yang, Y.Q., Tong, X.: Spline Positional Encoding for Learning 3D Implicit Signed Distance Fields. In: Zhou, Z.H. (ed.) *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*. pp. 1091–1097. International Joint Conferences on Artificial Intelligence Organization (8 2021), main Track
48. Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Change Loy, C.: ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks. In: *Proceedings of the European Conference on Computer Vision (ECCV) Workshops* (September 2018)
49. Xu, K., Zhang, M., Li, J., Du, S.S., Kawarabayashi, K., Jegelka, S.: How Neural Networks Extrapolate: From Feedforward to Graph Neural Networks. In: *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net (2021)
50. Yang, J., Shen, S., Yue, H., Li, K.: Implicit Transformer Network for Screen Content Image Continuous Super-Resolution. In: Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems*. vol. 34, pp. 13304–13315. Curran Associates, Inc. (2021)
51. Zhang, Y., Li, K., Li, K., Wang, L., Zhong, B., Fu, Y.: Image Super-Resolution Using Very Deep Residual Channel Attention Networks. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (September 2018)
52. Zhang, Y., Tian, Y., Kong, Y., Zhong, B., Fu, Y.: Residual Dense Network for Image Super-Resolution. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2018)