

# CycleINR: Cycle Implicit Neural Representation for Arbitrary-Scale Volumetric Super-Resolution of Medical Data

Wei Fang<sup>1,2</sup> Yuxing Tang<sup>1</sup> Heng Guo<sup>1,2</sup> Mingze Yuan<sup>1,4</sup> Tony C. W. Mok<sup>1,2</sup>  
 Ke Yan<sup>1,2</sup> Jiawen Yao<sup>1,2</sup> Xin Chen<sup>3</sup> Zaiyi Liu<sup>3</sup> Le Lu<sup>1</sup> Ling Zhang<sup>1</sup> Minfeng Xu<sup>1,2</sup>

<sup>1</sup> DAMO Academy, Alibaba Group

<sup>2</sup> Hupan Lab, 310023, Hangzhou, China

<sup>3</sup> Guangdong Provincial People’s Hospital

<sup>4</sup> Peking University

lucas.fw@alibaba-inc.com

## Abstract

In the realm of medical 3D data, such as CT and MRI images, prevalent anisotropic resolution is characterized by high intra-slice but diminished inter-slice resolution. The lowered resolution between adjacent slices poses challenges, hindering optimal viewing experiences and impeding the development of robust downstream analysis algorithms. Various volumetric super-resolution algorithms aim to surmount these challenges, enhancing inter-slice resolution and overall 3D medical imaging quality. However, existing approaches confront inherent challenges: 1) often tailored to specific upsampling factors, lacking flexibility for diverse clinical scenarios; 2) newly generated slices frequently suffer from over-smoothing, degrading fine details, and leading to inter-slice inconsistency. In response, this study presents CycleINR, a novel enhanced Implicit Neural Representation model for 3D medical data volumetric super-resolution. Leveraging the continuity of the learned implicit function, the CycleINR model can achieve results with arbitrary up-sampling rates, eliminating the need for separate training. Additionally, we enhance the grid sampling in CycleINR with a local attention mechanism and mitigate over-smoothing by integrating cycle-consistent loss. We introduce a new metric, Slice-wise Noise Level Inconsistency (SNLI), to quantitatively assess inter-slice noise level inconsistency. The effectiveness of our approach is demonstrated through image quality evaluations on an in-house dataset and a downstream task analysis on the Medical Segmentation Decathlon liver tumor dataset.

## 1. Introduction

Volumetric medical imaging, a cornerstone of diagnostic radiology involving techniques such as computed tomography

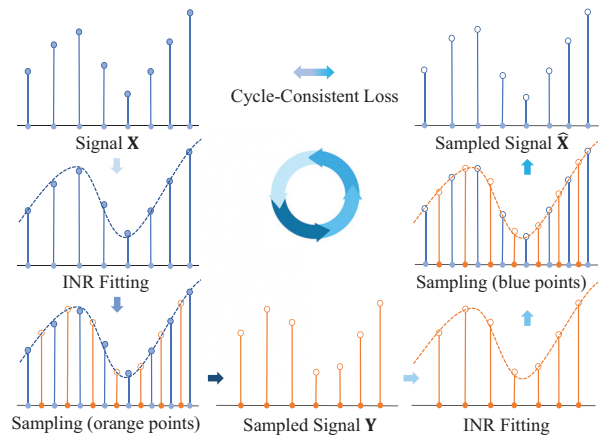


Figure 1. The core concept of CycleINR involves the initial use of signal  $X$  to fit a continuous Implicit Neural Representation (INR) function (shown in dotted blue curves). Subsequently, new points (depicted as hollow orange dots) are sampled from this function to create a new INR function (illustrated by dotted orange curves). The signal  $\hat{X}$  is then sampled from the new function at the same positions as  $X$ . The construction of a cycle-consistent loss is achieved by assessing the similarity between  $\hat{X}$  and  $X$ .

(CT) and magnetic resonance imaging (MRI), preserves the 3D three-dimensional characteristics of the body’s internal structures through the compilation of multiple cross-sectional images [47]. While high-resolution volumetric medical imaging offers detailed anatomical and functional information enhancing diagnosis, its widespread clinical adoption is hindered by prolonged acquisition time and elevated storage costs. As a common alternative, anisotropic volumes, characterized by high resolution within each slice and lower resolution between slices, have become prevalent [41]. However, this reduced inter-slice resolution introduces challenges: limiting detailed sagittal or coronal views, complicating lesion observation, and posing difficul-

ties for robust 3D medical image interpretation algorithms. Addressing this, there is a critical need for an accurate 3D super-resolution method, essentially framing it as a volumetric super-resolution task [33, 52]. It is noteworthy that video frame interpolation [35], which involves synthesizing intermediate images between a pair of input frames, exhibits notable similarities with volumetric interpolation.

Implementing volumetric super-resolution algorithms poses challenges on various fronts. Firstly, most super-resolution methods are constrained to specific super-resolution ratios, necessitating the training of multiple models for different super-resolution levels. Secondly, the newly generated slices often exhibit over-smoothing, creating a noticeable slice-wise inconsistency issue in volumetric scenarios. This becomes evident when scrolling through the slices and observing the contrast between the original and generated slices. In contrast, in 2D scenarios, newly generated pixels are usually interpolated between existing pixels, making them less perceptible.

To tackle these challenges, we propose CycleINR (cycle-consistent loss enhanced implicit neural representation) as a solution for volumetric super-resolution. CycleINR framework employs a unified implicit neural voxel function to represent both low-resolution (LR) and high-resolution (HR) images, each at different sampling rates. This enables treating volumetric super-resolution as an implicit neural function approximation problem. Leveraging the continuity of the implicit neural function, a single implicit voxel function achieves volumetric super-resolution with arbitrary up-sampling scales. To mitigate over-smoothing, we integrate cycle-consistent loss (CCL) into our method. This entails the utilization of pixel values at LR grid points to fit the implicit neural function, generating pixel values for initially empty points (forward path). Furthermore, we employ the newly generated HR images to reconstruct the LR images (backward path). The backward path ensures the generated image maintains a consistent noise level with the original images. These paths establish a cyclical process, ensuring both accurate and consistent HR image recovery. Additionally, in our sampling process from the implicit neural representation function, we design a local attention mechanism (LAM) to consider both spatial proximity and numerical similarity. This ensures accurate and reliable sampling by incorporating both local spatial relationships and pixel value similarities. Finally, we introduce a novel metric, slice-wise noise level inconsistency (SNLI), to quantitatively measure inter-slice inconsistency, tackling limitations in existing metrics such as PSNR (peak signal-to-noise ratio), SSIM (structural similarity index), and LPIPS (learned perceptual image patch similarity) [54]. These metrics are more favorable towards assessing the smoothness or single-image similarity, making SNLI crucial for evaluating inter-slice variations in volumetric data and showcasing the ef-

fectiveness of CCL.

Our contributions are outlined as follows:

1. We introduce CycleINR, a novel framework for volumetric super-resolution in 3D medical imaging. Leveraging its coordinate-based nature, our method offers flexibility for arbitrary upsampling scales and the ability to generate multiple super-resolution results without the need for additional training.
2. We propose a cycle-consistent loss within the implicit neural representation framework to address the slice-wise inconsistency problem, ensuring consistency between newly generated and existing images. Additionally, a local attention mechanism captures both spatial proximity and numerical similarity, enhancing the depiction of relationships between pixels.
3. We propose a new metric SNLI (slice-wise noise level inconsistency) to quantitatively measure inter-slice noise level inconsistency in 3D data.
4. We conduct extensive experiments, evaluating both image quality and downstream tasks, demonstrating the effectiveness of our proposed method.

## 2. Related Work

**Single Image Super-Resolution.** Single image super-resolution (SISR) has garnered significant attention in computer vision, evolving over decades [51]. Common interpolation-based methods like cubic and linear interpolation are widely employed for their simplicity [3]. Image prior like total variation can also enhance image super-resolution by formulating the task as a convex optimization problem [36]. In the era of deep learning, Dong et al. [8] introduced super-resolution convolutional neural networks (CNN) to learn the mapping from LR images to HR images. Subsequent SISR research has delved into refining deep learning models, exploring strategies like deeper network architectures [18] [13], recursive-supervision [19], and deep learning as priors [53]. Advanced upsampling methods in deep learning include deconvolution for direct HR image generation [9], sub-pixel convolution, and pixel shuffling for efficient upsampling at the output stage [37]. Residual learning [24] and the integration of generative adversarial networks [22, 44, 45] have further enriched the SISR landscape.

**Implicit Neural Representation.** Digitalized visual signals transition from their natural continuous form to discrete representations, such as 2D images or volumetric medical data. Implicit neural representation (INR) emerges as a coordinate-based strategy to address this transformation, employing a trainable continuous function often implemented with a straightforward fully connected network [40]. INR was first applied in 3D shape modeling [6, 27, 31]. Subsequently, it has been extensively employed across various 3D tasks, including object shape

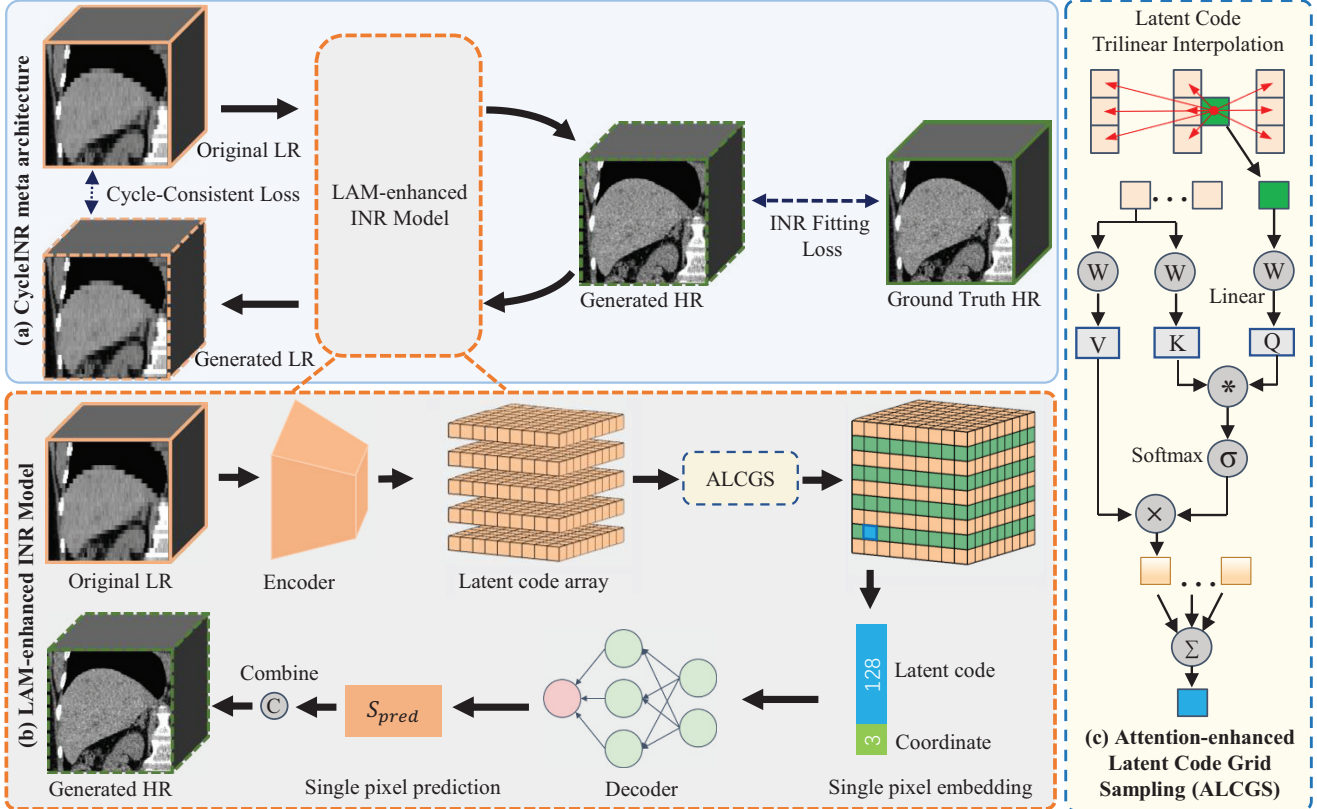


Figure 2. Schematic plot of our CycleINR framework for volumetric super-resolution. (a) Meta architecture of the proposed CycleINR. (b) Local Attention Mechanism (LAM) enhanced INR model. (c) Attention-enhanced Latent Code Grid Sampling (ALCGS) process.

modeling [2, 14], scene reconstruction [17, 39], and structure rendering [25, 29]. Notably, INR has made significant contributions to view synthesis, exemplified by the neural radiance fields (NeRF) [28]. INR has also been applied to 2D image super-resolution [4, 23, 26, 49, 50]. Chen et al. [5] used INR to learn a continuous function for 2D images. The function, taking an image coordinate and surrounding 2D deep features as inputs, predicts the RGB value at the given coordinate, enabling image generation at arbitrary resolutions. Sitzmann et al. [40] enhanced image representation by replacing the rectified linear units (ReLU) with periodic activations in multi-layer perceptron, achieving higher image quality.

**Volumetric Super-Resolution.** Volumetric super-resolution in 3D volumetric medical images, which aims to enhance spatial resolution between slices, has witnessed significant progress. DA-VSR [33] addressed domain adaptation challenges, SAINT [32] employed a two-stage framework with separate 2D CNNs for sagittal and coronal up-scaling, ArSSR [48] offered an INR-based arbitrary scale SISR model for 3D MR images, and TVSRN [52] introduced a transformer-based approach for CT scans. Additionally, Fang et al. [12] proposed a medical slice synthesis method using mutual distillation.

### 3. CycleINR

Volumetric super-resolution can be viewed as a z-axis upsampling process applied to 3D volume data. Let  $\mathbf{I}^{\text{lr}}(x, y, z) \in \mathbb{R}^{X \times Y \times Z}$  denote the original low-resolution (LR) CT scan, where  $X$ ,  $Y$ , and  $Z$  represent the number of slices along the  $x$ ,  $y$ , and  $z$  axes, respectively. The desired output of volumetric super-resolution is denoted as  $\mathbf{I}^{\text{vsr}}(x, y, z) \in \mathbb{R}^{X \times Y \times [Z \cdot r]}$ , where  $r > 1$  represents the upsampling factor along the  $z$ -axis. This implies that super-resolution is applied solely in the  $z$ -direction, maintaining the original resolution in the  $x$  and  $y$  directions, allowing more detailed information between slices in the volumetric data. For instance, setting the upsampling factor  $r$  to 5 facilitates the conversion from 5mm slice-thickness data to 1mm, or 4 for the conversion from 5mm to 1.25mm slice-thickness. In the context of super-resolution methods employing the coordinate-based INR (implicit neural representation) model,  $r$  can also be a decimal, given the availability of coordinates. The volumetric super-resolution transform can be represented as:

$$\mathcal{F} : \mathbf{I}^{\text{lr}} \in \mathbb{R}^{X \times Y \times Z} \rightarrow \mathbf{I}^{\text{vsr}} \in \mathbb{R}^{X \times Y \times [Z \cdot r]}. \quad (1)$$

### 3.1. Local Attention Mechanism-Enhanced INR

In the original INR model, 3D volumetric medical data can be represented using a coordinate-based implicit voxel function  $I = f_\theta(\mathbf{x})$ , where  $\theta$  represents the trainable parameters of the INR model,  $\mathbf{x} = (x, y, z)$  denotes any 3D voxel spatial coordinate, and  $I$  denotes the voxel intensity at the coordinate  $\mathbf{x}$  in the image. The spatial coordinates are normalized to the range  $[-1, 1]$  along the  $x$ ,  $y$ , and  $z$  dimensions, differing from physical coordinates. Both HR and LR images can be regarded as the explicit representation of the implicit function  $f_\theta(x)$  but at different sampling rates. Once the voxel function  $f_\theta(x)$  is well approximated, the desired super-resolution results at an arbitrary scale can be achieved by simply adjusting the sampling rate (Figure 1).

However, in  $I = f_\theta(\mathbf{x})$ , the image information is decoded into the parameters of the function, restricting the function’s specificity to the represented image. As training a unique model for each image is impractical, a more flexible function that can represent multiple images is essential. Inspired by [5], we reformulate the voxel function by incorporating a latent code using  $I = f_\theta(\mathbf{x}, \mathbf{z})$ , where  $\mathbf{z}$  is a vector that represents the latent code capturing pixel-specific spatial features, which is generated for each pixel through a convolutional encoder, capturing local semantic information around the pixels. Hence, it can be considered as a latent representation for the pixel in the latent space. The architecture of the INR model with local attention mechanism enhanced latent code grid sampling process, is illustrated in Figure 2 (b). The following paragraphs will delve into a detailed explanation of these designs.

**Encoder Network.** We first employ an encoder  $\mathcal{E}_\theta$  to extract semantic features from the LR volume. The encoder takes the LR image volume  $\mathbf{I}^{lr} \in \mathbb{R}^{h \times w \times d}$  as input and produces a feature map  $\mathbf{Z}^{lr} \in \mathbb{R}^{h \times w \times d \times |\mathbf{z}|}$ , referred to as the latent code. This implies that each voxel in the LR volume is transformed into a latent code with a length of  $|\mathbf{z}|$ . The encoding process can be represented as:

$$\mathbf{Z}^{lr} = \mathcal{E}_\theta(\mathbf{I}^{lr}), \quad (2)$$

where  $\theta$  is the trainable parameters of the encoder. We utilize a residual convolutional neural network encoder as described in [11]. Our strategy to extract voxel-wise latent space is crucial for facilitating the decoder network in seamlessly incorporating local image intensity information. This integration proves instrumental in effectively recovering fine details within the high-resolution image, especially when dealing with large upsampling scales.

**Attention-enhanced Latent Code Grid Sampling.** Following the encoding process, the obtained latent code corresponds to the voxels in the LR volume. To acquire the latent code for the voxels at coordinate  $\mathbf{x}$  in the HR volume  $\mathbf{Z}^{hr}$ , an additional operation known as grid sampling is required. Prior works [5, 48] performed grid sampling using

techniques such as trilinear interpolation or cell decoding, which can be expressed as:

$$\hat{\mathbf{z}}_{\mathbf{x}}^{hr} = \sum_{k=1}^L \left( \frac{S_k}{S} \cdot \mathbf{z}_k^{lr} \right), \quad (3)$$

where  $\mathbf{z}_k^{lr}$  ( $k = 1, \dots, L$ ) denotes the latent code of the  $L = 8$  surrounding coordinates in the feature map  $\mathbf{Z}^{lr}$ .  $S_k$  represents the volume of the cuboid formed using the coordinate  $\mathbf{x}$  and the coordinate diagonal to the  $k_{th}$  coordinate, and  $S$  is the sum of all  $S_k$ .

To better depict the relationships between voxels, we propose an attention-enhanced latent code grid sampling (ALCGS) mechanism to enhance the generated voxel latent code at coordinate  $\mathbf{x}$ , which is represented as follows:

$$\tilde{\mathbf{z}}_{\mathbf{x}}^{hr} = \sum_{i=1}^N \sigma(\mathcal{W}_Q(\hat{\mathbf{z}}_{\mathbf{x}}^{hr}) \cdot \mathcal{W}_K(\mathbf{z}_i^{lr})) \mathcal{W}_V(\mathbf{z}_i^{lr}), \quad (4)$$

where  $N$  is the number of neighboring latent codes used for ensembling, and  $\sigma$  denotes the Softmax function. Specifically, the initial latent code  $\hat{\mathbf{z}}_{\mathbf{x}}^{hr}$  of the voxel at coordinate  $\mathbf{x}$  is firstly computed by using trilinear interpolation. Subsequently,  $N$  nearest neighboring latent codes  $\mathbf{z}_i^{lr}$  are selected around  $\hat{\mathbf{z}}_{\mathbf{x}}^{hr}$  for aggregation. The initial latent code  $\hat{\mathbf{z}}_{\mathbf{x}}^{hr}$  is then mapped to the query space, while the  $N$  nearby latent codes  $\mathbf{z}_i^{lr}$  are mapped to the key and value space, respectively, through linear layers  $\mathcal{W}_Q$ ,  $\mathcal{W}_K$ , and  $\mathcal{W}_V$ . Following this, the similarities between the query vector and each key vector are computed, which are then converted to weight coefficients through the Softmax function. Finally, the improved latent code  $\tilde{\mathbf{z}}_{\mathbf{x}}^{hr}$  is obtained by weighted summation.

The local attention mechanism establishes pixel relationships not solely dependent on physical distance but also on visual information and latent code similarity. This aggregation strategy enhances trilinear interpolation and enables the generated latent codes to preserve richer details, similar to the functionality of bilateral filtering.

**Decoder Network.** The decoder network  $\mathcal{D}_\theta$  takes the query voxel coordinate  $\mathbf{x}$  and its latent code  $\tilde{\mathbf{z}}_{\mathbf{x}}^{hr}$  as input and outputs the estimated voxel value  $I_{hr}(\mathbf{x})$ , expressed as:

$$I_{hr}(\mathbf{x}) = \mathcal{D}_\theta(\mathbf{x}, \tilde{\mathbf{z}}_{\mathbf{x}}^{hr}). \quad (5)$$

Following the design of [48], the decoder network comprises eight standard fully connected (FC) layers, with a ReLU activation following each FC layer. A residual connection is employed between the input of the decoder network and the intermediate fourth FC layer.

### 3.2. Cycle-Consistent Loss for INR Model

Inspired by CycleGAN [55], we integrate a cycle-consistent loss into the INR model. This loss acts as a constraint to ensure consistency between the generated image and the original image in terms of image features, such as noise level.

Despite both being generative models, generative adversarial networks (GANs) and INR diverge in their objectives. GANs transform one probability distribution function into another, whereas the INR model learns a continuous function from discrete data. Arbitrary scaled super-resolution images can be sampled from this continuous function. The newly sampled image can also be used to fit a new continuous function. We assert that the newly generated function should remain consistent with the original discrete data, leading to the incorporation of cycle-consistent loss in the INR model.

The synergy between cycle-consistent loss and the INR model is highly effective, highlighting the INR model’s superior suitability for incorporating cycle-consistent loss compared to GANs. GANs are designed to transform between distinct input and output distributions, while INR models aim to model the same function for both the input LR image and the output HR image. It is expected that the newly generated continuous function should produce values matching the original discrete data when sampled.

Let’s represent the local attention mechanism (LAM)-enhanced INR model described in Figure 2 (b) as  $\mathcal{G}_\theta$ . In this context, the cycle-consistent loss is articulated as:

$$\hat{\mathbf{I}}_j^{hr} = \mathcal{G}_\theta(\mathbf{I}_j^{lr}, \mathbf{C}_j^{hr}), \quad (6)$$

$$\mathcal{L}_{Cycle} = \sum_j^M \|\mathcal{G}_\theta(\hat{\mathbf{I}}_j^{hr}, \mathbf{C}_j^{lr}) - \mathbf{I}_j^{lr}\|_1, \quad (7)$$

where  $j$  and  $M$  are the index and the number of training samples, respectively. The INR model  $\mathcal{G}_\theta$  takes the LR image  $\mathbf{I}_j^{lr}$  and the coordinates of the voxels to be generated  $\mathbf{C}_j^{hr}$  as input, producing the sampled HR image  $\hat{\mathbf{I}}_j^{hr}$ . In Equation 7, the newly generated image  $\hat{\mathbf{I}}_j^{hr}$ , along with the coordinates of the original LR image  $\mathbf{C}_j^{lr}$ , is once again fed as input into the INR model to generate an estimation for the LR image. The estimated LR image is subsequently compared for similarities with the original LR image  $\mathbf{I}_j^{lr}$  using the L1 loss, constituting the cycle-consistent loss.

It is worth noting that HR pixel coordinates may not cover all LR pixel coordinates, given the random sampling of scaling factors, which may involve decimals during training. Even with integers, especially when a scaling factor is an even number and the interpolation mode align-corner is set to False, there may be no overlap between the pixel positions in the upscaled image and those of the original LR image. This complexity makes the downsampling process in Equation 7 non-trivial, and obtaining it can not be accomplished by simply selecting pixels at regular intervals.

The forward INR fitting loss can be represented as:

$$\mathcal{L}_{INR} = \|\hat{\mathbf{I}}_j^{hr} - \mathbf{I}_j^{hr}\|_1. \quad (8)$$

The overall loss of training is represented as follows:

$$\mathcal{L}_{CycleINR} = \mathcal{L}_{INR} + \lambda \cdot \mathcal{L}_{Cycle}, \quad (9)$$

where  $\lambda$  serves as a regularization factor that helps to balance the INR fitting loss and cycle-consistent loss.

### 3.3. SNLI: Measure Slice-wise Noise Inconsistency

To quantitatively measure the inter-slice inconsistency caused by the over-smoothing in the newly generated slices, we propose a novel metric named slice-wise noise level inconsistency (SNLI). This inconsistency often manifests as a perceptible ‘slice discontinuity’ phenomenon during axial slice scrolling. Notably, there is currently no existing metric that adequately captures this phenomenon. The proposed SNLI metric is defined as follows:

$$\text{SNLI} = \Psi(\zeta(\mathbf{I}_t)), t = 0, 1, \dots, T, \quad (10)$$

where  $\mathbf{I}_t \in \mathbb{R}^{X \times Y}$  is the  $t_{th}$  axial slice of the image volume.  $\zeta$  is a robust wavelet-based estimator of the Gaussian noise standard deviation, described in [10].  $\Psi$  represents the process of calculating the standard deviation. Thus, the SNLI metric helps measure the slice-wise noise level inconsistency between axial slices, which was ignored by many deep learning-based volumetric super-resolution methods.

## 4. Experiments

### 4.1. Experimental Setup

**Data.** We employed two datasets in our experiment. The first, a privately collected chest CT data (referred to as the chest dataset hereafter) consists of 204 CT volumes with 1mm slice thickness. We randomly split it into 124 for training, 40 for validation, and 40 for testing. Following established practices in volumetric super-resolution [32, 48], we derived the LR volumes by down-sampling from the corresponding HR volume, creating LR-HR training pairs. The CycleINR model and other comparing learning-based methods were trained on these pairs. Testing involved down-sampling to 2mm, 3mm, and 5mm thickness using interval sampling, aiming to recover 1mm data using super-resolution approaches. Image quality was evaluated by comparing results with ground truth 1mm data. The second dataset is from the public Medical Segmentation Decathlon [1] Task03\_Liver training set (referred to as the MSD liver dataset). It serves to assess whether the HR images generated by the CycleINR model achieve performance comparable to the original HR images in downstream segmentation tasks. The MSD liver dataset consists of 131 volumes with liver and tumor masks. For evaluation, we selected 33 volumes with 1mm slice thickness as the test set and utilized the remaining 98 image-mask pairs to train a segmentation model using nnUNet [16]. The test data was downsampled and then restored to 1mm as on the chest CT dataset. We compared the segmentation results of different methods, setting the segmentation from the original 1mm data as references.




Figure 3. Visualization results of different super-resolution (x5) methods from axial, sagittal, and coronal views. The second and third rows display newly generated images at new positions, while the first and fourth rows show generated images at their original positions. In the bottom row, red arrows in TVSRN, ArSSR, and SAINT highlight horizontal lines in the coronal view, indicating slice-wise inconsistency.

**Implementation Details.** The encoder consists of three blocks, where each block includes six 3D convolutional layers followed by ReLU. The latent code vector length  $|\mathbf{z}|$  is set to 128. The feature channels of the intermediate layers are set to 64. The number of neighboring latent codes  $N$  used for ensembling is set to 27 to include all the  $3 \times 3 \times 3$  local regions around a specific voxel. The decoder is configured with 8 fully connected layers, where the intermediate layer has a dimension of 256. During CycleINR training, LR-HR pairs are randomly sampled on the fly from the original entire volume with an upsample scale ranging from 2.0 to 6.0, with one decimal point precision. The balancing factor  $\lambda$  for the INR fitting loss and cycle-consistent loss is set

to 1. The Adam optimizer [20] minimizes the total loss in Equation 9, with an initial learning rate of  $10^{-4}$ , and decays as half every 200 epochs. The training runs for a total of 3,000 epochs. During testing, Gaussian window-weighted sliding window inference is employed to mitigate stitching artifacts. Note that these hyper-parameters are set following empirical observations or prior works.

**Compared Methods.** For comparison, we include cubic interpolation, trilinear interpolation, and three additional deep learning-based volumetric super-resolution methods: TVSRN [52], ArSSR [48], and SAINT [32]. These methods, featuring distinct architectural designs, demonstrate the capability of achieving arbitrary upscaling ratios, ex-

Scale	Method	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS_alex( $\downarrow$ )	LPIPS_vgg( $\downarrow$ )	LPIPS_squeeze( $\downarrow$ )	SNLI( $\downarrow$ )
x2	Cubic	39.5039	0.9705	<b>0.0148</b>	0.0639	<b>0.0161</b>	<u>0.5318</u>
	Trilinear	40.6750	0.9757	<u>0.0201</u>	<b>0.0618</b>	<u>0.0189</u>	0.8450
	TVSRN [52]	<u>43.6167</u>	<u>0.9820</u>	0.0311	0.0824	0.0310	1.4307
	ArSSR [48]	42.6713	0.9799	0.0370	0.0913	0.0357	1.5738
	SAINT [32]	<b>44.3977</b>	<b>0.9833</b>	0.0361	0.0867	0.0354	1.6724
	CycleINR (Ours)	43.0137	0.9805	<u>0.0201</u>	<u>0.0625</u>	0.0206	<b>0.3527</b>
x3	Cubic	35.0140	0.9393	<u>0.0306</u>	0.1084	<u>0.0292</u>	<b>0.5322</b>
	Trilinear	36.5867	0.9515	0.0328	<u>0.0973</u>	0.0293	0.7422
	TVSRN [52]	<u>40.4857</u>	<u>0.9696</u>	0.0607	0.1419	0.0583	1.4389
	ArSSR [48]	39.3398	0.9659	0.0443	0.1135	0.0413	1.3026
	SAINT [32]	<b>40.8705</b>	<b>0.9711</b>	0.0625	0.1404	0.0599	1.7535
	CycleINR (Ours)	39.2748	0.9644	<b>0.0293</b>	<b>0.0902</b>	<b>0.0280</b>	<u>0.6674</u>
x5	Cubic	31.0470	0.8896	0.0562	0.1627	0.0480	<b>0.5336</b>
	Trilinear	32.6606	0.9106	<u>0.0525</u>	<u>0.1413</u>	<u>0.0433</u>	0.6682
	TVSRN [52]	<u>36.8459</u>	<u>0.9503</u>	0.0927	0.1989	0.0862	1.3593
	ArSSR [48]	35.1960	0.9394	0.0611	0.1485	0.0528	1.1714
	SAINT [32]	<b>36.9940</b>	<b>0.9519</b>	0.0996	0.2044	0.0951	1.6996
	CycleINR (Ours)	35.0022	0.9354	<b>0.0464</b>	<b>0.1289</b>	<b>0.0399</b>	<u>0.6050</u>

Table 1. Image quality evaluation of different methods. **bold** signifies the best, and underline represents the suboptimal. Note that PSNR and SSIM are not the sole determinants, and SNLI proves particularly valuable for comparing deep learning-based methods.

cept for TVSRN, which requires training separate models for super-resolution factors of x2, x3, and x5, respectively.

**Evaluation Metrics.** For the image quality evaluation, we first use peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) [46] as fundamental metrics for different super-resolution methods. However, acknowledging the bias of PSNR and SSIM towards favoring blurry images, where an over-smoothed super-resolution result could yield high scores under these metrics despite significant loss of image details, we also include learned perceptual image patch similarity (LPIPS) [54] in the evaluation. LPIPS scores are calculated using three backbones (AlexNet [21], VGGNet [38], and SqueezeNet [15]). In addition, we calculate the SNLI metric to measure the inconsistency between slices, often ignored by most volumetric super-resolution methods. For downstream segmentation task evaluation, we use the Dice similarity coefficient (DSC) [7] and normalized surface Dice (NSD) [30] with a tolerance of 1mm.

## 4.2. Super-resolution Results Evaluation

**Image Quality Evaluation.** Figure 3 visualizes the super-resolution results produced by different methods, and Table 1 presents the corresponding image quality metrics. In summary, traditional cubic and trilinear interpolation methods exhibit limitations in recovering HR images, particularly at large upscaling scales such as x5, where bone structures are notably challenging to recover, resulting in low PSNR and SSIM scores. While the three deep learning-

based methods demonstrate relatively favorable results, a shared challenge is their oversight of over-smoothness in the generated slices, as evident in both the visualization results and the SNLI metric. In Figure 3, the newly generated axial slices are over-smoothed, leading to the loss of fine vessel details in the liver. Moreover, in the sagittal and coronal views, horizontal lines emerge due to noise level inconsistencies between slices. The SNLI metric for these three deep learning-based methods is markedly elevated compared to other methods. However, with the enhancement of CCL, CycleINR demonstrates significant improvements, effectively recovering fine details without over-smoothing in the newly generated slices, and revealing the absence of visible horizontal lines in sagittal and coronal views. Despite a slight decrease in PSNR and SSIM scores, it is important to note, as mentioned in [54], that this decline does not necessarily signify a negative outcome. Furthermore, the SNLI metric demonstrates the superiority of the proposed CycleINR in preserving consistency between axial slices compared to other deep learning-based methods. This emphasizes the substantial effectiveness of employing CCL in mitigating over-smoothing issues in newly generated slices.

**Downstream Task Evaluation.** To make a comprehensive comparison of the super-resolution results generated by various deep learning models, we adopt a segmentation evaluation strategy, aligning with approaches from prior works [34, 43]. This strategy allows for a thorough assessment of the SR results, moving beyond visual appear-

INR	LAM	CCL	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS_alex( $\downarrow$ )	LPIPS_vgg( $\downarrow$ )	LPIPS_squeeze( $\downarrow$ )	SNLI( $\downarrow$ )
✓			<b>35.1960</b>	<b>0.9394</b>	0.0611	0.1485	0.0528	1.1714
✓	✓		<u>35.1395</u>	<u>0.9391</u>	0.0603	0.1478	0.0522	1.1395
✓		✓	34.9610	0.9350	<u>0.0470</u>	<u>0.1298</u>	<u>0.0405</u>	<u>0.6134</u>
✓	✓	✓	35.0022	0.9354	<b>0.0464</b>	<b>0.1289</b>	<b>0.0400</b>	<b>0.6050</b>

Table 2. Ablation study on the Local Attention Mechanism (LAM) and Cycle-Consistent Loss (CCL).

Scale	Method	DSC_L $\uparrow$	NSD_L $\uparrow$	DSC_T $\uparrow$	NSD_T $\uparrow$
x2	Cubic	0.9868	0.9568	<b>0.9069</b>	<b>0.8632</b>
	Trilinear	<u>0.9892</u>	<u>0.9622</u>	<b>0.9069</b>	<u>0.8575</u>
	TVSRN [52]	0.9886	0.9514	0.8552	0.8065
	ArSSR [48]	0.9856	0.9459	0.8516	0.7952
	SAINT [32]	0.9884	0.9540	0.8476	0.7904
	CycleINR	<b>0.9911</b>	<b>0.9648</b>	0.8730	0.8226
x3	Cubic	0.9625	0.7862	0.7621	0.6891
	Trilinear	0.9783	0.9037	0.7858	0.7083
	TVSRN [52]	0.9783	0.9105	0.7672	0.6892
	ArSSR [48]	0.9777	0.9056	0.7811	0.6988
	SAINT [32]	<u>0.9820</u>	<u>0.9250</u>	<u>0.8042</u>	<u>0.7338</u>
	CycleINR	<b>0.9832</b>	<b>0.9294</b>	<b>0.8087</b>	<b>0.7381</b>
x5	Cubic	0.8678	0.4475	0.5644	0.4386
	Trilinear	0.9487	0.6873	<u>0.6755</u>	0.5385
	TVSRN [52]	0.9481	0.8228	0.6440	<u>0.5431</u>
	ArSSR [48]	0.9589	0.7907	0.6582	0.5297
	SAINT [32]	<u>0.9619</u>	<u>0.8561</u>	0.6006	0.5050
	CycleINR	<b>0.9620</b>	<b>0.8566</b>	<b>0.6881</b>	<b>0.5640</b>

Table 3. Downstream task evaluation on the MSD liver tumor dataset. Comparison of segmentation across different super-resolution methods with regard to the segmentation on the original 1mm data. ‘L’ and ‘T’ represent the liver and tumor, respectively. Results are compared to segmentation on the original 1mm image, not the ground truth mask, to evaluate the fidelity of the upscaled image to the original 1mm characteristics.

ance to provide a more clinically relevant measure of performance. As depicted in Table 3, the segmentation results of CycleINR align better with the reference compared to other methods, particularly superior at higher upsampling scales. While traditional interpolation methods might retain advantages in specific scenarios with smaller upscaling ratios (e.g., x2), it is crucial to highlight that our proposed CycleINR consistently exhibits superior performance compared to other deep learning methods across various upscaling ratios. This underscores the clinical relevance of our method and establishes its efficacy in achieving better results for downstream tasks.

### 4.3. Ablation Studies

We perform ablation studies to assess the impact of two key components in our CycleINR framework: the local atten-

tion mechanism (LAM) during latent code grid sampling, and the cycle-consistent loss (CCL) aimed at minimizing inter-slice consistency. The contributions of these components to the overall performance are shown in Table 2.

**Local Attention Mechanism (LAM).** The local attention mechanism in our process effectively combines physical distance and latent code similarities to aggregate weights for latent code representation. Drawing inspiration from bilateral filtering, our LAM employs attention mechanisms for improved expressive power. As shown in Table 2, integrating LAM into the latent code grid sampling process enhances the LPIPS metric [54] by preserving edges and mitigating excessive smoothing, similar to the principles of traditional bilateral filtering [42].

**Cycle-Consistent Loss (CCL).** The purpose of CCL is to impose a constraint ensuring consistency between the generated and original images in terms of image characteristics, such as noise level. As shown in Table 2, the integration of CCL leads to a substantial improvement in both the LPIPS and SNLI metrics. Visual results in Figure 3 further demonstrate that CCL enhances CycleINR by effectively mitigating over-smoothing, as evidenced by significant reductions in LPIPS and SNLI scores.

## 5. Conclusion

We present CycleINR, a novel 3D medical image super-resolution approach. Unlike traditional approaches, CycleINR reframes the task as function approximation, leveraging the continuity of the implicit neural representation function. This enables efficient high-resolution image generation without specific training for upsampling ratios. The introduced cycle-consistent loss mitigates over-smoothing, a local attention mechanism further enhances performance, and a novel metric, slice-wise noise level inconsistency, assesses the quality of 3D medical images. Overall, CycleINR shows promise for volumetric super-resolution, offering efficiency and flexibility for enhanced visual quality, and robust downstream analysis. For future work, extending CycleINR to diverse medical imaging modalities and evaluating its performance in larger clinical scenarios would offer insights into generalizability. Exploring applications beyond the medical field, like video interpolation, presents a compelling direction for adapting CycleINR to handle the challenges in temporal sequences.



## References

- [1] Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, Annette Kopp-Schneider, Bennett A Landman, Geert Litjens, Bjoern Menze, Olaf Ronneberger, Ronald M Summers, et al. The medical segmentation decathlon. *Nature communications*, 13(1):4128, 2022. 5
- [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. 3
- [3] Thierry Blu, Philippe Thévenaz, and Michael Unser. Linear interpolation revitalized. *IEEE Transactions on Image Processing*, 13(5):710–719, 2004. 2
- [4] Hao-Wei Chen, Yu-Syuan Xu, Min-Fong Hong, Yi-Min Tsai, Hsien-Kai Kuo, and Chun-Yi Lee. Cascaded local implicit transformer for arbitrary-scale super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18257–18267, 2023. 3
- [5] Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8628–8638, 2021. 3, 4
- [6] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019. 2
- [7] Lee R Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302, 1945. 7
- [8] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015. 2
- [9] Chao Dong, Chen Change Loy, and Xiaoou Tang. Accelerating the super-resolution convolutional neural network. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 391–407. Springer, 2016. 2
- [10] David L Donoho and Iain M Johnstone. Ideal spatial adaptation by wavelet shrinkage. *biometrika*, 81(3):425–455, 1994. 5
- [11] Jinglong Du, Zhongshi He, Lulu Wang, Ali Gholipour, Zexun Zhou, Dingding Chen, and Yuanyuan Jia. Super-resolution reconstruction of single anisotropic 3d mr images using residual convolutional neural network. *Neurocomputing*, 392:209–220, 2020. 4
- [12] Chaowei Fang, Liang Wang, Dingwen Zhang, Jun Xu, Yixuan Yuan, and Junwei Han. Incremental cross-view mutual distillation for self-supervised medical ct synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20677–20686, 2022. 3
- [13] Chaowei Fang, Dingwen Zhang, Liang Wang, Yulun Zhang, Lechao Cheng, and Junwei Han. Cross-modality high-frequency transformer for mr image super-resolution. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1584–1592, 2022. 2
- [14] Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. Local deep implicit functions for 3d shape. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4857–4866, 2020. 3
- [15] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016. 7
- [16] Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2):203–211, 2021. 5
- [17] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, Thomas Funkhouser, et al. Local implicit grid representations for 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6001–6010, 2020. 3
- [18] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1646–1654, 2016. 2
- [19] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Deeply-recursive convolutional network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1637–1645, 2016. 2
- [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. 7
- [22] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017. 2
- [23] Jaewon Lee and Kyong Hwan Jin. Local texture estimator for implicit representation function. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1929–1938, 2022. 3
- [24] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 136–144, 2017. 2
- [25] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020. 3
- [26] Ying-Tian Liu, Yuan-Chen Guo, and Song-Hai Zhang. Enhancing multi-scale implicit learning in image super-resolution with integrated positional encoding. *arXiv preprint arXiv:2112.05756*, 2021. 3

- [27] 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*, pages 4460–4470, 2019. [2](#)
- [28] 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. *Communications of the ACM*, 65(1):99–106, 2021. [3](#)
- [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. [3](#)
- [30] Stanislav Nikolov, Sam Blackwell, Alexei Zverovitch, Ruheena Mendes, Michelle Livne, Jeffrey De Fauw, Yojan Patel, Clemens Meyer, Harry Askham, Bernardino Romera-Paredes, et al. Deep learning to achieve clinically applicable segmentation of head and neck anatomy for radiotherapy. *arXiv preprint arXiv:1809.04430*, 2018. [7](#)
- [31] 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/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019. [2](#)
- [32] Cheng Peng, Wei-An Lin, Haofu Liao, Rama Chellappa, and S Kevin Zhou. Saint: spatially aware interpolation network for medical slice synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7750–7759, 2020. [3](#), [5](#), [6](#), [7](#), [8](#)
- [33] Cheng Peng, S Kevin Zhou, and Rama Chellappa. Da-vs: domain adaptable volumetric super-resolution for medical images. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24*, pages 75–85. Springer, 2021. [2](#), [3](#)
- [34] Cheng Peng, Wei-An Lin, Haofu Liao, Rama Chellappa, and Shaohua Kevin Zhou. Deep slice interpolation via marginal super-resolution, fusion, and refinement. In *State of the Art in Neural Networks and Their Applications*, pages 133–145. Elsevier, 2023. [7](#)
- [35] Fitsum Reda, Janne Kontkanen, Eric Tabellion, Deqing Sun, Caroline Pantofaru, and Brian Curless. Film: Frame interpolation for large motion. In *European Conference on Computer Vision*, pages 250–266. Springer, 2022. [2](#)
- [36] Feng Shi, Jian Cheng, Li Wang, Pew-Thian Yap, and Dinggang Shen. Lrtv: Mr image super-resolution with low-rank and total variation regularizations. *IEEE transactions on medical imaging*, 34(12):2459–2466, 2015. [2](#)
- [37] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. 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*, pages 1874–1883, 2016. [2](#)
- [38] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [7](#)
- [39] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *Advances in Neural Information Processing Systems*, 32, 2019. [3](#)
- [40] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020. [2](#), [3](#)
- [41] Philippe Thévenaz, Thierry Blu, and Michael Unser. Chapter 28 - image interpolation and resampling. In *Handbook of Medical Image Processing and Analysis (Second Edition)*, pages 465–493. Academic Press, Burlington, second edition edition, 2009. [1](#)
- [42] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *Sixth international conference on computer vision (IEEE Cat. No. 98CH36271)*, pages 839–846. IEEE, 1998. [8](#)
- [43] Jiancong Wang, Yuhua Chen, Yifan Wu, Jianbo Shi, and James Gee. Enhanced generative adversarial network for 3d brain mri super-resolution. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3627–3636, 2020. [7](#)
- [44] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, pages 0–0, 2018. [2](#)
- [45] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1905–1914, 2021. [2](#)
- [46] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. [7](#)
- [47] Lauren H Williams and Trafton Drew. What do we know about volumetric medical image interpretation?: A review of the basic science and medical image perception literatures. *Cognitive Research: Principles and Implications*, 4:1–24, 2019. [1](#)
- [48] Qing Wu, Yuwei Li, Yawen Sun, Yan Zhou, Hongjiang Wei, Jingyi Yu, and Yuyao Zhang. An arbitrary scale super-resolution approach for 3d mr images via implicit neural representation. *IEEE Journal of Biomedical and Health Informatics*, 27(2):1004–1015, 2022. [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [49] Xingqian Xu, Zhangyang Wang, and Humphrey Shi. Ultrasr: Spatial encoding is a missing key for implicit image function-based arbitrary-scale super-resolution. *arXiv preprint arXiv:2103.12716*, 2021. [3](#)
- [50] Jingyu Yang, Sheng Shen, Huanjing Yue, and Kun Li. Implicit transformer network for screen content image continuous super-resolution. *Advances in Neural Information Processing Systems*, 34:13304–13315, 2021. [3](#)
- [51] Wenming Yang, Xuechen Zhang, Yapeng Tian, Wei Wang, Jing-Hao Xue, and Qingmin Liao. Deep learning for single

- image super-resolution: A brief review. *IEEE Transactions on Multimedia*, 21(12):3106–3121, 2019. 2
- [52] Pengxin Yu, Haoyue Zhang, Han Kang, Wen Tang, Corey W Arnold, and Rongguo Zhang. Rplhr-ct dataset and transformer baseline for volumetric super-resolution from ct scans. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 344–353. Springer, 2022. 2, 3, 6, 7, 8
- [53] Kai Zhang, Wangmeng Zuo, Shuhang Gu, and Lei Zhang. Learning deep cnn denoiser prior for image restoration. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3929–3938, 2017. 2
- [54] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 2, 7, 8
- [55] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 4