

This WACV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

SeCo-INR: Semantically Conditioned Implicit Neural Representations for Improved Medical Image Super-Resolution

Mevan Ekanayake, Zhifeng Chen, Gary Egan, Mehrtash Harandi, Zhaolin Chen Monash University, Melbourne, Australia

{mevan.ekanayake, zhifeng.chen, gary.egan, mehrtash.harandi, zhaolin.chen}@monash.edu

Abstract

Implicit Neural Representations (INRs) have recently advanced the field of deep learning due to their ability to learn continuous representations of signals without the need for large training datasets. Although INR methods have been studied for medical image super-resolution, their adaptability to localized priors in medical images has not been extensively explored. Medical images contain rich anatomical divisions that could provide valuable local prior information to enhance the accuracy and robustness of INRs. In this work, we propose a novel framework, referred to as the Semantically Conditioned INR (SeCo-INR), that conditions an INR using local priors from a medical image, enabling accurate model fitting and interpolation capabilities to achieve super-resolution. Our framework learns a continuous representation of the semantic segmentation features of a medical image and utilizes it to derive the optimal INR for each semantic region of the image. We tested our framework using several medical imaging modalities and achieved higher quantitative scores and more realistic super-resolution outputs compared to state-of-the-art methods.

1. Introduction

Advancements in deep learning (DL) research have revolutionized medical image super-resolution [3, 22, 35, 36]. These DL methods often rely on supervised learning techniques to transform low-resolution images into highresolution ones, using large cohorts of datasets. While this approach demonstrates remarkable success in specific scenarios, its performance deteriorates when faced with outof-distribution cases, as many models struggle to generalize beyond the extent of the training data. Furthermore, for medical imaging applications, it is challenging to obtain massive amounts of training data due to privacy concerns and high acquisition costs. As a result, exploring unsupervised or semi-supervised DL methods for image superresolution has become important with the ultimate goal of advancing the precision and reliability in medical imaging [9].

Implicit Neural Representations (INRs) have recently advanced the field of signal representation due to their ability to learn continuous representations of signals using a discrete set of samples without the requirement of large training datasets [27]. In contrast to conventional approaches that discretely store signal values on coordinate grids, INRs train neural networks, particularly Multilayer Perceptrons (MLPs), to represent continuous signals [18]. The objective is to approximate the intricate connection between coordinates and their corresponding signal values, ultimately producing a continuous representation of the signal. Due to their immense success in signal representation and solving inverse problems, INRs have been utilized in many tasks recently [2, 5, 19]. Some of the recent medical imaging-specific advancements include arbitrary scale 3D super-resolution [30, 31], assessing tumor progression with sparsely sampled images [25], jointly learning different medical imaging modalities of complementary views [16] and reconstruction of isotropic images from anisotropic images [34] belonging to a variety of medical imaging modalities like magnetic resonance imaging (MRI) and computed tomography (CT) imaging.

While INR techniques have been recently explored to enhance medical image super-resolution [16, 33], their potential to benefit from localized anatomical priors has not been thoroughly investigated. Medical imaging modalities contain valuable local priors that could effectively enhance the representation of images [4, 21]. For instance, different cortical regions of a brain image exhibit various contrasts and geometrically varied anatomical structures due to the unique signal attributes captured by the imaging system [6, 14]. In this work, we hypothesize that this semantic contrast and shape information can benefit INRs in learning the continuous representation of a medical image. To this end, we proposed a framework, referred to as the Semantically Conditioned INR (SeCo-INR), which learns a continuous representation of the semantic segmentation features of an image and utilizes it to derive the optimal INR parameters for each semantic region of the image through a data-driven process. Through comprehensive testing on various medical imaging datasets, our framework demonstrated higher quantitative scores and produced realistic super-resolution outputs, outperforming the state-ofthe-art methods.

2. Methods

2.1. Implicit Neural Representations

Implicit Neural Representations encode an underlying continuous signal, G via a deep neural network, $f_{\theta} : \mathbb{R}^M \to$ \mathbb{R}^{P} , where M is the dimension of input coordinates and P is the dimension of signal value at that particular coordinate. As the deep neural network, often an MLP with 3-4 layers is utilized with suitable activation functions in between. This network learns a functional mapping between the coordinates, $v \in \mathbb{R}^M$ (e.g., v = (x, y) for 2D, v = (x, y, z) for 3D, etc.) and the underlying signal, $G(v) \in \mathbb{R}^{P}$ (e.g., intensity, color, occupancy in 3D space, camera view, etc.) [32]. This can be achieved by optimizing the discrepancy between a set of discrete samples from G(v) and the corresponding network output $f_{\theta}(v)$, given by $||f_{\theta}(v) - G(v)||_2^2$. Each layer of f_{θ} is activated using $y_l = \sigma(W_l y_{l-1} + b_l)$, where l = 1, 2, ..., L - 1, and L is the total number of layers in the MLP. Here $W_l \in \mathbb{R}^{d_{l-1} \times d_l}$ and $b_l \in \mathbb{R}^{d_l}$ denote the weights matrix and the bias vector at the l^{th} layer of the network. $y_l \in \mathbb{R}^{d_l}$ is the output of the l^{th} layer and σ is the activation function.

When Rectified Linear Unit (ReLU) is used as the activation function, *i.e.*, $\sigma(x) = \max(0, x)$, the network tends to favor the representation of low-frequency signals which often leads to lower-quality signal representations [24]. To address this issue, Sitzmann *et al.* [27] proposed sinusoidal representation networks (SIRENs), which utilized periodic activation functions to fit the signal as well as its higher-order derivatives, *i.e.*, $\sigma(x) = \sin(\omega_0 x)$. Although SIREN outperforms ReLU in terms of image representation and various other tasks, there remains room for improvement in managing the sine activation function more effectively. Subsequent works demonstrated that the choice of ω_0 is critical to achieving higher representation accuracies [24] and explored control over the parameters of an extended sinusoidal function for better representation quality [11].

2.2. SeCo-INR

Our proposed SeCo-INR framework contains three major components: *Adaptive SIREN Network*, *Pixel Class Representation Network*, and *Conditioner Network* which we discuss in detail below. The overall framework is demonstrated in Fig. 1.

2.2.1 Adaptive SIREN Network

This is the central part of the framework that maps the input coordinates to their corresponding output intensities of the image. We define this network, $f_{\theta} : \mathbb{R}^M \to \mathbb{R}^P$ as an MLP and periodic nonlinear activation functions with additional parametrization as below:

$$y_l = p \sin(q\omega_0(W_l y_{l-1} + b_l) + r) + s.$$
 (1)

Note that the above function is reduced to SIREN [27] if p = 1, q = 1, r = 0, s = 0. In our framework, the values of p, q, r, s are dynamically learned in every iteration of the training process, with assistance from the semantic prior information of the underlying image. Each parameter uniquely influences the characterization of the representation. The parameter p sets the amplitude of the periodic signal and affects the strength of the activation, aiding in noise suppression. The parameter q controls the frequency scaling of the periodic signal, focusing on the fine details of the representation and reducing high-frequency noise. The parameter r adjusts the phase shift of the periodic signal, influencing the spatial arrangement in the representation. Lastly, the parameter s establishes the vertical shift of the periodic signal, managing the overall contrast [11].

2.2.2 Pixel Class Representation Network

In our work, we introduce a Pixel Class Representation Network to learn a continuous representation of the semantic segmentation map of a medical image. Such a segmentation map contains pixel-wise physical localization pertinent to the target image. The Pixel Class Representation Network $g(\phi)$, which is an INR model with a softmax output layer is trained to iteratively learn a continuous representation of segmentation mask S. It takes the coordinates v as the input to generate an output semantic class distribution of those coordinates, $\hat{S}(v) = g_{\phi}(v)$. We train the weights of this network using a cross-entropy loss as $-\sum_i S(v)_i \log(\hat{S}(v)_i)$, considering S(v) as the ground truth semantic class distribution for each coordinate v. The details about the ground truth segmentation masks are given in Section 3.1.

2.2.3 Conditioner Network

This network $h_{\psi} : \mathbb{R}^h \to \mathbb{R}^{4L}$ utilizes an MLP structure and takes the local prior latent code $\hat{S}(v) \in \mathbb{R}^h$ (derived by the Pixel Class Representation Network) as input and produces the optimal parameters for each semantic region of the underlying image, *i.e.*, p(v), q(v), r(v), s(v) = $h_{\psi}(\hat{S}(v))$ which is then utilized to condition each layer of the Adaptive SIREN Network separately. The Conditioner Network learns to generate the optimum parameter values adaptively from the localized prior information given by the learned semantic segmentation map, \hat{S} .



Figure 1. Illustration of SeCo-INR: Semantically Conditioned Implicit Neural Representations framework for Medical Image Super-Resolution. The Pixel Class Representation Network learns a continuous representation of the semantic segmentation mask, which provides conditioning to the Adapting SIREN Network via the Conditioner Network. This process enables the learning of a more accurate implicit neural representation of the underlying medical image, allowing robust super-resolution capabilities.

2.2.4 Loss Function

The Loss function for the training of the overall framework is based on the reconstruction loss which enforces the Adaptive SIREN Network to fit the target image and the classification loss which enforces the Pixel Class Representation Network to learn the semantic class distributions of the coordinates. We further add a regularization term to the loss function, which mandates that the parameters p, q, r, smaintain non-negative values. This type of regularization has been proven to steer the model toward more applicable solutions, accelerate model convergence, and diminish the chance of the model getting stuck in local minima during training. All trainable networks of the framework are trained in a data-driven end-to-end approach for a given target image, I(v), and the overall loss function can be presented below:

$$\underset{\theta,\psi,\phi}{\operatorname{argmin}} \mathbb{E}_{v \in V} \left[\|\hat{I}(v) - I(v)\|_{2}^{2} + \beta \sum_{i} S(v)_{i} \log(\hat{S}(v)_{i}) \right]$$
(2)

$$\text{s.t.} \quad p(v) \geq 0, \quad q(v) \geq 0, \quad r(v) \geq 0, \quad s(v) \geq 0, \quad \forall v$$

where $\hat{I}(v) = f_{\theta}(v|p(v), q(v), r(v), s(v)),$ $p(v), q(v), r(v), s(v) = h_{\psi}(\hat{S(v)})$ and $\hat{S}(v) = g_{\phi}(v).$ The hyperparameter β determines the trade-off between reconstruction loss and the classification loss.

2.2.5 Super-resolution Inference

Having trained the framework and obtained the optimum model parameters, h_{ψ}^* , f_{θ}^* , and g_{ϕ}^* on a given low-resolution image until convergence, the super-resolution inference can be obtained by first generating a set of extended pixel coordinates (*i.e.* super-resolution), v'. Next, the Pixel Class Representation Network can be used to generate a highresolution semantic segmentation mask using:

$$\hat{S}(v') = g_{\phi}^{*}(v')$$
 (3)

Then, the Conditioner Network can be used to generate the optimum p, q, r, s parameters for each layer of the Adaptive SIREN Network using:

$$p(v'), q(v'), r(v'), s(v') = h_{\psi}^*(\hat{S}(v')).$$
 (4)

Finally, the image intensities of the super-resolution image can be derived using the optimized Adaptive SIREN Network conditioned on the optimized parameters:

$$I(v') = f_{\theta}^*(v'|p(v'), q(v'), r(v'), s(v')).$$
(5)

The final image, I' can be formed by ordering the pixel intensities I(v'). The detailed algorithm of our framework is given in Algorithm 1.

Algorithm 1 SeCo-INR

- 1: **Input:** low-resolution image I, Semantic Segmentation map S, Coordinates $v \in V$, hyperparameter β
- 2: **Output:** high-resolution image I'
- 3: Initialize the Adaptive SIREN Network f_{θ} , Conditioner Network h_{ψ} , Pixel Class Representation Network g_{ϕ}
- 4: for each epoch do
- Generate semantic segmentation class distributions: $\hat{S}(v) = g_{\phi}(v)$ 5:
- Generate p, q, r, s parameters: $p(v), q(v), r(v), s(v) = h_{\psi}(\hat{S}(v))$ 6:
- Generate pixel intensities: $\hat{I}(v) = f_{\theta}(v|p(v), q(v), r(v), s(v))$ 7:
- Compute loss $\mathcal{L} = \underset{v \in V}{\mathbb{E}} \left[\|\hat{I}(v) I(v)\|_{2}^{2} + \beta \sum_{i} S(v)_{i} \log(\hat{S}(v)_{i}) \right]$ Update θ, ψ, ϕ to minimize \mathcal{L} 8:
- 9:

10: end for

- 11: Generate extended coordinates for super-resolution: v'
- Generate optimized semantic segmentation class distributions: $\hat{S}(v') = g_{\phi}^*(v')$ 12:
- 13: Generate optimized p, q, r, s parameters: $p(v'), q(v'), r(v'), s(v') = h_{\psi}^*(\hat{S}(v'))$
- 14. Generate pixel intensities: $I(v') = f_{\theta}^*(v'|p(v'), q(v'), r(v'), s(v'))$
- Arrange I(v'); $\forall v'$ to form final image, I'15:
- 16: return I'

2.3. Rationale behind feeding local semantic features

The rationale behind the utilization of semantic information is to condition the adaptive SIREN Network with the local prior knowledge, thereby learning distinct sets of p, q, r, s parameters for each semantic region of the image. For instance, the segmentation masks of a brain MRI (e.g., gray and white matter divisions, brain tumor segmentation, abnormalities, sub-cortical structures derived from a brain atlas, etc.) play a pivotal role in distinguishing different cortical regions of a brain image that exhibit various contrasts due to proton density, biophysical relaxation parameters, and geometrically varied anatomical structures [6, 14]. Similarly, segmentation masks of an abdominal CT image, which delineate various organs (e.g., liver, kidneys, stomach, spleen, pancreas, bowel, gallbladder, etc.) are essential for accurately analyzing the anatomical structures within the abdomen that exhibit various contrasts and geometrically varied organs [15].

3. Experiments

3.1. Datasets

Our proposed framework was validated using multiple medical imaging datasets including brain MRI datasets and abdominal CT datasets. The brain MRI datasets included fluid-attenuated inversion recovery (FLAIR) images from the fastMRI dataset [20] as well as T1-weighted (T1w) images from the Brain Tumor Segmentation (BraTS) datasets [17]. The abdominal CT datasets are from the multi-atlas labeling challenge which had been captured during portal venous contrast phase [12]. For the experiments in this paper, we use the axial brain slices from the fastMRI data, which have a spatial dimension of 320×320 ; axial brain slices from the BraTS MRI data, which have a spatial dimension of 240×240 ; and axial abdominal CT slices which have a spatial dimension of 512×512 . To obtain the ground truth segmentation masks, S, for the fastMRI brain data, we employed the FSL's Brain Extraction Tool (BET) and the Automated Segmentation Tool (FAST) [10, 28] which uses a combination of intensity-based thresholding, surface deformation, and atlas-guided techniques. For abdominal CT data, we used the multi-atlas labeled data [12] which had used manually labeled 13 abdominal organs verified by a radiologist on a volumetric basis. For BraTS brain MRI data, we used the labeled data [17] segmented manually and approved by experienced neuro-radiologists.

3.2. Experiments

For super-resolution experiments, we obtained images with the full spatial resolution and retrospectively resized the images by factors of $1.5\times$, $2\times$, and $2.5\times$ for fastMRI brain data, $1.5 \times$ and $2 \times$ for BraTS MRI brain data, and $2 \times$, $3\times$, and $4\times$ for abdominal CT data. These retrospectively obtained low-resolution images contained pixelated artifacts at various degrees. Then, we implemented the superresolution methods and assessed the outputs both qualitatively and quantitatively (w.r.t. to root-mean-square-error (RMSE), peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) following the standard definitions in [20]). To compare the performance of our framework, we implemented several state-of-the-art INR-based algorithms: ReLU with Positional-Encoding (ReLU+P.E.) [29], Gauss [23]. Multiplicative Filter Networks (MFN) [7]. SIREN [27], Wavelet INR (WIRE) [24], and Fourier Reparameterization (FR) [26]. For the abdominal CT data,



Figure 2. Qualitative evaluation between different INR methods for the Abdominal CT data at resolution factors of $2\times$, $3\times$, and $4\times$. SeCo-INR produces realistic outputs of abdominal organs with sharp edges, less noise, and less blurring.

fastMRI brain data, and BraTS brain tumor data, the qualitative results are illustrated in Fig. 2, Fig. 3, and Fig. 4, respectively, and the quantitative scores are tabulated in Table 1, Table 2, and Table 3, respectively.

3.3. Implementation

In our experiments, a 5-layer MLP was utilized as the Adaptive SIREN Network consisting of 256 neurons in each layer. To implement the constraints on the p, q, r, s parameters, we enforced ReLU functions and added to the overall loss to penalize the generation of negative values [11]. We implemented the framework on a single NVIDIA GeForce RTX 3060 GPU with 12GB memory using the PyTorch framework. The optimization process involved using the

Adam optimizer with a learning rate scheduler, *i.e.* decreasing the learning rate by 0.1 after each epoch. We set $\beta = 1$. For convergence, ReLU+P.E., Gauss, and MFN models were trained for 2000 epochs, whereas SIREN, WIRE, FR, and SeCo-INR models were trained for 1000 epochs. All codes for reproducibility are attached to the Supplementary Materials.

3.4. Ablation Studies

We conducted an ablation study to demonstrate that the performance gains are indeed due to the incorporation of semantic segmentation information in the INR framework. We utilized the extended SIREN network without feeding semantic features into the model. Table 5 presents the quan-



Figure 3. Qualitative evaluation between different INR methods for the fastMRI brain data at resolution factors of $1.5 \times$, $2 \times$, and $2.5 \times$. SeCo-INR produces realistic outputs of the brain with intricate anatomical details preserved (see yellow arrow), and less noise and blurring.

resolution	quantitative	Method						
factor	metric	ReLU+P.E [29]	Gauss [23]	MFN [7]	SIREN [27]	FR [26]	SeCo-INR (Ours)	
	RMSE (\downarrow)	0.0207	0.0121	0.0152	0.0153	0.0147	0.0099	
$2 \times$	SSIM (†)	0.9335	0.9564	0.9292	0.9559	<u>0.9571</u>	0.9836	
	PSNR (†)	33.8140	<u>38.4625</u>	36.5725	36.4029	36.7377	40.2139	
	RMSE (\downarrow)	0.0233	0.0199	0.0219	0.0183	0.0181	0.0160	
$3 \times$	SSIM (†)	0.9312	0.9018	0.8778	0.9510	<u>0.9518</u>	0.9660	
	PSNR (†)	32.7672	34.1146	33.3290	34.8399	<u>34.9703</u>	36.0422	
	RMSE (\downarrow)	0.0249	0.0306	0.0362	0.0217	0.0216	0.0211	
$4 \times$	SSIM (†)	0.9257	0.8046	0.7588	0.9439	0.9445	0.9340	
	PSNR (†)	32.1903	30.3293	28.9740	33.3819	<u>33.4216</u>	33.6130	

Table 1. Quantitative evaluation between different INR methods for the Abdominal CT data at resolution factors of $2\times$, $3\times$, and $4\times$. (average metrics across the entire dataset are presented). The best performance is in **boldface** and the second best is <u>underlined</u>.

titative scores for the fastMRI data at a resolution factor of $2 \times$ and they clearly indicate that the performance gains stem from the semantic features. This validates that a key strength of SeCo-INR is its ability to leverage additional strong medical image segmentation information.

We also compared the computational costs of each algorithm on the fastMRI brain dataset at a resolution factor of $2 \times$ by calculating the time taken on average, per image to reach a PSNR level of 30.0 OR an SSIM level of 0.90. The results are presented in Table 4.



Figure 4. Qualitative evaluation between different INR methods for the BraTS brain MRI data at resolution factors of $1.5 \times$ and $2 \times$. SeCo-INR produces realistic outputs of the tumor regions with less noise and blurring.

resolution	quantitative			Method				
factor	metric	ReLU+P.E [29]	Gauss [23]	MFN [7]	SIREN [27]	WIRE [24]	FR [26]	SeCo-INR (Ours)
	RMSE (\downarrow)	0.0255	0.0224	0.0217	0.0206	0.0217	0.0202	0.0154
$1.5 \times$	SSIM (†)	0.9004	0.9146	0.9196	0.9195	0.9202	<u>0.9219</u>	0.9606
	PSNR (†)	32.0217	33.1032	33.3939	33.8255	33.3758	34.0221	36.3731
	RMSE (\downarrow)	0.0270	0.0258	0.0267	0.0226	0.0239	0.0222	0.0210
$2\times$	SSIM (†)	0.8951	0.8982	0.8893	0.9163	0.9117	<u>0.9180</u>	0.9340
	PSNR (†)	31.5009	31.8457	31.5917	33.0093	32.5529	<u>33.1658</u>	33.6511
	RMSE (\downarrow)	0.0310	0.0334	0.0364	0.0270	0.0282	0.0259	0.0263
$2.5 \times$	SSIM (†)	0.8823	0.8565	0.8295	0.9024	0.8861	0.9030	0.9030
	PSNR (†)	30.3102	29.5691	28.8997	31.4860	31.0746	31.8485	<u>31.6760</u>

Table 2. Quantitative Evaluation between different INR methods for the fastMRI brain data at resolution factors of $1.5 \times$, $2 \times$, and $2.5 \times$. (average metrics across the entire dataset are presented). The best performance is in **boldface** and the second best is underlined.

Our overall framework simultaneously learns the semantic segmentation mask as well as the underlying image using the Pixel Class Representation Network and the Adaptive SIREN Network, respectively. Figure 5 demonstrates an example segmentation mask from the fastMRI brain dataset learned iteratively.

4. Results and Discussion

As seen in Fig. 2, SeCo-INR is able to achieve highquality CT super-resolution with minimal noise and blurring. The $2\times$ example highlights the kidney region where Seco-INR is able to reconstruct the edges and contrast information accurately, whereas Gauss and MFN methods produce noisy outputs, and SIREN and FR methods pro-



Figure 5. An example from the fastMRI brain dataset shows the evolution of the segmentation mask learning process using the Pixel Class Representation Network.

duce blurry outputs. The $3 \times$ example highlights the spinal region where most of the pixelated artifacts are rectified in the SeCo-INR output compared to other methods. The $4 \times$ example highlights the spleen region where SeCo-INR

resolution	quantitative			Method				
factor	metric	ReLU+P.E [29]	Gauss [23]	MFN [7]	SIREN [27]	WIRE [24]	FR [26]	SeCo-INR (Ours)
	RMSE (↓)	0.0248	0.0220	0.0234	0.0213	0.0207	0.0208	0.0191
$1.5 \times$	SSIM (†)	0.9436	0.9556	0.9409	0.9546	<u>0.9626</u>	0.9562	0.9732
	PSNR (†)	32.1597	33.1880	32.6338	33.4810	<u>33.7224</u>	33.6702	34.4158
	RMSE (↓)	0.0281	0.0279	0.0350	0.0205	0.0225	0.0201	0.0205
$2 \times$	SSIM (†)	0.9219	0.9170	0.8797	0.9559	0.9470	<u>0.9577</u>	0.9604
	PSNR (†)	31.0791	31.0991	29.1585	33.7836	32.9758	33.9721	<u>33.7852</u>

Table 3. Quantitative evaluation between different INR methods for the BraTS brain MRI data at resolution factors of $1.5 \times$ and $2 \times$. (average metrics across the entire dataset are presented). The best performance is in **boldface** and the second best is underlined.

ReLU+P.E [29]	Gauss [23]	MFN [7]	SIREN [27]	WIRE [24]	FR [26]	SeCo-INR (Ours)
4.19	<u>0.51</u>	3.39	0.65	4.66	0.66	0.45

Table 4. Comparison of computational time (in seconds) of each algorithm for $2 \times$ super-resolution on the fastMRI brain data. The best performance is in **boldface** and the second best is underlined.

	RMSE (\downarrow)	SSIM (†)	PSNR (†)
w/o semantic info	0.0231	0.9112	32.8125
SeCo-INR	0.0210	0.9340	33.6511

Table 5. Ablation study: Comparison of quantitative scores on the fastMRI data at a resolution factor of $2 \times$ with and without feeding semantic information to the extended INR model.

demonstrates superiority in mitigating noise during superresolution and producing realistic outputs. These observations align with the quantitative scores in Table 1, where SeCo-INR holds the best performance in almost all superresolution levels and metrics.

As depicted in Fig. 3 and Fig. 4, our framework demonstrates the overall best image resolution and SNR for the super-resolution task in MRI, especially at higher resolution levels. While all methods successfully rectify the pixelated artifacts of the low-resolution MRI images, MFN, Gauss, WIRE, and FR methods produce noisier images and ReLU+P.E. and SIREN methods tend to overly smooth the image, resulting in failure to reconstruct intricate anatomical structures of the brain (see yellow arrow in the $2 \times$ example). In contrast, SeCo-INR accurately reconstructs intricate edge details and eliminates pixelation artifacts, enabling higher-resolution MRI images with minimal noise. These results are further confirmed by the quantitative scores presented in Table 2, where SeCo-INR outperformed all other methods at $1.5 \times$ and $2 \times$ super-resolution, w.r.t. RMSE, PSNR, and SSIM. The brain tumor images in Fig. 4 show similar observations where SeCo-INR demonstrates the overall best image resolution and SNR specifically in the tumor region of the brain. Methods like ReLU+P.E., SIREN, and FR methods tend to overly smooth the tumor region whereas the other methods produced noisy outputs hindering the accurate visibility of the brain tumor,

especially at $2 \times$ super-resolution. The quantitative scores confirmed these observations where SeCo-INR dominated in all the metrics computed.

As evident by the results in Table 4, our SeCo-INR framework has the added benefit of faster convergence where it reaches certain threshold quantitative scores using minimal training time. This confirms SeCo-INR as a viable solution in practical medical diagnostic environments where faster computation is critical. Not only in training, SeCo-INR provides the additional benefit of automatically selecting the activation hyperparameters in contrast to other methods like SIREN and WIRE which are extremely sensitive to the activation hyperparameters [24] and are difficult to fine-tune for a given dataset.

One of the limitations of SeCo-INR is the requirement of ground truth semantic segmentation masks. For the experiments in this paper, we utilized the readily available segmentation masks or generated masks using existing tools. Nevertheless, many other tools are available to obtain segmentation masks for a variety of anatomical regions, including simple thresholding and clustering techniques, atlasbased methods [8], or recent AI tools such as MedSAM [13] and MONAI [1].

5. Conclusion

This paper presents SeCo-INR, an INR-based medical image super-resolution method that leverages local prior information derived from a semantic segmentation mask of the underlying image. SeCo-INR enables an INR network to learn local region-specific parameters for the underlying medical image, thereby achieving accurate superresolution. The results have been validated both qualitatively and quantitatively on MRI and CT data, and the method appears to be a promising solution for medical image super-resolution.

References

- [1] M Jorge Cardoso, Wenqi Li, Richard Brown, Nic Ma, Eric Kerfoot, Yiheng Wang, Benjamin Murrey, Andriy Myronenko, Can Zhao, Dong Yang, et al. Monai: An open-source framework for deep learning in healthcare. arXiv preprint arXiv:2211.02701, 2022. 8
- [2] 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. 1
- [3] Zhaolin Chen, Kamlesh Pawar, Mevan Ekanayake, Cameron Pain, Shenjun Zhong, and Gary F Egan. Deep learning for image enhancement and correction in magnetic resonance imaging—state-of-the-art and challenges. *Journal of Digital Imaging*, 36(1):204–230, 2023. 1
- [4] Aram Danielyan, Vladimir Katkovnik, and Karen Egiazarian. Bm3d frames and variational image deblurring. *IEEE Transactions on image processing*, 21(4):1715–1728, 2011.
- [5] Emilien Dupont, Adam Goliński, Milad Alizadeh, Yee Whye Teh, and Arnaud Doucet. Coin: Compression with implicit neural representations. *arXiv preprint arXiv:2103.03123*, 2021. 1
- [6] Alan C Evans, Andrew L Janke, D Louis Collins, and Sylvain Baillet. Brain templates and atlases. *Neuroimage*, 62(2):911–922, 2012. 1, 4
- [7] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Multiplicative filter networks. In *International Conference on Learning Representations*, 2020. 4, 6, 7, 8
- [8] Albert Gubern-Merida and Robert Marti. Atlas based segmentation of the prostate in mr images. In MICCAI: Segmentation Challenge Workshop, 2009. 8
- [9] Batuhan Gundogdu, Milica Medved, Aritrick Chatterjee, Roger Engelmann, Avery Rosado, Grace Lee, Nisa C Oren, Aytekin Oto, and Gregory S Karczmar. Self-supervised multicontrast super-resolution for diffusion-weighted prostate mri. *Magnetic Resonance in Medicine*, 2024. 1
- [10] Mark Jenkinson, Mickael Pechaud, Stephen Smith, et al. Bet2: Mr-based estimation of brain, skull and scalp surfaces. In *Eleventh annual meeting of the organization for human brain mapping*, volume 17, page 167. Toronto., 2005. 4
- [11] Amirhossein Kazerouni, Reza Azad, Alireza Hosseini, Dorit Merhof, and Ulas Bagci. Incode: Implicit neural conditioning with prior knowledge embeddings. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pages 1298–1307, 2024. 2, 5
- [12] Bennett Landman, Zhoubing Xu, J Igelsias, Martin Styner, Thomas Langerak, and Arno Klein. Miccai multi-atlas labeling beyond the cranial vault–workshop and challenge. In Proc. MICCAI Multi-Atlas Labeling Beyond Cranial Vault–Workshop Challenge, volume 5, page 12, 2015. 4
- [13] Jun Ma, Yuting He, Feifei Li, Lin Han, Chenyu You, and Bo Wang. Segment anything in medical images. *Nature Communications*, 15(1):654, 2024. 8
- [14] Jürgen K Mai, Milan Majtanik, and George Paxinos. *Atlas of the human brain*. Academic Press, 2015. **1**, 4

- [15] Daniele Marin, Daniel T Boll, Achille Mileto, and Rendon C Nelson. State of the art: dual-energy ct of the abdomen. *Radiology*, 271(2):327–342, 2014. 4
- [16] Julian McGinnis, Suprosanna Shit, Hongwei Bran Li, Vasiliki Sideri-Lampretsa, Robert Graf, Maik Dannecker, Jiazhen Pan, Nil Stolt-Ansó, Mark Mühlau, Jan S Kirschke, et al. Single-subject multi-contrast mri super-resolution via implicit neural representations. In *International Conference* on Medical Image Computing and Computer-Assisted Intervention, pages 173–183. Springer, 2023. 1
- [17] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014. 4
- [18] 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.
- [19] Amirali Molaei, Amirhossein Aminimehr, Armin Tavakoli, Amirhossein Kazerouni, Bobby Azad, Reza Azad, and Dorit Merhof. Implicit neural representation in medical imaging: A comparative survey. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2381– 2391, 2023. 1
- [20] Matthew J Muckley, Bruno Riemenschneider, Alireza Radmanesh, Sunwoo Kim, Geunu Jeong, Jingyu Ko, Yohan Jun, Hyungseob Shin, Dosik Hwang, Mahmoud Mostapha, et al. Results of the 2020 fastmri challenge for machine learning mr image reconstruction. *IEEE transactions on medical imaging*, 40(9):2306–2317, 2021. 4
- [21] Mariusz Oszust, Adam Piórkowski, and Rafał Obuchowicz. No-reference image quality assessment of magnetic resonance images with high-boost filtering and local features. *Magnetic Resonance in Medicine*, 84(3):1648–1660, 2020.
- [22] Defu Qiu, Shengxiang Zhang, Ying Liu, Jianqing Zhu, and Lixin Zheng. Super-resolution reconstruction of knee magnetic resonance imaging based on deep learning. *Computer methods and programs in biomedicine*, 187:105059, 2020. 1
- [23] Sameera Ramasinghe and Simon Lucey. Beyond periodicity: Towards a unifying framework for activations in coordinatemlps. In *European Conference on Computer Vision*, pages 142–158. Springer, 2022. 4, 6, 7, 8
- [24] Vishwanath Saragadam, Daniel LeJeune, Jasper Tan, Guha Balakrishnan, Ashok Veeraraghavan, and Richard G Baraniuk. Wire: Wavelet implicit neural representations. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 18507–18516, 2023. 2, 4, 7, 8
- [25] Liyue Shen, John Pauly, and Lei Xing. Nerp: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. 1
- [26] Kexuan Shi, Xingyu Zhou, and Shuhang Gu. Improved implicit neural representation with fourier reparameterized

training. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 25985– 25994, 2024. 4, 6, 7, 8

- [27] 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. 1, 2, 4, 6, 7, 8
- [28] Stephen M Smith. Fast robust automated brain extraction. *Human brain mapping*, 17(3):143–155, 2002. 4
- [29] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. Advances in Neural Information Processing Systems, 33:7537–7547, 2020. 4, 6, 7, 8
- [30] Qing Wu, Yuwei Li, Yawen Sun, Yan Zhou, Hongjiang Wei, Jingyi Yu, and Yuyao Zhang. An arbitrary scale superresolution approach for 3d mr images via implicit neural representation. *IEEE Journal of Biomedical and Health Informatics*, 27(2):1004–1015, 2022. 1
- [31] Qing Wu, Yuwei Li, Lan Xu, Ruiming Feng, Hongjiang Wei, Qing Yang, Boliang Yu, Xiaozhao Liu, Jingyi Yu, and Yuyao Zhang. Irem: High-resolution magnetic resonance image reconstruction via implicit neural representation. 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 65–74. Springer, 2021. 1*
- [32] Shaowen Xie, Hao Zhu, Zhen Liu, Qi Zhang, You Zhou, Xun Cao, and Zhan Ma. Diner: Disorder-invariant implicit neural representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6143–6152, 2023. 2
- [33] Siqi Ye, Liyue Shen, Md Tauhidul Islam, and Lei Xing. Super-resolution biomedical imaging via reference-free statistical implicit neural representation. *Physics in Medicine & Biology*, 68(20):205020, 2023. 1
- [34] Haonan Zhang, Yuhan Zhang, Qing Wu, Jiangjie Wu, Zhiming Zhen, Feng Shi, Jianmin Yuan, Hongjiang Wei, Chen Liu, and Yuyao Zhang. Self-supervised arbitrary scale superresolution framework for anisotropic mri. arXiv preprint arXiv:2305.01360, 2023. 1
- [35] Can Zhao, Blake E Dewey, Dzung L Pham, Peter A Calabresi, Daniel S Reich, and Jerry L Prince. Smore: a selfsupervised anti-aliasing and super-resolution algorithm for mri using deep learning. *IEEE transactions on medical imaging*, 40(3):805–817, 2020. 1
- [36] Can Zhao, Muhan Shao, Aaron Carass, Hao Li, Blake E Dewey, Lotta M Ellingsen, Jonghye Woo, Michael A Guttman, Ari M Blitz, Maureen Stone, et al. Applications of a deep learning method for anti-aliasing and super-resolution in mri. *Magnetic resonance imaging*, 64:132–141, 2019. 1