Cubic Structure-Aware Neural Radiance Fields for Zero-Shot Super-Resolution of Susceptibility-Weighted Imaging

Hongwei Bi[1], Rong Zhang[1]✉, Lijun Guo[1], Fengchun Liu[1], Guomin Xie[2], Zhongyue Lv[2], and Xiping Wu[2]

[1] Faculty of Electrical Engineering and Computer Science, Ningbo University, Ningbo, 315211, China
zhangrong@nbu.edu.cn
[2] Department of Neurology, Ningbo Medical Center Lihuili Hospital, Ningbo, 315040, China

Abstract. Susceptibility-Weighted Imaging (SWI) is essential for detecting Cerebral Microbleeds (CMBs), a key biomarker in neurodegenerative diseases. However, clinical SWI scans are often acquired with thick slices within clinically acceptable scan durations, resulting in severe anisotropy that degrades through-plane resolution and leads to missed small lesions. To address this, this paper proposes Cubic Structure-Aware Neural Radiance Fields (CSANeRF), a zero-shot Super-Resolution (SR) framework that integrates NeRF with Transformer architectures for high-fidelity SWI SR reconstruction. To preserve the critical biomarkers in SWI, we rethink the construction of implicit radiance fields by incorporating explicit structural priors from local 3D neighborhoods. The core innovation lies in a cube-based spatial modeling strategy. CSANeRF is validated on two independent, expert-annotated SWI datasets and consistently outperforms state-of-the-art 3D SR approaches in both visual quality and quantitative metrics. Critically, when employed as a preprocessing step for CMB segmentation, CSANeRF reconstructions substantially improve segmentation accuracy, achieving performance closest to the HR ground truth and surpassing all baseline methods.This finding underscores its potential to enhance diagnostic reliability in neurodegenerative disease evaluation.

Keywords: Susceptibility-weighted imaging · Neural Radiance Field · Super-Resolution · Zero-Shot.

1 Introduction

As an emerging neuroimaging technique, Susceptibility Weighted Imaging (SWI) has demonstrated substantial clinical value in detecting CMBs, evaluating cerebrovascular malformations, diagnosing neurodegenerative diseases, and assessing traumatic brain injury [11, 19, 24]. For CMB detection, SWI can identify a greater number of smaller lesions than conventional T2*-weighted sequences [21],

thereby providing critical imaging evidence for stroke risk stratification, anticoagulation therapy decision-making, and differential diagnosis of dementia [9, 6].

However, medical image acquisition inherently involves a trade-off among scan time, spatial resolution, and signal-to-noise ratio [22, 17]. To obtain sufficient diagnostic information within clinically acceptable scan durations, routine imaging protocols typically employ large slice thickness in the through-plane direction, resulting in significantly anisotropic acquisitions characterized by high in-plane resolution but markedly reduced through-plane resolution [29]. To address this issue, SR reconstruction strategies have been explored across multiple imaging modalities [14, 20]. In recent years, with the advancement of deep learning, representative approaches mainly include supervised Deep Neural Networks (DNNs) [4, 8] and continuous signal modeling methods based on Implicit Neural Representations (INRs) [27].

CMBs in SWI serve as a key biomarker for neurodegenerative diseases. A critical challenge in their clinical identification lies in distinguishing true lesions from mimicking artifacts, which depends on the spatial continuity of microvascular structures [9]. However, the through-plane resolution deficiency in routine SWI not only obscures small CMBs but also fragments the 3D continuity of these microvascular networks, hindering reliable lesion tracking and increasing diagnostic uncertainty. Most existing general-purpose SR frameworks [4, 15] fail to account for the unique magnetic susceptibility-induced contrast mechanism of SWI. They may blur or distort fine structural details, exacerbating discontinuities and introducing spurious patterns. Moreover, supervised SR methods [4, 8, 27] require paired Low-Resolution(LR)-High-Resolution(HR) training data. Yet, HR SWI entails prolonged acquisition times that are poorly tolerated by elderly or cognitively impaired patients, rendering large-scale collection of such paired data clinically infeasible. Compounding this issue, domain shifts across scanner protocols may cause model miscalibration and misdiagnosis [1, 10].

So far, dedicated SR methods specifically designed for SWI remain scarce. To address the above-mentioned limitations, this paper proposes Cubic Structure-Aware Neural Radiance Field (CSANeRF), a zero-shot SR reconstruction framework tailored for SWI. Built upon the Neural Radiance Field (NeRF) [18] architecture, the proposed method utilizes solely the patient's own LR SWI volumetric data as the supervisory signal, requiring neither HR training data nor external priors. This self-supervised, subject-specific reconstruction paradigm removes the reliance on paired HR ground-truth images, bypassing a major barrier to clinical deployment.

To the best of our knowledge, this is the first zero-shot SR approach for SWI to enhance CMB detection. To preserve the critical biomarkers in SWI, we rethink the construction of implicit radiance fields by incorporating structural priors from local 3D neighborhoods. The main contributions are three folds:

1. We propose a cube-based spatial modeling strategy that constrains NeRF's point queries to geometrically coherent 3D neighborhoods.

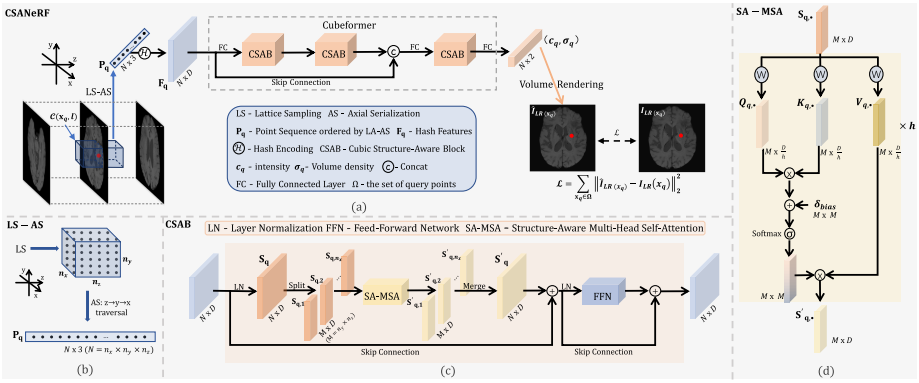

Fig. 1: **(a)** Overall framework of the proposed CSANeRF. **(b)** Lattice Sampling and Axial Serialization (LS-AS). **(c)** Cubic Structure-Aware Block (CSAB). **(d)** Structure-Aware Multi-Head Self-Attention (SA-MSA).

2. We introduce Cubeformer, a structure-aware Transformer module that replaces the standard Multi-Layer Perceptron (MLP) based radiance field modeling in NeRF.
3. Extensive evaluation on CMB-annotated SWI datasets demonstrates reconstruction quality that preserves critical CMB biomarkers better than state-of-the-art SR methods.

Related Work. Recent studies have explored the application of NeRF to medical image SR [23, 5]. NeRF [18] optimizes an implicit neural representation on a single scene, enabling the reconstruction of continuous 3D structures from sparse observations and offering a promising avenue for eliminating dependence on external training data. However, the original NeRF employed a point-wise inferring mechanism, where each spatial coordinate is independently processed through an MLP, lacking explicit modeling of neighborhood context or structural priors. CuNeRF [5] adapted NeRF to 3D medical volumetric data by replacing the original ray-based sampling strategy with cube-based sampling. Despite its zero-shot advantage, CuNeRF retains the point-independent MLP architecture, neglecting the pervasive local structural correlations inherent in medical images.

Transformer architectures [25] have demonstrated significant advantages in medical image analysis owing to its capacity for modeling long-range dependencies [12, 13]. For 3D medical volumetric data, SuperFormer [8] extended the Swin Transformer [16] to 3D settings, demonstrating the potential of voxel-level Transformers in MRI SR tasks. SAX-NeRF [3] leveraged Transformer backbones to capture structural dependencies for sparse-view X-ray to CT 3D reconstruction, outperforming conventional NeRF methods.

2 Method

Fig. 1(a) illustrates the overall framework of the proposed CSANeRF model. Given a LR SWI volume $I_{\mathrm{LR}} \in \mathbb{R}^{H\times W\times L}$ with spatial coordinates normalized to $[-1,1]^3$, the pipeline comprises three main modules: Lattice Sampling and Axial Serialization (LS-AS) to constrain NeRF's point queries to geometrically coherent 3D neighborhoods, Cubeformer for Transformer-based structure-aware radiance field learning, and Volume Rendering (VR) for intensity prediction.

During training, the model learns a continuous coordinate-to-intensity mapping for volume synthesis. During inference, given a SR factor s, we first construct a dense HR target grid based on the original LR volume $I_{\mathrm{LR}} \in \mathbb{R}^{H\times W\times L}$. The coordinates of this HR grid are then normalized to $[-1,1]^3$ and queried from the trained CSANeRF model. Finally, the predicted intensities are assembled to reconstruct the super-resolved volume $\hat{I}_{\mathrm{HR}} \in \mathbb{R}^{H\times W\times (s\cdot L)}$.

LS-AS. Unlike the ray-based sampling in NeRF or the cube-based sampling in CuNeRF, we propose LS-AS for structure-preserving spatial modeling (Fig. 1(b)).

Let $\Omega = \{\mathbf{x}_q\}_{q=1}^{Q}$ denote the set of Q query points sampled from a selected slice. For each $\mathbf{x}_q = (x_q, y_q, z_q) \in \Omega$, we construct a local cubic neighborhood $\mathcal{C}(\mathbf{x}_q, l)$ centered at $\mathbf{x}_q$ with side length l. This cube is uniformly discretized along the spatial axes into n_x, n_y, and n_z grid nodes, forming a regular 3D lattice comprising a total of $N = n_x \times n_y \times n_z$ sampling points. It is worth noting that n_x, n_y, and n_z need not be identical, allowing for anisotropic sampling densities along each axis.

To adapt to Transformer-based attention mechanisms, we introduce AS to convert the 3D lattice into an ordered sequence while preserving spatial adjacency. The N sampling points are traversed in a lexicographical order with respect to the indices (i, j, k), where i, j, and k correspond to the grid positions along the x-, y-, and z-axes. This z-fastest traversal strategy explicitly models inter-slice structural continuity, which is particularly beneficial for SWI imaging where vessels exhibit strong longitudinal correlations across adjacent slices.

Consequently, for the N sampling points in $\mathcal{C}(\mathbf{x}_q, l)$, an ordered coordinate sequence $\mathbf{P}_q = \{\mathbf{p}_{q,n}\}_{n=0}^{N-1}$ is generated, where $n = i \cdot (n_y \cdot n_z) + j \cdot n_z + k$, $i \in \{0, \ldots, n_x - 1\}$, $j \in \{0, \ldots, n_y - 1\}$, and $k \in \{0, \ldots, n_z - 1\}$, and $\mathbf{p}_{q,n} \in \mathbb{R}^3$.

Cubeformer. Inspired by SAX-NeRF [3], we propose a structure-aware Transformer, termed Cubeformer. The Cubeformer consists of three stacked Cubic Structure-Aware Blocks (CSABs) (Fig. 1(a)). The coordinate sequence $\mathbf{P}_q$ is first hash-encoded [3] to generate initial features $\mathbf{F}_q = \mathcal{H}(\mathbf{P}_q) \in \mathbb{R}^{N\times D}$, where D denotes the feature dimension. The features $\mathbf{F}_q$ are then fed into Cubeformer, whose outputs comprise intensity values $\mathbf{c}_q \in \mathbb{R}^N$ and volume densities $\boldsymbol{\sigma}_q \in \mathbb{R}^N$.

CSAB is the core unit of Cubeformer. As shown in Fig. 1(c), it adopts a residual architecture with a Structure-Aware Multi-head Self-Attention (SA-MSA) module and a Feed-Forward Network (FFN). The SA-MSA captures structural dependencies among sampling points within each cubic neighborhood (Fig. 1(d)).

Within each CSAB, the input is first layer-normalized into $\mathbf{S}_q \in \mathbb{R}^{N\times D}$. To improve computational efficiency, $\mathbf{S}_q$ is partitioned into n_x non-overlapping local segments $\{\mathbf{S}_{q,m}\}_{m=1}^{n_x}$, where each segment $\mathbf{S}_{q,m} \in \mathbb{R}^{M\times D}$, $M = n_y \cdot n_z$. Each segment can be independently processed by SA-MSA to model local structural dependencies. The outputs are merged to reconstruct the full sequence as $\mathbf{S'}_q = \text{Merge}\big(\{\text{SA-MSA}(\mathbf{S}_{q,m})\}_{m=1}^{n_x}\big)$.

To encode geometric relationships between sampling points, a relative position bias is incorporated into SA-MSA (Fig. 1(d)). A lightweight MLP maps the 3D coordinate difference between any two points $\mathbf{p}_q^{(u)}$ and $\mathbf{p}_q^{(v)}$ to an h-dimensional bias vector $\boldsymbol{\delta}_{\text{bias}} = \alpha \cdot \tanh\big(\text{MLP}(\mathbf{p}_q^{(u)} - \mathbf{p}_q^{(v)})\big)$, where h denotes the number of attention heads, α is a learnable scaling factor, and u and v are the indices of two points within a segment, respectively. The bias enables explicit modeling of relative spatial configurations.

VR. Unlike NeRF and CuNeRF, which employ hierarchical coarse-to-fine sampling, CSANeRF utilizes a single-stage sampling strategy for VR. Based on the predictions of Cubeformer, the final voxel intensity at $\mathbf{x}_q$ is approximated via discrete volume rendering integration [5]:

$$\hat{I}\left(\mathbf{x}_q, l\right) = 4\pi \sum_{i=1}^{N} \frac{r_i^2 \left(1 - \exp\left(-\sigma_{q,i}\left(r_{i+1} - r_i\right)\right)\right)}{\exp\left(4\pi \sum_{j=1}^{i} r_j^2 \sigma_{q,j}\left(r_{j+1} - r_j\right)\right)} c_{q,i} \tag{1}$$

where r_i denotes the radial distance from the i-th sampling point to the cube center $\mathbf{x}_q$, and $r_1 \leq r_2 \leq \cdots \leq r_N$.

3 Experiments and Results

Dataset. To evaluate the SR capacity of CSANeRF in preserving critical biomarkers, experiments were conducted on two SWI datasets: (1) HKCU-SWI-CMB [7], a public SWI dataset ($512 \times 512 \times 150$, voxel spacing $0.45 \times 0.45 \times 1.0$ mm^3). (2) Private-SWI-CMB, an in-house clinical dataset collected from the neurology department of a tertiary hospital with ethical approval and informed consent. It comprises 34 subjects ($384 \times 384 \times 72$, voxel spacing $0.6 \times 0.6 \times 1.5$ mm^3). All CMB annotations were validated by experienced neuroradiologists. Additionally, to further validate the SR reconstruction capability of our method on other modalities, we also evaluated it on MSD-T1 [2], a T1-weighted brain tumor dataset from the Medical Segmentation Decathlon ($240 \times 240 \times 155$, voxel spacing $1 \times 1 \times 1$ mm^3).

Comparison Methods. We compared CSANeRF with: (1) Interpolation methods (cubic, trilinear); (2) NeRF-based methods (NeRF [18], CuNeRF [5]). All methods followed the original network configurations reported in the literature.

Table 1: Quantitative comparison on two SWI datasets and one T1-weighted dataset. Best results are bold, second best underlined.

	SR Scale	2×		3×		4×	
	Method	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑
HKCU-SWI-CMB	Cubic	30.32	0.9009	28.29	0.8500	26.95	0.8132
	Trilinear	30.60	0.9043	28.75	0.8640	27.58	0.8363
	NeRF	24.32	0.8623	23.54	0.8235	22.67	0.8014
	CuNeRF	30.35	0.9401	28.04	0.9092	26.53	0.8735
	CSANeRF(Ours)	**31.63**	**0.9571**	**29.28**	**0.9304**	**27.75**	**0.9026**
Private-SWI-CMB	Cubic	35.12	0.9261	33.12	0.8837	31.84	0.8538
	Trilinear	35.41	0.9291	33.61	0.8959	32.48	0.8724
	NeRF	31.70	0.8896	30.20	0.8576	28.81	0.8095
	CuNeRF	36.04	0.9412	34.25	0.9253	33.15	0.9088
	CSANeRF(Ours)	**37.72**	**0.9587**	**35.41**	**0.9414**	**33.92**	**0.9289**
MSD-T1	Cubic	35.27	0.9652	33.61	0.9442	32.13	0.9231
	Trilinear	35.53	0.9703	33.81	0.9530	32.39	0.9380
	NeRF	30.25	0.8651	27.06	0.8424	26.23	0.8257
	CuNeRF	41.02	0.9876	39.46	0.9831	37.67	0.9768
	CSANeRF(Ours)	**42.83**	**0.9915**	**40.91**	**0.9862**	**39.02**	**0.9824**

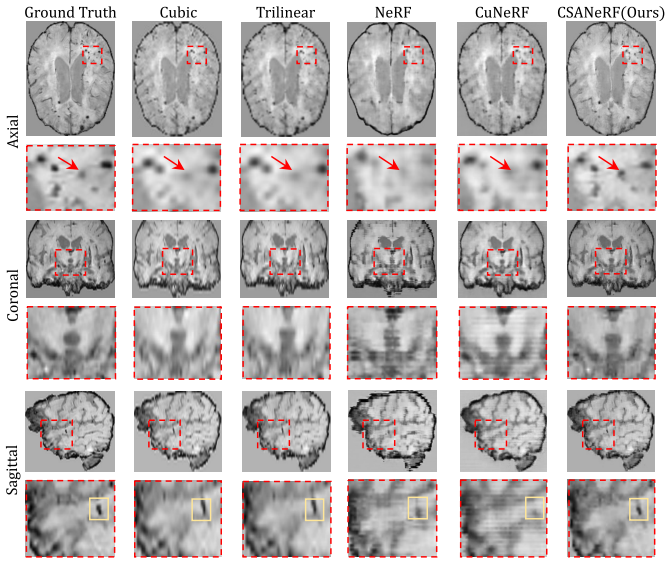

Fig. 2: Qualitative comparison of 2× SR reconstructions on the Private-SWI-CMB dataset. Red arrows highlight regions where CSANeRF preserves fine structural details better than competing methods; yellow boxes delineate CMB lesions where CSANeRF better maintains lesion morphology.

Training Protocol. LR training data were generated by z-axis downsampling while preserving in-plane resolution. All learning-based methods (NeRF, CuNeRF and CASNeRF) were trained solely on LR data without ground truth HR access during training. Original HR volumes were used exclusively for evaluation as reference standards.

Table 2: Ablation study on lattice sampling numbers (N) and segmented attention.

N (n_x, n_y, n_z)	100 (5,4,5)		120 (6,4,5)		125 (5,5,5)		150 (6,5,5)		180 (6,6,5)	
Segmented Attn.	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o
PSNR (↑)	37.09	37.51	37.72	**37.97**	37.15	37.58	37.49	37.60	37.64	37.89
SSIM (↑)	0.9541	0.9569	0.9587	**0.9598**	0.9558	0.9574	0.9571	0.9577	0.9578	0.9593

Table 3: Ablation study on different axis traversal strategies.

Axis orders	$x \to y \to z$	$z \to x \to y$	$z \to y \to x$
PSNR (↑)	36.87	37.46	**37.72**
SSIM (↑)	0.9483	0.9563	**0.9587**

Table 4: Ablation study on relative position bias ($\boldsymbol{\delta}_{\text{bias}}$).

Method	w/o $\boldsymbol{\delta}_{\text{bias}}$	w/ $\boldsymbol{\delta}_{\text{bias}}$
PSNR↑	37.16	**37.72**
SSIM↑	0.9526	**0.9587**

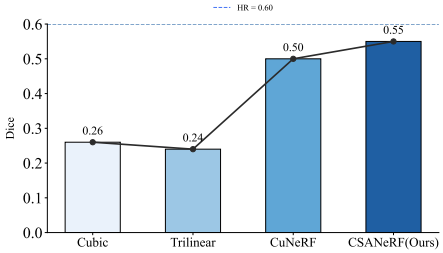

Fig. 3: CMB segmentation on SR reconstructions from different methods.

Evaluation Metrics. SR outputs were compared against the corresponding HR ground truth for fair comparison. We quantitatively evaluated reconstruction quality using Peak Signal-to-Noise Ratio (PSNR) for pixel-level accuracy, and Structural Similarity Index Measure (SSIM) [26] for perceptual structural fidelity. Higher values indicate better performance for both metrics. All metrics were computed within the brain mask to exclude background regions.

Implementation Details. This study was implemented based on the PyTorch framework. All experiments were conducted on a server equipped with 24 GB NVIDIA RTX 4090 GPUs. The Adam optimizer was employed with a weight decay set to 1×10^{-6}. The maximum training iterations were set to 25000, with a batch size of 4. The learning rate decayed exponentially from 2×10^{-3} to 2×10^{-5}. The number of query points Q was set to 1024.

Quantitative Analysis. Table 1 presents the quantitative performance of five SR methods on the HKCU-SWI-CMB, Private-SWI-CMB and MSD-T1 datasets at SR factors of 2×, 3×, and 4×. Across all scales, our method consistently achieves the highest PSNR and SSIM. It is worth noting that on the Private-SWI-CMB dataset, the PSNR scores of CSANeRF at 2×, 3×, and 4× scales are significantly higher than the second-best method (CuNeRF) by 1.68, 1.16, and 0.77 dB, respectively. Similarly, remarkable improvements are observed on the HKCU-SWI-CMB and MSD-T1 datasets, indicating strong robustness across imaging devices, patient cohorts, and MRI modalities. Compared with other methods, NeRF performs relatively poorly, possibly due to its ray-based sampling strategy, which provides limited spatial coverage.

Qualitative Analysis. Fig. 2 presents visual comparisons of axial, coronal, and sagittal slices reconstructed at a 2× scale on the Private-SWI-CMB dataset. In the axial view, the enlarged regions indicated by red arrows show that CSANeRF recovers sharper fine-grained structural details than interpolation-based and NeRF-based methods, which suffer from over-smoothing and blurring artifacts. In the coronal view, CSANeRF produces more coherent anatomical contours with improved inter-slice structural continuity along the through-plane direction. In the sagittal view, the yellow boxes highlight CMB lesion regions where CSANeRF better preserves the compact morphology and susceptibility contrast of microbleeds, demonstrating its clinical relevance for CMB detection.

Ablation Study. We conducted a comprehensive ablation study on the Private-SWI-CMB dataset to validate the effectiveness of key design choices in our proposed method. **(1) Lattice sampling and segmented attention.** Table 2 shows that reconstruction quality peaks at $N = 120$, beyond which redundancy leads to performance degradation. Regarding the attention mechanism in CSAB, global attention (without segmentation) consistently yields slightly higher accuracy than segmented attention across all N. However, it incurs prohibitive costs (e.g., 2.5 h vs. 1.7 h training time for $N = 120$ on a $384 \times 384 \times 72$ volume). We adopt a trade-off: $N = 120$ with segmented attention as our default configuration. For practical applications, configurations can be flexibly adjusted based on specific resource. **(2) Axis traversal strategies for AS.** Table 3 examines the impact of different axis traversal orders on model performance. The $z \to y \to x$ order achieves the best performance, validating that prioritizing the z-axis enables the model to effectively learn inter-slice structural continuity and anatomical consistency. This is particularly critical for SWI imaging with anisotropic resolution. **(3) Relative position bias.** Table 4 reports that introducing $\boldsymbol{\delta}_{\text{bias}}$ in SA-MSA improves PSNR by 0.56 dB and SSIM by 0.0061, indicating its effectiveness in enhancing fine-grained detail reconstruction.

Applicability to Downstream Tasks. To evaluate the practical utility of our method, we applied its SR outputs to CMB lesion segmentation. Using E^{2}U-Net [28] trained on original HR images (80% of the Private-SWI-CMB dataset), we evaluated performance on the remaining 20% testing set. For evaluation, LR images in the testing set were generated via 2× downsampling and reconstructed using various SR methods. As shown in Fig 3, CSANeRF achieves a Dice score of 0.55, closest to the HR ground truth (0.60), and significantly outperforms both interpolation-based and deep SR baselines. This confirms that our approach better preserves lesion morphology, offering more reliable support for automated CMB detection.

4 Discussion and Conclusion

In this work, we presented CSANeRF, a novel structure-aware NeRF-Transformer framework for SR reconstruction of SWI images. A key advantage of our ap-

proach is its self-supervised learning paradigm, which eliminates the dependency on paired HR ground truth data during training. This feature offers substantial flexibility for diverse clinical scenarios where HR references are constrained.

Overall, CSANeRF demonstrates outstanding ability in preserving anatomical structures. Extensive quantitative and qualitative evaluations confirm that our method significantly outperforms existing state-of-the-art approaches. We further applied CSANeRF to a downstream CMB detection task to validate its capacity of preserving critical biomarkers. We anticipate that CSANeRF will facilitate the integration of deep learning into clinical workflows, particularly by enabling faster scan protocols and enhancing the visibility of subtle biomarkers for CMB detection and neurological disease assessment.

Acknowledgments. This research work was partly supported by the Ningbo Key Research and Development Program (Grant No. 2023Z196), the Ningbo Municipal Public Welfare Technology Research Project (No. 2023S035), the Guangxi Key Technologies R & D Program (Grant No. FN2504240023), the Neurology Department of the National Key Clinical Speciality Construction Project, and the Project of Ningbo Leading Medical & Health Discipline (Grant No. 2022-F05).

Disclosure of Interests. The authors have no competing interests to declare.

References

1. AlBadawy EA, Saha A, M.M.: Deep learning for segmentation of brain tumors: Impact of cross-institutional training and testing. Medical Physics **45**(3), 1150–1158 (2018). https://doi.org/10.1002/mp.12752
2. Antonelli, M., Reinke, A., Bakas, S., et al.: The medical segmentation decathlon. Nature Communications **13**, 4128 (2022)
3. Cai, Y., Wang, J., Yuille, A., Zhou, Z., Wang, A.: Structure-aware sparse-view x-ray 3d reconstruction. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11174–11183 (2024)
4. Chen, Y., Xie, Y., Zhou, Z., Shi, F., et al.: Brain mri super resolution using 3d deep densely connected neural networks. 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018) pp. 739–742 (2018), https://api.semanticscholar.org/CorpusID:2103644
5. Chen, Z., Yang, L., Lai, J., Xie, X.: Cunerf: Cube-based neural radiance field for zero-shot medical image arbitrary-scale super resolution. In: ICCV (2023)
6. Cordonnier, C., van der Flier, W.: Brain microbleeds and alzheimer's disease: Innocent observation or key player? Brain **134**(2), 335–344 (2011). https://doi.org/10.1093/brain/awq321
7. Dou, Q., Chen, H., Yu, L., et al.: Automatic detection of cerebral microbleeds from mr images via 3d convolutional neural networks. IEEE Transactions on Medical Imaging **35**(5), 1182–1195 (2016)
8. Forigua, C., Escobar, M., Arbelaez, P.: Superformer: Volumetric transformer architectures for mri super-resolution. In: Simulation and Synthesis in Medical Imaging (2022)
9. Greenberg, S., Vernooij, M., Cordonnier, C., et al.: Cerebral microbleeds: A guide to detection and interpretation. The Lancet Neurology **8**(2), 165–174 (2009). https://doi.org/10.1016/S1474-4422(09)70013-4

10. Guan, H., Liu, M.: Domain adaptation for medical image analysis: A survey. IEEE Transactions on Biomedical Engineering **69**(3), 1173–1185 (2022). https://doi.org/10.1109/TBME.2021.3117407
11. Haacke, E., Xu, Y., Cheng, Y., Reichenbach, J.: Susceptibility weighted imaging (swi). Magnetic Resonance in Medicine **52**(3), 612–618 (2004). https://doi.org/10.1002/mrm.20198
12. Hatamizadeh, A., Tang, Y., Nath, V., et al.: Unetr: Transformers for 3d medical image segmentation. In: WACV (2022)
13. Hu, D., Zhang, Y., Zhu, J., Liu, Q., Chen, Y.: Trans-net: Transformer-enhanced residual-error alternative suppression network for mri reconstruction. IEEE Transactions on Instrumentation and Measurement **71**, 1–13 (2022). https://doi.org/10.1109/TIM.2022.3205684
14. Lehmann, T.M., Gonner, C., Spitzer, K.: Survey: interpolation methods in medical image processing. IEEE Transactions on Medical Imaging **18**(11), 1049–1075 (1999). https://doi.org/10.1109/42.816070
15. 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. pp. 1833–1844 (2021)
16. Liu, Z., Lin, Y., Cao, Y., et al.: Swin transformer: Hierarchical vision transformer using shifted windows. In: ICCV (2021)
17. Lustig, M., Donoho, D., Pauly, J.: Sparse mri: The application of compressed sensing for rapid mr imaging. Magnetic Resonance in Medicine **58**(6), 1182–1195 (2007). https://doi.org/10.1002/mrm.21391
18. Mildenhall, B., Srinivasan, P., Tancik, M., et al.: Nerf: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM **65**(1), 99–106 (2021)
19. Mittal, S., Wu, Z., Neelavalli, J., Haacke, E.: Susceptibility-weighted imaging: Technical aspects and clinical applications, part 2. American Journal of Neuroradiology **30**(2), 232–252 (2009). https://doi.org/10.3174/ajnr.A1461
20. Muhammad, A., Aramvith, S., Duangchaemkarn, K., Sun, M.T.: Brain mri image super-resolution reconstruction: A systematic review. IEEE Access **12**, 156347–156362 (2024). https://doi.org/10.1109/ACCESS.2024.3478829
21. Nandigam, R., Viswanathan, A., Delgado, P., et al.: Mr imaging detection of cerebral microbleeds: Effect of susceptibility-weighted imaging, section thickness, and field strength. American Journal of Neuroradiology **30**(2), 338–343 (2009). https://doi.org/10.3174/ajnr.A1355
22. Plenge, E., Poot, D., Bernsen, M., et al.: Super-resolution methods in mri: Can they improve the trade-off between resolution, signal-to-noise ratio, and acquisition time? Magnetic Resonance in Medicine **68**(6), 1983–1993 (2012). https://doi.org/10.1002/mrm.24187
23. Shen, L., Pauly, J., Xing, L.: Nerp: Implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. IEEE Transactions on Neural Networks and Learning Systems **35**(1), 770–782 (2022)
24. Tong, K., Ashwal, S., Holshouser, B., et al.: Susceptibility-weighted mr imaging: A review of clinical applications in children. American Journal of Neuroradiology **29**(1), 9–17 (2008). https://doi.org/10.3174/ajnr.A0786
25. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., et al.: Attention is all you need. In: NeurIPS (2017)
26. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE Transactions on Image Processing **13**(4), 600–612 (2004). https://doi.org/10.1109/TIP.2003.819861

27. Wang X, Wang S, X.H., et al.: Spatial attention-based implicit neural representation for arbitrary reduction of mri slice spacing. Medical Image Analysis **94**, 103158 (2024). https://doi.org/10.1016/j.media.2024.103158
28. Xia, T., Zhang, R., Chen, Z., et al.: Progressive learning based knowledge distillation for low resolution cerebral microbleed segmentation. In: ICASSP (2024)
29. Y. Jia, A. Gholipour, Z.H., Warfield, S.K.: A new sparse representation framework for reconstruction of an isotropic high spatial resolution mr volume from orthogonal anisotropic resolution scans. IEEE Transactions on Medical Imaging **36**(5), 1182–1193 (2017). https://doi.org/10.1109/TMI.2017.2656907