Faster, Self-Supervised Super-Resolution for Anisotropic Multi-View MRI Using a Sparse Coordinate Loss

Maja Schlereth[1], Moritz Schillinger[1], and Katharina Breininger[1,2]

[1] Department Artificial Intelligence in Biomedical Engineering, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany
`maja.schlereth@fau.de`
[2] Center for AI and Data Science (CAIDAS), Julius-Maximilians-Universität Würzburg, Würzburg, Germany

Abstract. Acquiring images in high resolution is often a challenging task. Especially in the medical sector, image quality has to be balanced with acquisition time and patient comfort. To strike a compromise between scan time and quality for Magnetic Resonance (MR) imaging, two anisotropic scans with different low-resolution (LR) orientations can be acquired. Typically, LR scans are analyzed individually by radiologists, which is time consuming and can lead to inaccurate interpretation. To tackle this, we propose *tripleSR*, a novel approach for fusing two orthogonal anisotropic LR MR images, to reconstruct anatomical details in a unified representation. Our multi-view neural network is trained in a self-supervised manner, without requiring corresponding high-resolution (HR) data. To optimize the model, we introduce a sparse coordinate-based loss, enabling the integration of LR images with arbitrary scaling. We evaluate our method on MR images from two independent cohorts. Our results demonstrate comparable or even improved super-resolution (SR) performance compared to state-of-the-art (SOTA) self-supervised SR methods for different upsampling scales. By combining a patient-agnostic offline and a patient-specific online phase, we achieve a substantial speed-up of up to ten times for patient-specific reconstruction while achieving similar or better SR quality. Code is available at `https://github.com/MajaSchle/tripleSR`.

Keywords: MR imaging · Super-Resolution · Multi-View.

1   Introduction

Magnetic Resonance (MR) imaging allows the 3-D assessment of bone and soft tissue anomalies without ionizing radiation. To increase patient comfort and minimize motion artifacts, it is important to keep image acquisition times as low as reasonably possible [12,20]. One important factor that impacts acquisition time is the inter-slice (out-of-plane) resolution. Acquiring multiple anisotropic 3-D images with a high in-plane resolution can effectively reduce and split up

scanning time, especially in populations who struggle to remain still [15]. Multi-view anisotropic acquisitions are already used in rheumatology and abdominal imaging [13]. Conventionally, the individually acquired images then need to be assessed separately, which is time consuming and error prone as lesions might only be visible in one of the scans, potentially leading to misinterpretations. Multi-view super-resolution (SR) enables the generation of a single high-resolution (HR) image given several low-resolution (LR) images for a simpler diagnostic process [3].

In this work, we introduce a novel self-supervised MR SR method that combines pretraining on a small, anisotropic dataset and per-patient adaptation during inference. Incorporating multi-patient information enables the use of cross-patient similarities while maintaining applicability to clinical scenarios with limited data availability. Accordingly, we split our optimization process into two phases: The "offline phase" which constitutes patient-agnostic feature extraction, and the "online phase" for additional patient-specific adaptation and subsequent inference to generate the final SR image. Minimizing runtimes for the online phase is critical for efficient deployment in production settings, such as hospitals, where fast image processing is essential. Specifically, we incorporate two orthogonal anisotropic LR MR scans to form a joint representation that facilitates the generation of accurate SR images. Following clinical routine, where typically no HR data is available, we train our network only on the available LR data using a fully self-supervised approach, circumventing the need for HR images as reference for training. By integrating a sparse coordinate-based loss function, we are able to 1) use images with varying resolution for training and 2) allow for arbitrary-scale upsampling. For evaluation, we perform extensive evaluations on two publicly available datasets [1,8,16] and demonstrate that our proposed method generalizes across datasets and to unseen MR sequences. We specifically highlight that our approach speeds up the SR process substantially compared to state-of-the-art (SOTA) self-supervised SR methods while preserving image quality.

1.1   Related Work

Several methods have been proposed to improve the resolution of MR image data [23,2,9,10]. They can be divided into single-image super-resolution (SISR) and multi-image super-resolution (MISR) methods [22,10]. MISR utilizes different (anisotropic) views or several sequentially acquired LR scans by merging them to generate a higher-resolution image. For MR SR specifically, multiple learning-based methods have been proposed that utilize matching LR and HR data as reference for training SR models. Lyu et al. [5] performed MR SR by matching features of the target LR contrast and the available reference HR image of a different domain. They used a transformer with a texture-preserving branch and contrastive learning to enhance the textural details of the SR image. To overcome the issue of generalization of SR approaches to different scales, Tan et al. [14] proposed an arbitrary scale SR approach for brain MR images. They

combined a Weight Prediction Network with SRGAN and achieved comparable performance to SOTA SR techniques. However, their approach relies on a large training dataset with HR images and they performed evaluation only on 2-D slices. Similar works propose SISR methods for MR or Computed Tomography (CT) SR but still rely on ground truth HR data [11,19]. Recently, implicit neural representations (INRs) have become more prominent in the context of image SR. In the work of Wu et al. [18], an INR for MR image SR was proposed. This method allows implicit modeling of a continuous function of sparsely sampled data points and projects the learned intensity distribution onto a more densely sampled grid. They achieved high reconstruction performance for several upsampling scales but still relied on paired isotropic LR and HR data. Distinct from the previously presented methods, other approaches allow for reference-free training of SR network while integrating arbitrary-scale upsampling. Zhao et al. propose a reference free SISR for MR image data [23]. McGinnis et al. [7] used two anisotropic LR MR images of different MR sequences to perform INR-based SR. While achieving very good SR reconstruction, their approach requires patient-specific training in the order of 10 to 20 minutes per image, which may not be clinically feasible.

2   Methodology

Our model, *tripleSR* for self-supervised SR, takes N paired LR images I_{ax}^i and I_{cor}^i from a training set $T = \{I_{ax}^i \in \mathbb{R}^{h \times w \times \frac{d}{e_i}}, I_{cor}^i \in \mathbb{R}^{h \times \frac{w}{e_i} \times d}\}_{i=1}^N$ as inputs, with e_i being the upsampling scale of the i-th LR-image pair. In Fig. 1 the complete training process is shown. Building on the methodology of [18], a convolutional encoder extracts feature maps $V_{ax}^i \in \mathbb{R}^{h \times w \times \frac{d}{e_i} \times 128}$ and $V_{cor}^i \in \mathbb{R}^{h \times \frac{w}{e_i} \times d \times 128}$ which are subsequently used for reconstruction. We adapt a residual dense network (RDN) [21], which is specifically developed for image SR, by removing the upsampling layer and applying 3-D convolution.

Each element in the feature map corresponds to a voxel in the respective LR image. In the HR image, coordinates may lie between those of the LR counterparts, along with the corresponding voxel data. To obtain a feature vector at a corresponding high-resolution coordinate x_{HR}, we perform trilinear interpolation on the feature maps V_{ax}^i and V_{cor}^i and generate two HR feature maps with the shape $h \times w \times d \times 128$ each. From this, we sample the feature vectors $v_{HR(ax)}^i$ and $v_{HR(cor)}^i$ at a specific HR coordinate x_{HR}. The two extracted feature vectors $v_{HR(ax)}^i$ and $v_{HR(cor)}^i$ corresponding to the HR coordinate are concatenated and fed to the decoder. The decoder consists of eight fully connected layers, each followed by ReLU activation with a residual connection after the fourth layer. Each inner layer has 512 input and output features. The decoder outputs the predicted voxel intensity $\hat{I}$ of the respective spatial coordinate x_{HR}.

We split our SR process into two phases. The first part is a training phase on a disjoint dataset which allows patient-agnostic feature extraction. We call this the "offline phase" as it only needs to be performed once. The second part
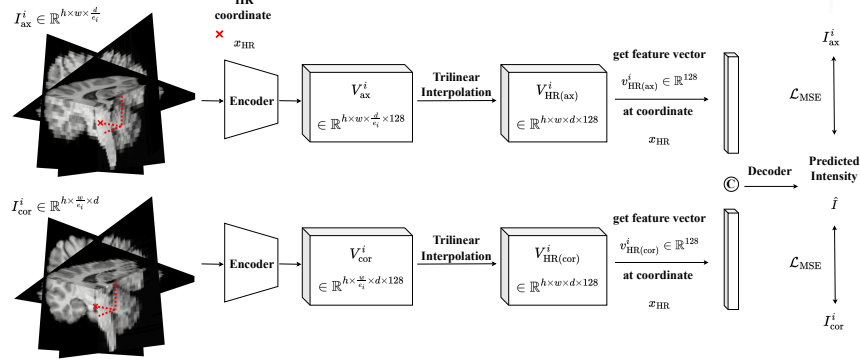
Fig. 1. Overview of the network architecture proposed in this study.

is an optional patient-specific fine-tuning and inference, which we refer to as the "online phase". Here, additional patient-specific details and anomalies can be learned. During both phases, no HR ground truth data is available. We optimize our model by using a loss function based on sparse continuous coordinates. Therefore, two orthogonal LR images of the same object are viewed as sparse representations of the same underlying instance. Instead of fixed voxel coordinates, we compute the loss for suitable locations in the continuous coordinate space where the coordinates are normalized to the range [-1, 1]. This enables the seamless SR of data with varying resolutions, as demonstrated in our experiments section. The applied loss computes the mean squared error ($\mathcal{L}_{\mathrm{MSE}}$) of the voxel intensity of reference LR images and the voxel intensity of the predicted image $\hat{I}$ at matching coordinates M_{ax} and M_{cor}, where

$$\mathcal{M}_{\mathrm{ax}} = \mathbf{C}_{\mathrm{HR}} \cap \mathbf{C}_{\mathrm{ax}} \text{ , and } \mathcal{M}_{\mathrm{cor}} = \mathbf{C}_{\mathrm{HR}} \cap \mathbf{C}_{\mathrm{cor}}. \tag{1}$$

Here, $\cap$ represents the intersection of available normalized coordinates. Coordinates corresponding to regions without measurements in the LR image are excluded from the loss calculation.

3   Experiments

3.1   Datasets

To assess performance, we simulated LR images from HR images and compared the SR results to the corresponding HR. We emphasize that the HR data was used only as a ground truth for evaluation and was not used during model training. We used two publicly available datasets for our evaluation: The Brain Tumor Segmentation (BraTS) dataset [1,8], consisting of brain MR images with an isotropic resolution of $1 \times 1 \times 1\,\mathrm{mm}^3$ and the HCP dataset [16], which includes 1200 MR images with an isotropic resolution of $0.7 \times 0.7 \times 0.7\,\mathrm{mm}^3$. For each

dataset and each MR sequence, we randomly selected 170 subjects (100 for training, 30 for validation, and 40 for final testing).

3.2   Preprocessing

For fair comparison, all images were registered using the NMI152 template with an isotropic resolution of $1\,\mathrm{mm}$. Background areas were removed, and images were normalized to $[0, 1]$. To generate the anisotropic images in the *training dataset*, we randomly downsampled each image in both axial and coronal direction using a random downsampling scale in the range $[2, 4]$ once at the beginning of the training. We applied a commonly used downsampling method where HR images are downsampled by cropping in the frequency domain [4,6].

For evaluation, we generated two anisotropically sampled LR MR images in the axial and coronal dimension each with an in-slice resolution two and four times the between-slice resolution. As the LR images were simulated from one HR scan, they can be assumed to be rigidly registered. The axial images were resampled to a resolution of $1{\times}1{\times}2\,\mathrm{mm}^3$ and $1{\times}1{\times}4\,\mathrm{mm}^3$, while the coronal images were resampled to a resolution of $1{\times}2{\times}1\,\mathrm{mm}^3$ and $1{\times}4{\times}1\,\mathrm{mm}^3$. All LR image coordinates were normalized to a reference frame in the range $[-1, 1]$.

3.3   Implementation Details

During the offline phase, we trained our network for 35 epochs with a batch size B of 10 (lr=0.0001, Adam optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.999$), with hyperparameters tuned on the validation set. Training and evaluation were executed on an NVIDIA A40 with Python 3.9 and PyTorch 1.13.1. In each batch, B randomly cropped LR patches p of one image were selected, where B equals the training batch size. The axial and coronal image patches were of size $10 \times (e_i \times e_i \times 1)$ and $10 \times (e_i \times 1 \times e_i)$, respectively. To enable efficient batch processing for different upsampling scales, 8000 random samples are selected from each patch. Each sample corresponds to one HR voxel coordinate. For the online phase, we used the trained model from the offline phase. For the optional patient-specific training, the model undergoes 10 additional epochs on LR patches from a single patient before the final inference.

During the online phase, two different settings were evaluated. First, patient-specific online training was performed to update the model before subsequent inference. Second, inference without additional training was conducted, which is referred to "ours w/o FT". We quantitatively evaluated models using Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) [17], which assess pixel accuracy and perceptual quality. To ensure a fair comparison, we computed the metrics only within the brain region and excluded the background, as BISR is only trained on the brain area.

3.4   Reference Approaches

As we trained our models without HR ground truth data, we performed evaluations by comparing them with other unsupervised/self-supervised methods. We

Table 1. Quantitative results for all MR sequences and SR methods on the BraTS and the HCP test set (trained and evaluated on the same MR sequence). Best results are bold, second best underlined. "Ours w/o FT" refers to results without online training.

		Resampling Scale	×2		×4	
			PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑
BraTS	T1 CE	cubic spline	33.794 ± 2.408	0.981 ± 0.006	29.93 ± 2.285	0.945 ± 0.01
		SMORE	37.939 ± 1.483	0.981 ± 0.003	30.678 ± 1.348	0.915 ± 0.009
		BISR	$\underline{42.727 \pm 2.576}$	$\underline{0.993 \pm 0.003}$	35.008 ± 2.183	0.963 ± 0.007
		Ours	$\mathbf{43.166 \pm 2.105}$	$\mathbf{0.994 \pm 0.001}$	$\mathbf{35.598 \pm 1.813}$	$\mathbf{0.966 \pm 0.005}$
		Ours w/o FT	41.358 ± 2.701	0.993 ± 0.002	$\underline{35.531 \pm 1.4655}$	$\underline{0.964 \pm 0.004}$
	T1	cubic spline	35.774 ± 3.745	0.991 ± 0.004	29.624 ± 3.629	0.962 ± 0.009
		SMORE	34.333 ± 2.64	0.982 ± 0.003	26.627 ± 1.817	0.906 ± 0.008
		BISR	$\mathbf{39.003 \pm 3.984}$	0.993 ± 0.006	$\underline{31.078 \pm 4.378}$	$\underline{0.967 \pm 0.017}$
		Ours	$\underline{38.832 \pm 4.121}$	$\mathbf{0.995 \pm 0.002}$	$\mathbf{31.169 \pm 3.33}$	$\mathbf{0.969 \pm 0.006}$
		Ours w/o FT	37.780 ± 4.061	$\underline{0.993 \pm 0.003}$	30.473 ± 3.632	0.965 ± 0.009
	T2	cubic spline	33.890 ± 1.876	0.986 ± 0.004	30.643 ± 2.168	0.959 ± 0.014
		SMORE	35.645 ± 1.388	0.984 ± 0.004	27.799 ± 1.217	0.91 ± 0.015
		BISR	$\mathbf{42.460 \pm 1.853}$	$\mathbf{0.997 \pm 0.001}$	$\mathbf{34.74 \pm 1.661}$	$\mathbf{0.976 \pm 0.007}$
		Ours	$\underline{42.218 \pm 1.750}$	$\mathbf{0.997 \pm 0.001}$	33.427 ± 1.640	0.973 ± 0.009
		Ours w/o FT	40.531 ± 2.190	0.995 ± 0.002	33.073 ± 1.412	0.969 ± 0.008
HCP	T1	cubic spline	29.820 ± 3.785	0.982 ± 0.006	22.341 ± 3.947	0.930 ± 0.023
		SMORE	30.603 ± 4.099	0.973 ± 0.007	22.533 ± 3.177	0.881 ± 0.019
		BISR	$\underline{31.249 \pm 4.574}$	$\underline{0.985 \pm 0.007}$	$\underline{22.758 \pm 4.636}$	0.931 ± 0.034
		Ours	$\mathbf{34.140 \pm 5.410}$	$\mathbf{0.990 \pm 0.004}$	$\mathbf{23.301 \pm 5.274}$	$\mathbf{0.938 \pm 0.025}$
		Ours w/o FT	29.950 ± 4.655	0.984 ± 0.007	22.523 ± 4.860	$\underline{0.933 \pm 0.026}$
	T2	cubic spline	30.561 ± 1.275	0.973 ± 0.004	27.336 ± 0.864	0.930 ± 0.008
		SMORE	31.444 ± 0.814	0.967 ± 0.004	25.817 ± 0.597	0.877 ± 0.010
		BISR	$\underline{35.321 \pm 0.858}$	$\underline{0.985 \pm 0.002}$	$\mathbf{29.862 \pm 0.531}$	$\mathbf{0.947 \pm 0.004}$
		Ours	$\mathbf{35.482 \pm 0.877}$	$\mathbf{0.986 \pm 0.001}$	29.372 ± 1.231	0.941 ± 0.006
		Ours w/o FT	34.841 ± 1.155	$\underline{0.985 \pm 0.002}$	$\underline{29.681 \pm 1.055}$	$\underline{0.942 \pm 0.006}$

used cubic spline interpolation, a SISR method SMORE [23] and an INR MISR approach referred to as BISR [7]. Originally, [7] used two sequences as input and optimized the SR jointly. We adapted the approach to directly predict only one SR image instead of two images with different MR sequences. Cubic spline interpolation and the SISR method SMORE [23] are performed individually on the axial and coronal LR images of each patient. The coronal and axial metric values are then averaged.

4   Quantitative and Qualitative Results

We evaluated on 40 brain MR images each for T1, T1 contrast-enhanced (CE), and T2 sequences for BraTS and T1 and T2 for HCP which are trained either on images from the same sequence or from any of the other sequences (for 2× and 4×

Table 2. Duration of the offline and online phase for each method in minutes. Best, i.e., shortest times are highlighted in bold.

|	Resampling Scale	×2		×4	
		offline	online	offline	online
BraTS T1	SMORE	0	9.09 ± 0.29	0	9.16 ± 0.32
	BISR	0	19.25 ± 1.32	0	10.59 ± 0.67
	Ours	130.99 ± 5.15	3.32 ± 0.31	130.99 ± 5.15	1.55 ± 0.15
	Ours w/o FT	130.99 ± 5.15	**1.94 ± 0.22**	130.99 ± 5.15	**0.73 ± 0.08**

Table 3. Quantitative results for different training settings with patient-specific online training. "Train on" refers to the used training data and "test on" refers to the specific test set. "Br" is used as abbreviation for BraTS.

Resampling Scale		×2		×4	
Train on	Test on	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑
Br T1 CE	Br T1 CE	43.166 ± 2.105	0.994 ± 0.001	35.598 ± 1.813	0.966 ± 0.005
Br T1	Br T1 CE	42.680 ± 2.435	0.994 ± 0.001	35.275 ± 1.818	0.963 ± 0.005
Br T2	Br T1 CE	42.762 ± 2.285	0.994 ± 0.001	35.476 ± 1.66	0.963 ± 0.005
Br T1	Br T1	38.832 ± 4.121	0.995 ± 0.002	31.169 ± 3.33	0.969 ± 0.006
Br T1 CE	Br T1	40.178 ± 4.388	0.995 ± 0.003	31.457 ± 3.382	0.970 ± 0.006
Br T2	Br T1	39.981 ± 4.259	0.995 ± 0.002	31.474 ± 3.245	0.969 ± 0.007
HCP T1	Br T1	39.410 ± 4.200	0.995 ± 0.002	31.450 ± 3.674	0.970 ± 0.006
Br T2	Br T2	42.218 ± 1.750	0.997 ± 0.001	33.427 ± 1.640	0.973 ± 0.009
Br T1 CE	Br T2	42.360 ± 1.962	0.997 ± 0.001	34.290 ± 1.359	0.976 ± 0.007
Br T1	Br T2	42.226 ± 1.958	0.997 ± 0.002	33.610 ± 1.467	0.973 ± 0.008
HCP T2	Br T2	42.456 ± 2.311	0.997 ± 0.002	33.548 ± 1.819	0.975 ± 0.009

upsampling). Table 1 summarizes the quantitative results for all SR approaches, including mean and standard deviation. We achieve higher or comparable PSNR and SSIM values on all MR sequences and upsampling scales. Using the additional patient-specific online training helps to improve performance in all cases. Still, without additional training (ours w/o FT), the performance is comparable to the best reference approach at approx. 10-15× faster inference times. Table 3 shows results obtained from testing on a dataset distinct from the one used for training. Even when the network is trained on a completely different dataset, the SR performance is comparable or even better compared to training on the same dataset. A reason for the better performance of the model trained on Br T1 CE for the BR T1 test data may be that the contrast-enhanced data potentially provides richer features and patterns, which allows the model to generalize to non-enhanced images.

The time required to generate a SR reconstruction for each method can be divided into offline and online phases. The results are shown in Table 2. Our model is pre-trained once offline which takes approximately 130 minutes. We average the time for the offline phase for three runs. During the online phase,
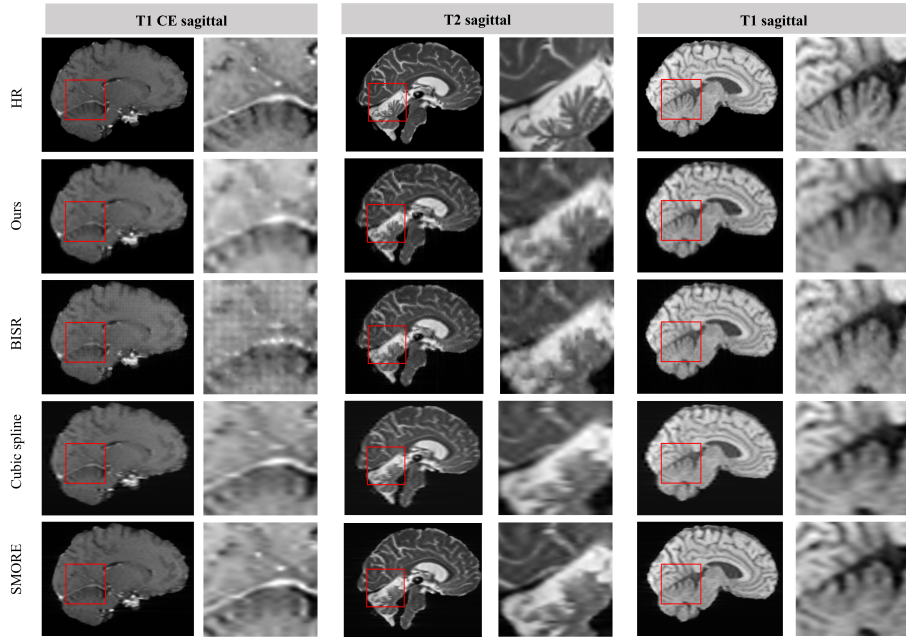
Fig. 2. Reference HR image and qualitative SR results for all MR sequences in the sagittal plane of the BraTS test set where no in-plane HR images are available.

our approach is approximately 10-15× faster when using the offline model w/o patient-specific fine-tuning and still 6× faster when fine-tuning for a specific patient, compared to the reference approach BISR.

In Fig. 2, qualitative examples of the BraTS dataset for each MR sequence and SR method can be seen in comparison to the actual HR image. The full image and a zoomed-in version can be seen for each sample. Overall, the SR results obtained using SMORE and cubic spline interpolation appear blurred or noisy. In some cases, BISR introduces noticeable blocking artifacts. In contrast, our method preserves accurate anatomical structures while minimizing noise.

5   Discussion and Conclusion

In this work, we presented a novel approach for MR image SR. In both quantitative and qualitative evaluations, our method produces HR images that match or surpass the quality of current reference approaches. A key strength lies in its substantially reduced patient-specific training time during the online phase, which is critical for translating such techniques into clinical practice. Notably, after generating only 14 SR images with a resampling factor of 4 using BISR, the time invested in our offline phase is already offset. This efficiency gain enables a ten-fold increase in throughput for the generation of SR data. This could help to

further drive acceptance of deep learning-supported diagnostic imaging in clinical workflows. Additionally, the sparse coordinate loss enables the integration of LR images with varying resolution scales, further enhancing the versatility and applicability of our approach. The patient-specific training allows strong adaptation and has lower risk of traditional overfitting due to the inherent inductive bias of the network.

Finally, the presented method demonstrates promising capabilities by consistently achieving high-quality results, regardless of whether training is performed on data from the same cohort or on a disjoint dataset. Further experiments are needed to validate our findings and demonstrate the benefits of using the reconstructed SR images in diagnostic assessments.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bakas, S., et al.: Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. Scientific data **4**, 170117 (2017). `https://doi.org/10.1038/sdata.2017.117`
2. Chen, Y., et al.: Brain MRI super resolution using 3D deep densely connected neural networks. In: 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018). pp. 739–742. IEEE, Piscataway, NJ (2018). `https://doi.org/10.1109/ISBI.2018.8363679`
3. Jia, Y., Gholipour, A., He, Z., 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`, `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5534179/`
4. Li, G., et al.: Rethinking multi-contrast MRI super-resolution: Rectangle-window cross-attention transformer and arbitrary-scale upsampling. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 21173–21183. IEEE (2023). `https://doi.org/10.1109/ICCV51070.2023.01941`
5. Lyu, J., et al.: Multicontrast MRI super-resolution via transformer-empowered multiscale contextual matching and aggregation. IEEE transactions on neural networks and learning systems **PP** (2023). `https://doi.org/10.1109/TNNLS.2023.3250491`
6. Masutani, E.M., Bahrami, N., Hsiao, A.: Deep learning single-frame and multiframe super-resolution for cardiac MRI. Radiology **295**(3), 552–561 (2020). `https://doi.org/10.1148/radiol.2020192173`
7. McGinnis, J., et al.: Single-subject multi-contrast MRI super-resolution via implicit neural representations. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2023, Lecture Notes in Computer Science, vol. 14227, pp. 173–183. Springer Nature Switzerland and Imprint Springer, Cham (2023). `https://doi.org/10.1007/978-3-031-43993-3{_}17`
8. Menze, B.H., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). IEEE Transactions on Medical Imaging **34**(10), 1993–2024 (2015). `https://doi.org/10.1109/TMI.2014.2377694`

9. Nasrollahi, K., Moeslund, T.B.: Super-resolution: a comprehensive survey. Machine Vision and Applications **25**(6), 1423–1468 (2014). `https://doi.org/10.1007/s00138-014-0623-4`, `https://link.springer.com/article/10.1007/s00138-014-0623-4#Sec3`

10. Ooi, Y.K., Ibrahim, H.: Deep learning algorithms for single image super-resolution: A systematic review. Electronics **10**(7), 867 (2021). `https://doi.org/10.3390/electronics10070867`, `https://www.mdpi.com/2079-9292/10/7/867#B6-electronics-10-00867`

11. Peng, C., et al.: Saint: Spatially aware interpolation network for medical slice synthesis. pp. 7750–7759 (2020), `https://openaccess.thecvf.com/content_CVPR_2020/html/Peng_SAINT_Spatially_Aware_Interpolation_NeTwork_for_Medical_Slice_Synthesis_CVPR_2020_paper.html`

12. Riek, J.K., Tekalp, A.M., Smith, W.E., Kwok, E.: Out-of-plane motion compensation in multislice spin-echo MRI. IEEE Transactions on Medical Imaging **14**(3), 464–470 (1995). `https://doi.org/10.1109/42.414611`

13. Shetty, A.S., et al.: Body MRI: Imaging protocols, techniques, and lessons learned. Radiographics : a review publication of the Radiological Society of North America, Inc **42**(7), 2054–2074 (2022). `https://doi.org/10.1148/rg.220025`

14. Tan, C., Zhu, J., Lio', P.: Arbitrary scale super-resolution for brain MRI images. Artificial Intelligence Applications and Innovations **583**, 165–176 (2020). `https://doi.org/10.1007/978-3-030-49161-1{_}15`, `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7256400/`

15. Thrower, S.L., et al.: The effect of slice thickness on contours of brain metastases for stereotactic radiosurgery. Advances in radiation oncology **6**(4), 100708 (2021). `https://doi.org/10.1016/j.adro.2021.100708`, `https://pubmed.ncbi.nlm.nih.gov/34124413/`

16. van Essen, D.C., et al.: The WU-Minn human connectome project: an overview. NeuroImage **80**, 62–79 (2013). `https://doi.org/10.1016/j.neuroimage.2013.05.041`, `https://www.sciencedirect.com/science/article/pii/s1053811913005351`

17. 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`

18. Wu, Q., et al.: 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 (2023). `https://doi.org/10.1109/JBHI.2022.3223106`, `https://arxiv.org/pdf/2110.14476`

19. Yu, P., et al.: RPLHR-CT dataset and transformer baseline for volumetric super-resolution from CT Scans. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2022, Lecture Notes in Computer Science, vol. 13436, pp. 344–353. Springer Nature Switzerland and Imprint Springer, Cham (2022). `https://doi.org/10.1007/978-3-031-16446-0{_}33`

20. Zaitsev, M., Maclaren, J., Herbst, M.: Motion artifacts in MRI: A complex problem with many partial solutions. Journal of magnetic resonance imaging : JMRI **42**(4), 887–901 (2015). `https://doi.org/10.1002/jmri.24850`

21. Zhang, Y., et al.: Residual dense network for image super-resolution. pp. 2472–2481 (2018), `https://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_Residual_Dense_Network_CVPR_2018_paper.html`

22. Zhao, C., et al.: Applications of a deep learning method for anti-aliasing and super-resolution in MRI. Magnetic Resonance Imaging **64**, 132–141 (2019). `https://doi.org/10.1016/j.mri.2019.05.038`, `https://www.sciencedirect.com/science/article/pii/S0730725X18306507`

23. Zhao, C., et al.: Smore: A self-supervised anti-aliasing and super-resolution algorithm for MRI using deep learning. IEEE Transactions on Medical Imaging **40**(3), 805–817 (2021). `https://doi.org/10.1109/TMI.2020.3037187`