Reconstructing Isotropic 3D Cervical MRI from Anisotropic Clinical Scans via Anatomy-Style Decoupled INRs

Qi Zhang[1,*], Xiuyuan Chen[2,*], Yibo Hu[1], Hongxing Shen[2,✉], and Jianqi Sun[1,✉]

[*] These authors contributed equally to this work.
[✉] Co-corresponding authors.
[1]School of Biomedical Engineering, Shanghai Jiao Tong University, Shanghai, China
[2]Department of Spine Surgery, Renji Hospital, School of Medicine, Shanghai Jiao Tong University, Shanghai, China
milesun@sjtu.edu.cn

Abstract. Clinical cervical Magnetic Resonance Imaging (MRI) routinely relies on multi-view anisotropic 2D scans with large slice gaps, causing severe through-plane information loss that hinders 3D diagnosis. Reconstructing an isotropic 3D volume from such scans is challenging due to heterogeneous intensity distributions and partially mismatched fields of view (FOVs). We propose an anatomy-style decoupled hybrid implicit neural representation (INR) framework for multi-view super-resolution without paired high-resolution supervision. Our approach aligns orthogonal scans within a unified coordinate system, utilizes a two-stage feature extraction module to inject anatomical priors, and introduces a stochastic view dropout strategy with learnable style embeddings to explicitly disentangle shared anatomy from view-specific intensities. Extensive experiments on 120 clinical cervical MRI pairs demonstrate artifact-free isotropic 3D reconstruction with 50% rated as excellent by a blinded clinical expert, achieving axial and sagittal FIDs of 45.14 and 55.50 and a Lesion Dice of 0.673. Our code is available: https://github.com/zhibaishouheilab/JointSR

Keywords: Magnetic Resonance Imaging · Super-Resolution · Implicit Neural Representation · Multi-view Fusion · Image Decoupling.

1 Introduction

Magnetic Resonance Imaging (MRI) is the clinical gold standard for assessing cervical pathologies. To mitigate acquisition time and motion artifacts, clinical protocols typically acquire anisotropic 2D slices in orthogonal planes, such as sagittal and axial views [9,1]. While providing high in-plane resolution, large slice gaps cause severe through-plane information loss. This degradation often obscures critical peripherally located lesions (e.g., far-lateral disc herniations

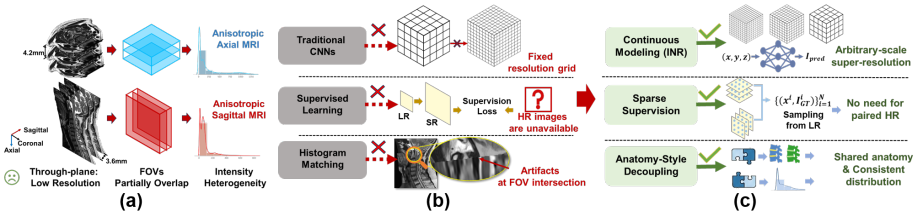

Fig. 1. Illustration of the (a) clinical challenges in cervical MRI scans, (b) Limitations of existing super-resolution methods, and (c) Solutions of the proposed method.

and foraminal stenosis) whose delayed identification can lead to severe radicular pain and irreversible functional decline. Consequently, radiologists must cross-reference these scans, which is a demanding task complicated by patient micro-motions, distinct intensity distributions, and partially overlapping fields of view (FOVs) [5]. Furthermore, the lack of isotropic 3D data hinders downstream computer-aided diagnosis (CAD) systems, such as surgical navigation and 3D segmentation [7]. Thus, developing multi-view super-resolution (SR) techniques to integrate low-resolution (LR) anisotropic scans into a cohesive high-resolution (HR) volume is of great importance.

Reconstructing isotropic volumes directly from orthogonal anisotropic scans is an ill-posed inverse problem. Deep learning-based SR methods, particularly convolutional neural networks (CNNs), have emerged for clinical applications [17,3,4]. While effective on 2D images or single-view volumes, traditional CNNs are constrained by fixed spatial grids, rendering them inflexible to varying clinical resolutions (Fig. 1(b)). Implicit Neural Representations (INRs) offer a compelling alternative [12,7]. By utilizing coordinate-based Multi-Layer Perceptrons (MLPs), INRs model continuous intensity fields for arbitrary-scale reconstruction. Recent advances have extended INRs to slice-to-volume reconstruction, e.g., NeSVoR [13] for fetal MRI. However, optimizing pure INRs from scratch is slow. Hybrid INRs address this by combining CNN-based feature extractors with implicit decoders, injecting local image features to accelerate training and enhance arbitrary-scale SR performance [12].

Furthermore, many existing methods rely on paired HR images for supervision [12,6], severely restricting their clinical applicability. To overcome this, self-supervised models like SMORE extract high-resolution 2D patches from in-plane slices [16]. Advancing toward multi-plane self-supervised learning, recent approaches utilize a sparse coordinate loss derived from orthogonal LR scans, eliminating the need for HR ground truths [7,9]. However, primarily developed for brain MRI, these techniques assume strictly overlapping FOVs and identical intensity distributions. In contrast, clinical cervical MRI scans (Fig. 1(a)) exhibit highly heterogeneous intensities and mismatched FOVs—differences in frequency-encoding directions and coil sensitivity profiles across slice orientations further exacerbate this heterogeneity. Simple fusion techniques (e.g., histogram matching) fail to achieve distribution consistency, leading to intensity artifacts

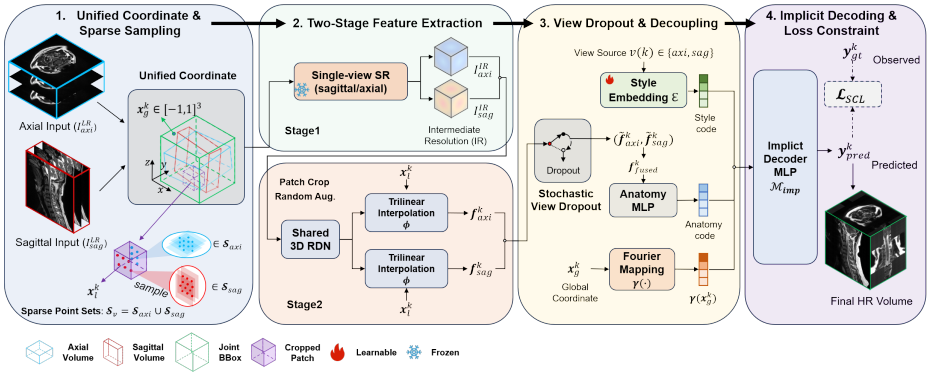

Fig. 2. Workflow of our proposed anatomy-style decoupled hybrid INR framework.

and blurred anatomical details at intersections. Fundamentally, these methods lack the capability to explicitly decouple shared anatomical geometry from view-specific intensity characteristics.

To bridge these challenges, we propose a novel anatomy-style decoupled hybrid implicit neural representation framework. Our main contributions are threefold: (1) **Unified Coordinate System:** We introduce a joint bounding box-based coordinate normalization strategy, enabling continuous arbitrary-scale querying across orthogonal multi-view scans. (2) **Two-Stage Prior Enhancement:** A pre-trained single-view SR module paired with a 3D Residual Dense Network (RDN) extracts dense features to accelerate convergence and restore high-frequency details. (3) **Anatomy-Style Decoupling & FOV Fusion:** A stochastic view dropout strategy and a learnable style embedding explicitly disentangle shared anatomy from view-specific intensities. Optimized via sparse coordinate loss, the proposed method ensures seamless FOV fusion and robust 3D SR, generating isotropic volumes that significantly improve downstream tasks such as 3D segmentation and lesion detection.

2 Methodology

2.1 Problem Formulation and Overall Architecture

Given a pair of orthogonal, anisotropic cervical MRI scans from the same patient, denoted as the sagittal view I_{sag} and axial view I_{axi}, we aim to reconstruct a high-resolution, isotropic, and field-of-view (FOV) fused 3D volume I_{iso}. To address inherent spatial misalignment, distinct intensity distributions, and partially overlapping FOVs, we propose an anatomy-style decoupled hybrid implicit neural representation framework (Fig. 2).

2.2 Unified Coordinate System Construction and Sparse Sampling

After rigidly registering I_{sag} to the reference coordinate system of I_{axi}, we compute a joint bounding box $[p_{\min}, p_{\max}] \in \mathbb{R}^3$ encompassing both physical volumes. For any coordinate x, its global coordinate is normalized as $x_g = 2(x - p_{\min})/(p_{\max} - p_{\min}) - 1 \in [-1, 1]^3$. To ensure computational feasibility across the large 3D volume, we adopt a patch-based training strategy. Specifically, we partition the bounded box into localized patches. For a given patch defined by its physical bounds $[p_{\min}^{\text{patch}}, p_{\max}^{\text{patch}}]$, the coordinate of any point within it is independently normalized to a local coordinate $x_l = 2(x - p_{\min}^{\text{patch}})/(p_{\max}^{\text{patch}} - p_{\min}^{\text{patch}}) - 1 \in [-1, 1]^3$. Within each patch, we extract valid sparse point sets $\mathcal{S}_v = \{(x_g^i, x_l^i, y_v^i)\}_{i=1}^{N_v}$ for $v \in \{\text{sag}, \text{axi}\}$ from the LR inputs, where y is the observed intensity. During training, we randomly sample a batch of N points $\mathcal{B} \subset \mathcal{S}_{\text{sag}} \cup \mathcal{S}_{\text{axi}}$ for sparse supervision. Let $v(k) \in \{\text{sag}, \text{axi}\}$ denote the native source view of any sampled point $k \in \mathcal{B}$.

2.3 Two-Stage Anatomical Feature Extraction

To bridge the spatial domain gap between orthogonal anisotropic inputs, raw LR images are first upsampled to isotropic intermediate-resolution (IR) volumes utilizing a pre-trained single-view SR model (SA-INR) [11]. Directly learning 3D features from raw LR scans with large slice gaps is challenging due to severe through-plane undersampling; the single-view SR provides an essential "warm-up" that recovers anatomically plausible structures. This prior is used solely for feature extraction—it does not provide direct supervision to the reconstruction.

We crop patches from IR volumes and apply random intensity augmentations (gamma, contrast, and brightness perturbations). This forces the network to learn robust anatomical features rather than memorizing absolute intensities as a shortcut to the decoder. A shared 3D Residual Dense Network (RDN) [15] extracts dense feature volumes F_{sag} and F_{axi}. For each point $k \in \mathcal{B}$, view-specific features are extracted via differentiable trilinear interpolation Φ: $f_v^k = \Phi(F_v, x_l^k)$ for $v \in \{\text{sag}, \text{axi}\}$.

2.4 Stochastic View Dropout and Anatomy-Style Decoupling

Since clinical FOVs only partially overlap, the network must handle single-view FOV during inference. In overlapping regions, we dynamically construct effective features during training by sampling from three dropout states:

$$(\tilde{f}_{\text{sag}}^k, \tilde{f}_{\text{axi}}^k) \sim \mathcal{U}\left\{(f_{\text{sag}}^k, f_{\text{axi}}^k), (f_{\text{sag}}^k, f_{\text{sag}}^k), (f_{\text{axi}}^k, f_{\text{axi}}^k)\right\} \tag{1}$$

where $\mathcal{U}$ denotes a discrete uniform distribution. This forces the model to extrapolate missing cross-view features solely from one available view. These are then concatenated and fused via an MLP: $f_{\text{fused}}^k = \mathcal{M}_{\text{fuse}}([\tilde{f}_{\text{sag}}^k, \tilde{f}_{\text{axi}}^k])$.

An anatomy MLP extracts an anatomical latent code $z_{\text{anat}}^k = \mathcal{M}_{\text{anat}}(f_{\text{fused}}^k)$ to decouple shared anatomy from intensity. Simultaneously, a patient-specific,

continuously learnable style embedding $\mathcal{E}$ is queried to extract a style code $s_{v(k)} = \mathcal{E}(v(k)) \in \mathbb{R}^d$ (d=64), globally optimized to represent view-specific intensity characteristics. Since $\mathcal{M}_{\text{imp}}$ is continuous, linear interpolation in the learned style space yields smooth intensity transitions, contributing to the interpretability of the disentangled representation. Finally, mapping the global coordinate to high-frequency Fourier features [10,7] $\gamma(x_g^k)$, the shared implicit decoder $\mathcal{M}_{\text{imp}}$ predicts the final intensity: $I_{\text{pred}}^k = \mathcal{M}_{\text{imp}}(z_{\text{anat}}^k, \gamma(x_g^k), s_{v(k)})$.

Loss Function: The network is optimized end-to-end without paired HR ground truths via a sparse coordinate loss [9], which computes the Mean Squared Error (MSE) strictly on valid observed points: $\mathcal{L}_{\text{SCL}} = \frac{1}{|\mathcal{B}|} \sum_{k \in \mathcal{B}} \left| I_{\text{pred}}^k - y^k \right|^2$.

3 Experiments and Results

3.1 Datasets and Preprocessing

We collected 120 paired dual-view (sagittal and axial) anisotropic cervical MRI scans from an in-house cohort (randomly split into 100/20 for training/testing). This retrospective study was approved by the Institutional Review Board (IRB) of Renji Hospital, Shanghai Jiao Tong University School of Medicine, and written informed consent was obtained from all participants. In-plane resolutions range from 0.25×0.25 mm^2 to 0.75×0.75 mm^2, with 3.3–4.5 mm slice gaps. We target a $0.4 \times 0.4 \times 0.4$ mm^3 isotropic volume reconstruction. To establish spatial correspondence, sagittal scans were rigidly registered to the axial coordinate system via Mutual Information and Euler 3D transformations, following expert radiological verification. All images were min-max normalized to [0, 1].

3.2 Training, Inference, and Implementation Details

The Stage 1 single-view SA-INR was pre-trained on the Spine Generic dataset [2] (270 isotropic $0.8 \times 0.8 \times 0.8$ mm^3 MRIs), simulating LR inputs via random downsampling (up to 10x). This pre-trained SA-INR offline-generates $0.8 \times 0.8 \times 0.8$ mm^3 intermediate-resolution (IR) volumes from native inputs to provide anatomical features, without providing direct supervision.

Experiments were conducted on an RTX 4090 GPU (24 GB) using Adam (batch size 4, learning rates: 5×10^{-5} for 3D RDN, 2×10^{-4} for decoder and style embeddings). Training used $32 \times 32 \times 32$ mm^3 patches and 16,000 randomly sampled points per patch for sparse supervision. To handle clinical heterogeneity, we employ instance-level Test-Time Optimization (TTO) initialized from the model pre-trained on the training set: a 2-epoch *Warm-up* stage freezes the 3D RDN to optimize only the decoder and style codes for rapid adaptation, followed by a 10-epoch *Full-tune* stage of all weights. Inference targets a $0.4 \times 0.4 \times 0.4$ mm^3 grid within the joint bounding box. To eliminate boundary artifacts under memory constraints, we padded sliding-window inputs to $48 \times 48 \times 48$ mm^3 for broader context, retaining only central $32 \times 32 \times 32$ mm^3 predictions for final reconstruction.

3.3 Comparative Experiments and Evaluation Metrics

We compared our method against three baselines: Cubic Interpolation, SA-INR [11], and SCL-INR [9]. To accommodate dual-view intensities, SCL-INR was modified by duplicating its MLP decoder. SCL-INR was excluded from clinical expert evaluation due to poor reconstruction quality ($\text{Dice}_S = 0.513$). Ref. [7] was not compared as it requires identical input dimensions, incompatible with variable-FOV cervical data. Lacking native multi-view fusion, Cubic and SA-INR were augmented via a post-processing pipeline: the sagittal SR was histogram-matched to the axial reference and non-overlapping regions appended into the axial volume. Critically, Cubic's axial output is identical to the native axial in-plane data in the overlapping region, constituting a self-comparison that yields artificially low axial FID/KID while collapsing sagittal fidelity (FID = 100.96 vs. Ours 55.50). Final 3D predictions used the axial decoder branch for SCL-INR and the axial style embedding for our method. Ablation studies validated core components: (1) w/o Style Code (distinct MLP decoders per view), (2) w/o Prior (features directly from LR inputs), (3) w/o Stochastic View Dropout (SVD) (disabling dropout).

Due to the absence of paired HR ground truths—a field-wide challenge as isotropic 3D cervical MRI is rarely acquired and no public 3D HR dataset exists—we employed a triple-validation strategy: (1) FID/KID against native in-plane slices for distribution-level fidelity; (2) downstream segmentation Dice (sagittal: spine/lesion by [14], axial: muscles by [8], all labels manually verified by a radiologist) as an anatomical faithfulness proxy; (3) blinded clinical expert evaluation on all 20 test cases using a 5-point Likert scale assessing clinically meaningful artifacts (hallucination, stitching, intensity inconsistencies).

3.4 Results

Visualization and Downstream Validation. Figure 3 visualizes super-resolution and segmentation results. Original anisotropic inputs suffer from severe through-plane degradation (e.g., obscured disc herniations), and comparison methods fail to reconstruct both planes accurately. Single-view methods (Cubic, SA-INR) rely on histogram matching, producing stark artifacts at view intersections and blurred sagittal anatomy. SCL-INR lacks view dropout and decoupling, causing intensity shifts and blurred boundaries in non-overlapping regions. Conversely, our method achieves seamless FOV fusion and unified intensity distribution by explicitly decoupling style from anatomy, yielding clear boundaries and precise micro-lesion segmentation (yellow arrows).

Quantitative Results. Table 1 summarizes quantitative performance. Comparison methods exhibit extreme imbalance: Cubic achieves deceptively low axial scores (FID 1.641, KID 0.000) because its pipeline copies native axial data into the overlapping region (self-comparison), collapsing sagittal reconstruction (FID 100.96, Spine Dice 0.723). SA-INR and ablations similarly overfit the axial distribution at the expense of sagittal quality (e.g., w/o SVD Spine Dice drops to 0.558). In contrast, our framework achieves balanced fidelity across both planes

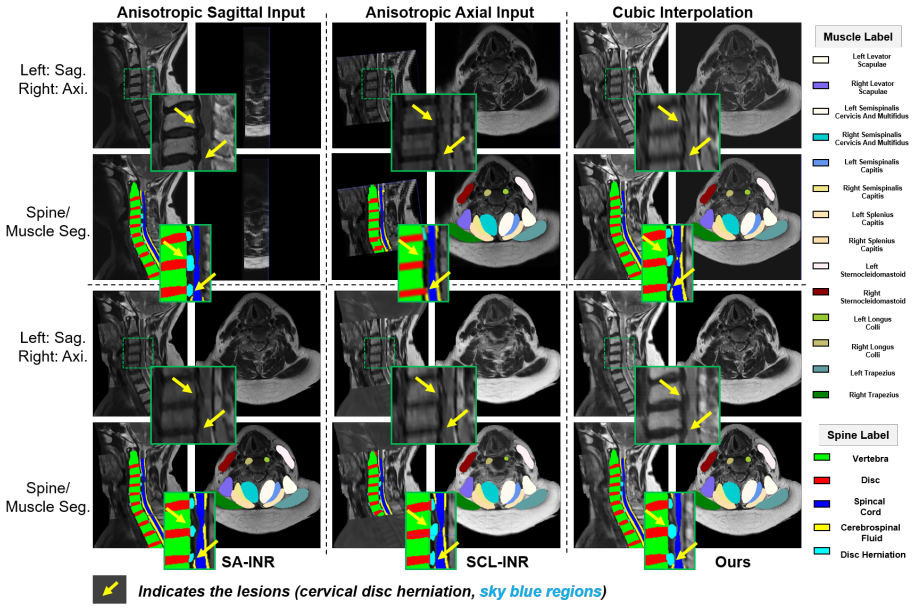

Fig. 3. Visual comparison of super-resolution outputs, downstream segmentation tasks, and lesion detection across different methods. The left panel shows the sagittal view (segmentation of spine structures and lesions), and the right panel shows the axial view (segmentation of corresponding paravertebral muscles).

Table 1. Quantitative comparison of super-resolution and downstream segmentation performance. Dice_M, Dice_S, Dice_L: Dice score for muscles, spine, and lesions

Method	Methods	Axial View			Sagittal View				Time (min)
		KID ↓	FID ↓	Dice_M ↑	KID ↓	FID ↓	Dice_S ↑	Dice_L ↑	
Comp.	Cubic	**0.000**	**1.641**	**0.941**	0.073	100.964	0.723	0.529	0.08
	SA-INR [11]	0.007	23.606	0.935	0.052	100.349	0.778	0.532	1.67
	SCL-INR [9]	0.005	21.023	0.936	0.193	188.446	0.513	0.337	3.88
Abla.	w/o Style Code	0.011	35.645	0.931	0.071	89.147	0.795	0.609	4.20
	w/o Prior	0.022	51.080	0.923	0.033	64.727	0.841	0.634	
	w/o SVD	0.008	28.560	0.933	0.115	124.320	0.558	0.562	
	Ours	0.017	45.145	0.928	**0.025**	**55.504**	**0.885**	**0.673**	

(KID: 0.017/0.025, FID: 45.145/55.504). For downstream segmentation, we attain Muscle Dice of 0.928 and Spine Dice of 0.885. Crucially, our Lesion Dice of 0.673 substantially outperforms all comparisons (Cubic: 0.529, SA-INR: 0.532, SCL-INR: 0.337)—the decisive metric for clinical utility. The marginal computational overhead (4.2 vs. 3.8 min per subject) is an acceptable trade-off.

Ablation Studies. Table 1 and Fig. 4(a) validate our design. **w/o Style Code:** Independent MLP decoders overfit native-view sparse points, penalizing cross-view structures; our shared MLP with style embeddings ensures balanced syn-

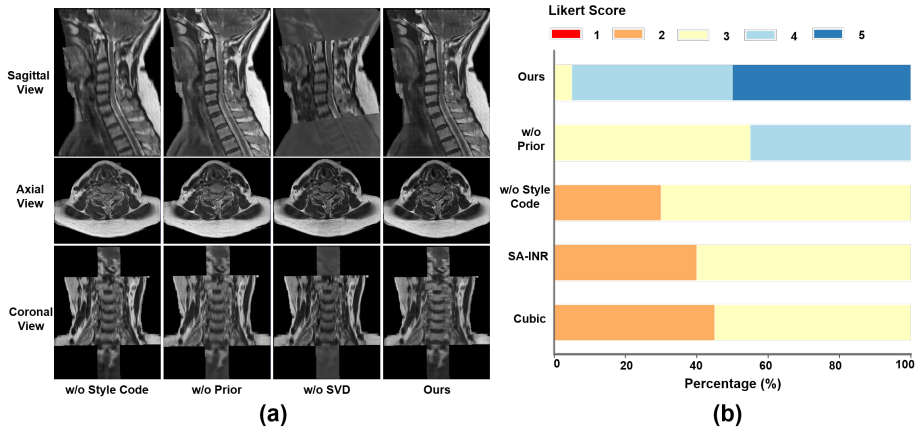

Fig. 4. Visual and clinical assessments. (a) Qualitative ablation results illustrating the central slices across sagittal, axial, and coronal planes. (b) Stacked bar chart detailing the Likert scale scores from the clinical evaluation (sample size: 20).

thesis. **w/o Prior:** Omitting the single-view prior uniformly degrades metrics, confirming its necessity for high-frequency detail preservation. **w/o SVD:** Disabling stochastic dropout causes catastrophic failure in non-overlapping sagittal FOVs (Spine Dice drops to 0.558), proving that forced single-view feature dropping during training is essential for cross-view FOV fusion.

Clinical Expert Evaluation. To validate diagnostic utility, an experienced radiologist blinded to the reconstruction method scored all 20 test cases on a 5-point Likert scale (1: poor, non-diagnostic; 5: excellent, artifact-free). The expert was specifically instructed to assess each reconstruction for clinically meaningful artifacts, including hallucinated structures, stitching discontinuities, and intensity inconsistencies, which directly addresses the Meta-Reviewer's concern about artifact verification. Fig. 4(b) shows our framework consistently scored ≥ 3, with 50% achieving a perfect 5, confirming great diagnostic value. The high proportion of artifact-free ratings demonstrates that the expert independently verified the clinical acceptability of our reconstructions.

4 Conclusion

We propose a hybrid INR framework, optimized without paired HR data, to reconstruct isotropic 3D volumes from orthogonal anisotropic cervical MRI scans. By integrating a unified coordinate system, stochastic view dropout, and learnable style embeddings, our method decouples anatomical geometry from view-specific intensities, facilitating 3D segmentation and micro-lesion detection. Several limitations warrant discussion. First, spatial misalignment between inputs can cause anatomical distortions; future work will develop a joint registration–super-resolution framework. Second, the current trilinear interpolation for fea-

ture querying may degrade at coarser grids due to the curse of dimensionality, motivating investigation of learnable interpolation strategies. Finally, recruiting subjects for dedicated isotropic 3D cervical MRI acquisitions would enable direct PSNR/SSIM computation, addressing the field-wide absence of paired 3D HR ground truth.

Acknowledgment. This work was supported in part by the National Natural Science Foundation of China Grant 62471297, in part by the Ministry of Science and Technology of the People's Republic of China under Grant 2023YFC2411401, in part sponsored by the Fundamental Research Funds for the Central Universities (No. YG2025ZD03), in part by the National Key Research and Development Program of China (2025YFE0214000), and in part by the Talent Cultivation Project of Renji Hospital (RJTJ25-QN-031).

Disclosure of Interests. The authors have no competing interests to declare relevant to this article's content.

References

1. Benisty, R., Shteynman, Y., Porat, M., Ilivitzki, A., Freiman, M.: Simple: Simultaneous multi-plane self-supervised learning for isotropic mri restoration from anisotropic data. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 551–561. Springer (2025)
2. Cohen-Adad, J., Alonso-Ortiz, E., Abramovic, M., Arneitz, C., Atcheson, N., Barlow, L., Barry, R.L., Barth, M., Battiston, M., Büchel, C., et al.: Open-access quantitative mri data of the spinal cord and reproducibility across participants, sites and manufacturers. Scientific data **8**(1), 219 (2021)
3. Ding, Y., Zhou, T., Wu, Y.: Cross-contrast fusion and aggregation network for multi-contrast mri super-resolution. In: Proceedings of the 2023 9th International Conference on Computing and Artificial Intelligence. pp. 102–107 (2023)
4. Du, J., He, Z., Wang, L., Gholipour, A., Zhou, Z., Chen, D., Jia, Y.: Super-resolution reconstruction of single anisotropic 3d mr images using residual convolutional neural network. Neurocomputing **392**, 209–220 (2020)
5. Graf, R., Platzek, P.S., Riedel, E.O., Kim, S.H., Lenhart, N., Ramschütz, C., Paprottka, K.J., Kertels, O.R., Möller, H.K., Atad, M., et al.: Generating synthetic high-resolution spinal stir and t1w images from t2w fse and low-resolution axial dixon. European Radiology **35**(4), 1761–1771 (2025)
6. Lei, P., Zhang, M., Fang, F., Zhang, G.: Robust deep convolutional dictionary model with alignment assistance for multi-contrast mri super-resolution. IEEE Transactions on Medical Imaging (2025)
7. McGinnis, J., Shit, S., Li, H.B., Sideri-Lampretsa, V., Graf, R., Dannecker, M., Pan, J., Stolt-Ansó, N., Mühlau, M., Kirschke, J.S., et al.: Single-subject multi-contrast mri super-resolution via implicit neural representations. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 173–183. Springer (2023)
8. McKay, M.J., Weber, K.A., Wesselink, E.O., Smith, Z.A., Abbott, R., Anderson, D.B., Ashton-James, C.E., Atyeo, J., Beach, A.J., Burns, J., et al.: Musclemap: An open-source, community-supported consortium for whole-body quantitative mri of muscle. Journal of imaging **10**(11), 262 (2024)

9. Schlereth, M., Schillinger, M., Breininger, K.: Faster, self-supervised super-resolution for anisotropic multi-view mri using a sparse coordinate loss. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 172–182. Springer (2025)
10. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. Advances in neural information processing systems **33**, 7537–7547 (2020)
11. Wang, X., Wang, S., Xiong, H., Xuan, K., Zhuang, Z., Liu, M., Shen, Z., Zhao, X., Zhang, L., Wang, Q.: Spatial attention-based implicit neural representation for arbitrary reduction of mri slice spacing. Medical Image Analysis **94**, 103158 (2024)
12. Wu, Q., Li, Y., Sun, Y., Zhou, Y., Wei, H., Yu, J., Zhang, Y.: An arbitrary scale super-resolution approach for 3d mr images via implicit neural representation. IEEE Journal of Biomedical and Health Informatics **27**(2), 1004–1015 (2022)
13. Xu, J., Moyer, D., Gagoski, B., Iglesias, J.E., Grant, P.E., Golland, P., Adalsteinsson, E.: Nesvor: implicit neural representation for slice-to-volume reconstruction in mri. IEEE transactions on medical imaging **42**(6), 1707–1719 (2023)
14. Zhang, Q., Chen, X., He, Z., Wu, L., Wang, K., Sun, J., Shen, H.: Pathology-guided ai system for accurate segmentation and diagnosis of cervical spondylosis. IEEE Journal of Biomedical and Health Informatics (2025)
15. Zhang, Y., Tian, Y., Kong, Y., Zhong, B., Fu, Y.: Residual dense network for image super-resolution. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2472–2481 (2018)
16. Zhao, C., Dewey, B.E., Pham, D.L., Calabresi, P.A., Reich, D.S., Prince, J.L.: Smore: a self-supervised anti-aliasing and super-resolution algorithm for mri using deep learning. IEEE transactions on medical imaging **40**(3), 805–817 (2020)
17. Zhong, J., Yang, B., Li, Y., Zhong, F., Chen, Z.: Image fusion and super-resolution with convolutional neural network. In: Chinese Conference on Pattern Recognition. pp. 78–88. Springer (2016)