



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Single-Subject Multi-View MRI Super-Resolution via Implicit Neural Representations

Heejong Kim^{1*}[†], Abhishek Thanki^{3*}, Roel van Herten¹, Daniel J. Margolis¹,
and Mert R. Sabuncu^{1,2}

¹ Department of Radiology, Weill Cornell Medicine, New York, USA

² School of Electrical and Computer Engineering, Cornell University
and Cornell Tech, New York, USA

³ Weill Cornell Graduate School of Medical Sciences, New York, USA

Abstract. Clinical MRI frequently acquires anisotropic volumes with high in-plane resolution and low through-plane resolution to reduce acquisition time. Multiple orientations are therefore acquired to provide complementary anatomical information. Conventional integration of these views relies on registration followed by interpolation, which can degrade fine structural details. Recent deep learning-based super-resolution (SR) approaches have demonstrated strong performance in enhancing single-view images. However, their clinical reliability is often limited by the need for large-scale training datasets, resulting in increased dependence on cohort-level priors. Self-supervised strategies offer an alternative by learning directly from the target scans. Prior work either neglects the existence of multi-view information or assumes that in-plane information can supervise through-plane reconstruction under the assumption of pre-alignment between images. However, this assumption is rarely satisfied in clinical settings. In this work, we introduce Single-Subject Implicit Multi-View Super-Resolution for MRI (SIMS-MRI), a framework that operates solely on anisotropic multi-view scans from a single patient without requiring pre- or post-processing. Our method combines a multi-resolution hash-encoded implicit representation with learned inter-view alignment to generate a spatially consistent isotropic reconstruction. We validate the SIMS-MRI pipeline on both simulated brain and clinical prostate MRI datasets. The code is publicly available for reproducibility[‡].

Keywords: Super-resolution · Implicit Neural Representation · MRI · Self-supervision

1 Introduction

In clinical practice, MRI acquisitions typically provide high in-plane resolution but lower through-plane resolution to reduce scan time, with multiple orientations acquired to offer complementary anatomical views. Integration of these

*Equal contribution.

[†]Corresponding author. Email: hek4004@med.cornell.edu

[‡]<https://github.com/abhshkt/SIMS-MRI>

views commonly relies on image registration followed by interpolation [27], which may blur fine structures, particularly along the low-resolution axis. To address this, super-resolution (SR) methods aim to reconstruct isotropic volumes. Early optimization-based approaches leveraged analytical priors such as low-rank and total variation regularization [24], but struggled under large anisotropy. More recently, deep learning methods trained on paired low- and high-resolution data have demonstrated strong quantitative performance [20,19,11]. However, their clinical applicability is often limited by domain shift, scanner variability, and the absence of standardized MRI intensity scales [21], as well as potential cohort-specific priors.

Starting from SSR [8], self-supervised super-resolution approaches aim to exploit high-resolution information inherently present in the acquired data to enhance the lower-resolution component without relying on external paired data or large-scale training datasets. ECLARE [22] models in-plane as the target distribution by leveraging ESPRESO [6] for slice excitation profile and FOV aware interpolation. Similarly, SMORE [35] demonstrates that high-resolution in-plane slices can supervise through-plane reconstruction. However, these methods primarily operate on a single-view anisotropic volume and overlook the fact that MRI data are often acquired in multi-view setups in many clinical domains [29,13]. SIMPLE [1] extends this to join multi-plane adversarial training with per-plane discriminators that use ATME-generated references, but requires a multi-subject training corpus.

More recently, implicit neural representations (INRs) have emerged as a continuous alternative to voxel-based models, providing greater flexibility without being constrained by a discrete voxel grid. By incorporating Fourier [26] and Sinusoidal encoding [25], INR based methods have shown promising performance in image reconstruction applications. In MRI, McGinnis et al. [16] demonstrated single-subject multi-contrast super-resolution using INRs with Fourier feature encoding, highlighting the promise of self-supervised continuous models for integrating complementary acquisitions from multi-modal MR images, and IREM [31] applied INR to multi-view MRI reconstruction. However, these frameworks assume pre-registered inputs and do not explicitly account for the inter-scan misalignments that are common in routine clinical acquisitions. NeSVoR [32] integrates an implicit volumetric representation with slice-to-volume reconstruction for fetal MRI but accounts only for slice-wise rigid motion.

In this work, we propose **Single-subject Implicit Multi-view Super-resolution for MRI (SIMS-MRI)**, a self-supervised framework that reconstructs isotropic volumes from multi-view anisotropic scans of a single subject without external training data or prior registration. The method follows a three-stage pipeline: (1) hash-encoded implicit image modeling of one view, (2) implicit deformable image registration (IDIR) [30] to align the second view, and (3) multi-view fusion to refine the implicit representation and reconstruct an isotropic volume. We validate the proposed approach on axial and coronal view images simulated

brain and clinical prostate MRI datasets, using an orthogonal sagittal view as ground truth for quantitative and qualitative evaluation.

2 Method

Our problem formulation considers two anisotropic MRI volumes of the same subject acquired from different views (e.g., axial and coronal view volumes), denoted as $I_{v1} : \Omega \rightarrow \mathbb{R}$ and $I_{v2} : \Omega \rightarrow \mathbb{R}$, where $\Omega \subset \mathbb{R}^3$ represents the spatial domain. Due to acquisition constraints, both volumes have high-resolution in-plane but thick slices along the through-plane directions. Our goal is to reconstruct an isotropic high-resolution volume $\hat{I} : \Omega \rightarrow \mathbb{R}$ by leveraging the complementary information available in both views without requiring additional acquisitions. We formulate this as learning a continuous implicit neural representation $f_\theta : \mathbb{R}^3 \rightarrow \mathbb{R}$ that maps a 3D coordinate $\mathbf{x} = (x, y, z)$ to an intensity value, enabling arbitrary-resolution sampling. To handle potential inter-view misalignment, our pipeline learns a deformable registration model $g_\phi : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ that maps coordinates from the one view to the coordinate space of the other.

2.1 SIMS-MRI

SIMS-MRI consists of three phases, where we train two INR models sequentially, each with a distinct objective. An image representation modeling, f_θ , is optimized for phases 1 and 3, and an image registration model, g_ϕ , for phase 2.

Phase 1: Initial Implicit Image Representation. The INR-based image representation model, f_θ , is a multi-layer perceptron (MLP) with ReLU activations mapping encoded coordinates to intensity values. The input coordinates are first transformed using a multi-resolution hash encoding before MLP to avoid spectral bias. Specifically, we adopt a coordinate-based INR inspired by Instant-NGP [18], where multi-resolution hash encoding replaces conventional positional encodings. Unlike fixed projection approaches such as Fourier feature mapping [26] and SIREN [25], the hash encoding learns adaptive multi-scale embeddings, enabling improved memory efficiency and faster convergence.

Phase 2: Implicit Image Registration. The f_θ is kept fixed while we train the registration model to align two views. Image registration model, g_ϕ , predicts a dense displacement field that maps coordinates from the one view to another view’s space. Following the IDIR [30], we use a SIREN network [25] that maps 3D coordinates to 3D displacements ($\delta = g_\phi(\mathbf{x})$), which are added to the input to obtain the transformed position: $\mathbf{x}' = \mathbf{x} + \delta$. The model was optimized using normalized cross-correlation (NCC) [12] as the similarity measure. NCC was used over MSE for registration loss because separately acquired MRI views may exhibit intensity-profile differences. For the simulated linear movement in two-view brain MRI experiments, we employed a bending energy penalty to encourage near-linear behavior.




Fig. 1. Overview of the proposed three-phase coordinate-based multi-view MRI fusion framework. **Phase 1:** A continuous implicit image representation model with multi-resolution hash encoding, f_{θ} , is trained on a reference (axial) view to map 3D world coordinates to voxel intensities. **Phase 2:** A SIREN-based registration network [30] g_{ϕ} learns a dense displacement field that aligns the second (coronal) view to the reference coordinate space while keeping f_{θ} fixed. **Phase 3:** The implicit image representation is fine-tuned using both the reference and registered views for information fusion. **Inference:** The learned continuous representation is queried on an isotropic 3D grid to generate a super-resolved volume, enabling reconstruction of unseen slice orientations.

Phase 3: Multi-view Fusion using Both Views. With the registration model frozen, the pre-trained image model in Phase 1 is trained on both axial and registered coronal views using the same optimizer as phase 1 to learn complementary information of both views.

Inference: Isotropic Reconstruction. To obtain the isotropic volume, we query a regular isotropic grid. The grid coordinates are normalized to $[-1, 1]$, consistent with the normalization applied during training. The final isotropic volume is then reconstructed using the original coordinate indices.

Implementation Details. The intrinsic image model consists of a ReLU-activated MLP with a 515-dimensional input layer, four hidden layers of 1024

units each, and a single-unit output layer. A multi-resolution hash encoder has $L = 16$ levels, $F = 32$ features per level, and a hash table size of $T = 2^{19}$ per level. Following IDIR [30], a SIREN network with three 256-unit hidden layers ($\omega_0 = 32$) is used for the image registration model [25]. In Phase 2, for the brain dataset, where inter-view differences were assumed to be predominantly rotational, we applied bending energy regularization with a high weighting factor ($\alpha = 1000$) to encourage near-linear transformations [23]. With this high weighting factor, the bending energy globally limits the curvature of the deformation field, resulting in a smooth, near-linear transformation. For the prostate dataset, where inter-view differences may include non-linear soft-tissue deformation, no bending energy regularization was applied ($\alpha = 0$). In our current implementation, per-subject optimization requires approximately 30 minutes on a single NVIDIA A100 GPU. The hash grid contains 268.4M trainable embeddings, while the decoder MLP and registration network contain 4.7M and 0.1M trainable parameters, respectively.

Input images are mapped to scanner coordinates using the affine transformation specified in the image header, and intensity values are jointly normalized to $[-1, 1]$. During training, all phases use batches of 50,000 voxel coordinates. Phases 1 and 3 are optimized using AdamW [14] (learning rate 1.2×10^{-3} , weight decay 5×10^{-5}) with a cosine annealing schedule, while Phase 2 employs Adam [9] with a learning rate of 10^{-5} following IDIR [30]. To mitigate overfitting to high-frequency noise, we adopt a progressive hash-level unlocking strategy following FreeNeRF [33], where hash levels are gradually activated during training to prioritize coarse structures before fine details.

3 Experiments

For both experiments, we compare SIMS-MRI against self-supervised super-resolution baselines, including SMORE [35] and ECLARE [22], both of which require image registration to integrate information from multiple views. After SMORE and ECLARE reconstruct view-specific isotropic volumes, image registration is applied to align the two views, which are then fused by averaging. In addition, classical interpolation methods, including B-spline and linear interpolation, are included for comparison under the assumption of perfectly aligned inputs. For linear registration, SimpleITK [15] was used, and non-linear registration was performed using FireANTs [7].

3.1 Simulated Two-View Brain MRI

The simulated dataset was generated by $4\times$ downsampling (e.g., 1 mm to 4 mm thickness) T1-weighted brain MRIs from 50 subjects in the OASIS-3 dataset [10]. For each subject, two anisotropic volumes were created: an axial volume with high in-plane and low through-plane resolution, and a coronal volume with high in-plane resolution in the coronal plane. To simulate inter-acquisition misalignment, the coronal volume was rotated by 0.1 rad about the x-axis prior to downsampling. For final isotropic volume of SIMS-MRI, we used isotropic volume grid

of ground truth for the inference. An orthogonal view (e.g., sagittal) extracted from the original isotropic volume is used as ground truth for validation. This experiment is designed to isolate anisotropic spatial sampling and inter-acquisition misalignment, while more detailed modeling of slice-selection profiles, slice gaps, and through-plane point-spread functions is left for future work.

We evaluate the proposed SIMS-MRI framework through comprehensive quantitative and qualitative experiments on a simulated two-view anisotropic 3D brain MRI dataset (e.g., axial and coronal views). While many recent deep learning approaches emphasize SSIM and PSNR as primary metrics for super-resolution, these measures do not capture perceptual realism [5]. For quantitative evaluation, we compare pixel-based metrics, including Mean Absolute Error (MAE), Structural Similarity Index Measure (SSIM) [28], and Peak Signal-to-Noise Ratio (PSNR), as well as perceptual metrics such as Learned Perceptual Image Patch Similarity (LPIPS) [34], Deep Image Structure and Texture Similarity (DISTS) [4], and a Medical Open Network for Artificial Intelligence model-based perceptual loss (MONAI) [2] for a more comprehensive assessment. For this perceptual loss, we employed a 3D MedicalNet ResNet-50 model [3], to quantify our simulated brain data experiments.

Table 1 reports quantitative results. SIMS-MRI achieves the best performance on five out of six metrics and the third-best on the DISTS, demonstrating strong improvements in both intensity-based and perceptual quality measures. Notably, on MAE, SSIM, and PSNR, SIMS-MRI outperforms competing methods by a substantial margin. These quantitative gains are reflected in the qualitative results in Fig. 2, where SIMS-MRI shows clearer structural delineation and improved intensity consistency compared to baseline methods. Baseline approaches exhibit noticeable grid artifacts (green arrow), whereas SIMS-MRI produces smoother reconstructions without such artifacts. Furthermore, fine anatomical structures, such as cortical gyri in the sagittal view, are better preserved by SIMS-MRI (yellow arrow), further validating its effectiveness for isotropic MRI reconstruction.

Beyond the comparison with SIMS-MRI, the table also highlights the benefit of multi-view fusion for the baseline methods. Even after image registration, which may introduce slight blurring, both ECLARE-fused and SMORE-fused outperform their single-view counterparts across all metrics, confirming that combining axial and coronal views improves reconstruction quality.

Table 2 presents the ablation results. Replacing hash encoding (HE) with Fourier mapping (FM) substantially degrades performance, underscoring the importance of multi-resolution hash encoding for capturing high-frequency anatomical details. On perceptual metrics (LPIPS, DISTS, and MONAI perceptual score), the SIMS-MRI achieves strong performance, indicating improved structural realism. Incorporating FreeNeRF further enhances distortion-based metrics (MAE, SSIM, and PSNR), improving voxel-wise fidelity while maintaining comparable perceptual quality.




Fig. 2. Qualitative results for the unseen view in simulated two-view brain MRI. The first row presents a comparison across methods, and the second row shows a zoomed-in view of the highlighted region. Yellow and green arrows indicate areas where SIMS-MRI achieves superior structural detail reconstruction.

Table 1. Quantitative comparison on simulated two-view brain MRI (unseen view evaluation). ‘axial’ denotes reconstruction using the axial view only, while ‘fused’ indicates co-registered and averaged axial/coronal volumes. (↑: higher is better and ↓: lower is better. Best in **bold** and second-best in underline.)

Method	MAE ↓	SSIM ↑	PSNR ↑	LPIPS ↓	DISTS ↓	MONAI ↓
Interpolation (B-spline)	0.0551	0.6723	18.8914	0.2842	0.1960	6.17×10^{-4}
Interpolation (Linear)	0.0541	0.6779	19.0907	0.2865	0.1991	7.13×10^{-4}
ECLARE-axial	0.0435	0.7616	20.5620	0.2687	0.2040	3.60×10^{-4}
ECLARE-fused	0.0412	0.7786	21.2464	0.2464	<u>0.1802</u>	<u>2.91×10^{-4}</u>
SMORE-axial	0.0431	0.7633	20.5709	0.2700	0.2037	3.57×10^{-4}
SMORE-fused	<u>0.0409</u>	<u>0.7816</u>	<u>21.2720</u>	<u>0.2445</u>	0.1787	3.10×10^{-4}
SIMS-MRI (ours)	0.0247	0.8877	24.6369	0.2362	0.1845	2.18×10^{-4}

Table 2. Ablation study on simulated two-view brain MRI. HE: Hash Encoding; FM: Fourier Mapping. Bold and underlined values indicate the best and second-best results.

Method	MAE ↓	SSIM ↑	PSNR ↑	LPIPS ↓	DISTS ↓	MONAI ↓
SIMS-MRI	0.0247	0.8877	24.6369	<u>0.2362</u>	<u>0.1845</u>	<u>2.18×10^{-4}</u>
w/o FreeNeRF & w/ HE	<u>0.0281</u>	<u>0.8461</u>	<u>23.6493</u>	0.2254	0.1744	1.64×10^{-4}
w/o FreeNeRF & w/ FM	0.0327	0.8077	23.3116	0.3577	0.2345	5.31×10^{-4}

3.2 Real-world Clinical Prostate MRI

For further evaluation, we assessed our method using routinely acquired clinical prostate MRI images from 50 patients at a single hospital, which typically are acquired in $0.5 \times 0.5 \text{ mm}^2$ in-plane with 3 mm slice thickness. The acquisition series included three in-plane acquisitions, resulting in axial, coronal, and sagittal volumes. We used the anisotropic axial and coronal volumes as inputs and the subsampled isotropic grid for the final isotropic inference grid. The sagittal volume was used as a held-out pseudo-reference for evaluation. Because it was

not explicitly co-registered with the reconstructed volumes, it provides only a weak reference for evaluating consistency in an unseen orientation rather than an exact voxel-wise ground truth. To avoid potential blurring artifacts that volume registration can introduce, we performed slice-wise comparison based on spatial proximity rather than voxel-level alignment. This procedure was applied consistently to all baseline methods as well as SIMS-MRI.

The MONAI perceptual metric, which is robust to minor spatial misalignment as other perceptual metrics and better reflects structural consistency in medical images, was computed on the sagittal images by matching each reconstructed slice to the closest corresponding slice in physical coordinate space. A 2D ResNet-50 model pretrained on RadImageNet [17] was used for prostate MRI experiments.




Fig. 3. Results for prostate MRI. The violin plots show the MONAI metric (MONAI: Medical Open Network for Artificial Intelligence model-based perceptual metric). Each row corresponds to a representative subject, selected from slices near the median of the distribution shown in the violin plots (color-coded and marked in the plots).

Figure 3 summarizes the qualitative and quantitative comparison between baseline methods and the proposed SIMS-MRI. Compared to baseline methods (a–d), SIMS-MRI (e) exhibits a clear downward shift in the distribution, indicating lower perceptual discrepancy in the unseen sagittal view. The highlighted representative subjects lie close to the median of each distribution, demonstrating that the qualitative examples shown on the left are representative rather than selectively chosen. In the first row, SIMS-MRI achieves improved structural similarity, with clearer anatomical boundaries and better preservation of spatial organization. The second row highlights improved textural consistency, where SIMS-MRI better maintains fine-grained intensity patterns. The third row demonstrates simultaneous improvements in both structural and textural fidelity.

4 Discussion and Conclusion

We presented SIMS-MRI, a single-subject, self-supervised implicit framework for multi-view MRI super-resolution that jointly models continuous image representation and inter-scan misalignment without external training data or prior registration. By combining multi-resolution hash encoding with implicit deformable alignment and multi-view refinement, the method integrates complementary anisotropic acquisitions into a spatially consistent isotropic reconstruction, reducing sensitivity to domain shift, scanner variability, and cohort-level priors. While the per-subject optimization is more computationally demanding than feed-forward approaches and clinical evaluation was limited to slice-level correspondence, SIMS-MRI consistently improved quantitative performance in simulated data and perceptual quality in real clinical prostate MRI. Future work will extend this framework to additional views and contrasts, evaluate robustness across different anisotropy ratios, and improve the efficiency of the pipeline for practical deployment. Overall, we believe SIMS-MRI represents a meaningful step toward data-efficient isotropic MRI reconstruction from complementary anisotropic MRI acquisitions.

Acknowledgments. Funding for this project was in part provided by the NIH grants R01AG053949, R01AG064027, R01AG070988, R01HL174863, U54DK144866, K25CA283145, and the NSF CAREER 1748377 grant. Parts of Figure 1 were designed using icons from Flaticon.com.

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

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. Cardoso, M.J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Murrey, B., Myronenko, A., Zhao, C., Yang, D., et al.: Monai: An open-source framework for deep learning in healthcare. arXiv preprint arXiv:2211.02701 (2022)
3. Chen, S., Ma, K., Zheng, Y.: Med3d: Transfer learning for 3d medical image analysis. arXiv preprint arXiv:1904.00625 (2019)
4. Ding, K., Ma, K., Wang, S., Simoncelli, E.P.: Image quality assessment: Unifying structure and texture similarity. arXiv preprint arXiv:2004.07728 (2020), <https://arxiv.org/abs/2004.07728>
5. Dohmen, M., Klemens, M.A., Baltruschat, I.M., Truong, T., Lenga, M.: Similarity and quality metrics for mr image-to-image translation. Scientific Reports **15**(1), 3853 (2025)
6. Han, S., Remedios, S.W., Schär, M., Carass, A., Prince, J.L.: Espresso: An algorithm to estimate the slice profile of a single magnetic resonance image. Magnetic Resonance Imaging **98**, 155–163 (2023)

7. Jena, R., Chaudhari, P., Gee, J.C.: Fireants: Adaptive riemannian optimization for multi-scale diffeomorphic matching. arXiv preprint arXiv:2404.01249 (2024)
8. Jog, A., Carass, A., Prince, J.L.: Self super-resolution for magnetic resonance images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 553–560. Springer (2016)
9. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
10. LaMontagne, P.J., Benzinger, T.L., Morris, J.C., Keefe, S., Hornbeck, R., Xiong, C., Grant, E., Hassenstab, J., Moulder, K., Vlassenko, A.G., et al.: Oasis-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease. medrxiv pp. 2019–12 (2019)
11. de Leeuw Den Bouter, M., Ippolito, G., O’Reilly, T., Remis, R., Van Gijzen, M., Webb, A.: Deep learning-based single image super-resolution for low-field mr brain images. Scientific Reports **12**(1), 6362 (2022)
12. Lewis, J.P.: Fast normalized cross-correlation. In: Vision interface. vol. 10, pp. 120–123 (1995)
13. Liberman, L., Menell, J.H.: Breast imaging reporting and data system (bi-rads). Radiologic Clinics **40**(3), 409–430 (2002)
14. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
15. Lowekamp, B.C., Chen, D.T., Ibáñez, L., Blezek, D.: The design of simpleitk. Frontiers in neuroinformatics **7**, 45 (2013)
16. 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)
17. Mei, X., Liu, Z., Robson, P.M., Marinelli, B., Huang, M., Doshi, A., Jacobi, A., Cao, C., Link, K.E., Yang, T., et al.: Radimagenet: an open radiologic deep learning research dataset for effective transfer learning. Radiology: Artificial Intelligence **4**(5), e210315 (2022)
18. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM transactions on graphics (TOG) **41**(4), 1–15 (2022)
19. Oktay, O., Bai, W., Lee, M., Guerrero, R., Kamnitsas, K., Caballero, J., de Marvao, A., Cook, S., O’Regan, D., Rueckert, D.: Multi-input cardiac image super-resolution using convolutional neural networks. In: International conference on medical image computing and computer-assisted intervention. pp. 246–254. Springer (2016)
20. Pham, C.H., Tor-Díez, C., Meunier, H., Bednarek, N., Fablet, R., Passat, N., Rousseau, F.: Multiscale brain mri super-resolution using deep 3d convolutional networks. Computerized Medical Imaging and Graphics **77**, 101647 (2019)
21. Remedios, S.W., Wei, S., Dewey, B.E., Carass, A., Pham, D.L., Prince, J.L.: Pushing the limits of zero-shot self-supervised super-resolution of anisotropic mr images. In: Medical Imaging 2024: Image Processing. vol. 12926, pp. 29–35. SPIE (2024)
22. Remedios, S.W., Wei, S., Han, S., Zhang, J., Carass, A., Schilling, K.G., Pham, D.L., Prince, J.L., Dewey, B.E.: Eclare: efficient cross-planar learning for anisotropic resolution enhancement. arXiv preprint arXiv:2503.11787 (2025)
23. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J.: Non-rigid registration using free-form deformations: application to breast mr images. IEEE transactions on medical imaging **18**(8), 712–721 (1999)

24. Shi, F., Cheng, J., Wang, L., Yap, P.T., Shen, D.: Lrtv: Mr image super-resolution with low-rank and total variation regularizations. *IEEE transactions on medical imaging* **34**(12), 2459–2466 (2015)
25. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. *Advances in neural information processing systems* **33**, 7462–7473 (2020)
26. 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)
27. Tsai, R.Y., Huang, T.S.: Multiframe image restoration and registration. *Multiframe image restoration and registration* **1**, 317–339 (1984)
28. 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)
29. Weinreb, J.C., Barentsz, J.O., Choyke, P.L., Cornud, F., Haider, M.A., Macura, K.J., Margolis, D., Schnall, M.D., Shtern, F., Tempany, C.M., et al.: Pi-rads prostate imaging-reporting and data system: 2015, version 2. *European urology* **69**(1), 16–40 (2016)
30. Wolterink, J.M., Zwienenberg, J.C., Brune, C.: Implicit neural representations for deformable image registration. In: *International Conference on medical imaging with deep learning*. pp. 1349–1359. PMLR (2022)
31. Wu, Q., Li, Y., Xu, L., Feng, R., Wei, H., Yang, Q., Yu, B., Liu, X., Yu, J., Zhang, Y.: Irem: High-resolution magnetic resonance image reconstruction via implicit neural representation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 65–74. Springer (2021)
32. 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). <https://doi.org/10.1109/TMI.2023.3236216>
33. Yang, J., Pavone, M., Wang, Y.: Freenerf: Improving few-shot neural rendering with free frequency regularization. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 8254–8263 (2023)
34. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 586–595 (2018)
35. 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)