



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.

Shared Gaussian Geometry for Brain MRI Spatial Resolution Harmonization Without External Training

Yifan Gao^{1,2}, Peiran Xu^{1,3}, Yimeng He^{1,4}, Haoran Li^{1,3}, Ziyang Long^{1,3}, and Debiao Li^{1,3}

¹ Biomedical Imaging Research Institute, Cedars-Sinai Medical Center, Los Angeles, CA, USA

{Yifan.Gao, Peiran.Xu, Yimeng.He, Haoran.Li, Ziyang.Long}@cshs.org, {Debiao.Li}@csmc.edu

² Tsinghua Medicine, Tsinghua University, Beijing, China

³ Department of Bioengineering, University of California-Los Angeles, Los Angeles, CA, USA

⁴ Department of Computational Medicine, Cedars-Sinai Medical Center, Los Angeles, CA, USA

Abstract. Clinical magnetic resonance imaging (MRI) protocols are spatially heterogeneous across modalities and clinical objectives, compromising multimodal joint analysis and arbitrary-view generation. To harmonize MRI resolution, we leverage the commonly acquired isotropic T1-weighted sequence in neuroimaging protocols. Our Gaussian Splatting (GS)-based shared geometry framework adopts a two-stage training strategy, in which an explicit, subject-specific Gaussian scaffold encoding anatomical geometry is first learned from the isotropic structural scan and then reused to fit appearance for target modalities acquired with sparse slices. Experiments on the UK Biobank, GBM, and ABCD datasets for through-plane super-resolution across multiple modalities (T2-weighted, FLAIR, DWI, ASL), degradation factors ($\times 3$, $\times 5$, $\times 7$), and pathological abnormalities (glioblastoma) demonstrate state-of-the-art reconstruction fidelity. The shared Gaussian geometry enables arbitrary-view generation for target modalities with strong structural consistency and further shows potential for self-supervised in-plane super-resolution. This work establishes explicit geometry-guided representations as a novel, flexible, and interpretable pathway toward retrospective multi-contrast MRI harmonization and reliable clinical reference construction. The GitHub code is available at: <https://github.com/yfgao76/AtlasGS>.

Keywords: Super-resolution · Multi-contrast MRI · Gaussian Splatting.

1 Introduction

Magnetic resonance imaging (MRI) spatial harmonization refers to establishing a consistent anatomical representation across heterogeneous acquisition pro-

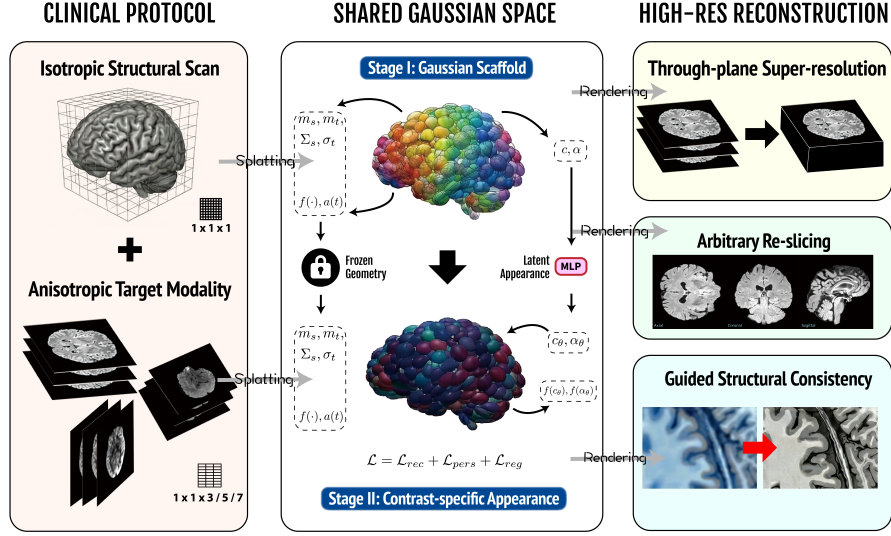


Fig. 1. Overview of the proposed shared Gaussian geometry framework for zero-shot multi-contrast MRI harmonization.

protocols, enabling reliable cross-contrast visualization, comparison, and measurement [9,1,27]. In this paper, we focus on clinical multi-contrast neuroimaging, where protocols comprise a high-resolution isotropic structural scan together with multiple disease-targeted sequences acquired using anisotropic slice thickness and modality-specific resolutions, trajectories, and acquisition parameters [17]. This variability results in inconsistent geometric representations across contrasts and limits coherent arbitrary-view reconstruction, particularly in retrospective datasets where acquisition settings are pre-arranged and hard to simulate precisely [8], while also indicates the potential of utilizing complementary contrast-specific information [16,21].

Conventional MRI protocol harmonization approaches face substantial challenges under diverse imaging modalities and sparse through-plane supervision [13]. Interpolation produces over-smoothed images and does not align geometric representations across contrasts. Deep learning-based super-resolution models rely on external training data and often struggle with clinical datasets that exhibit wide variability and individualized harmonization requirements [6]. Implicit Neural Representations (INR) and Gaussian Splatting (GS) have recently gained attention due to their zero-shot continuous coordinate modeling and resolution-free rendering capabilities [20,12]. However, these methods do not inherently encode semantic organ geometry, and their potential in bridging multiple contrasts remains under-explored.

Building upon GS, we propose a **shared Gaussian geometry representation** that learns an anatomical prior while accommodating modality-specific

signal formation. Following the MedGS framework [18], we apply video Gaussian splatting [23] to an isotropic structural scan to construct an explicit Gaussian scaffold. Contrast-specific appearance parameters are then fitted on this fixed scaffold to enable zero-shot through-plane arbitrary super-resolution across different modalities and resolutions. To enhance structural stability, we incorporate topology-aware Gaussian densification [22] and latent appearance regularization, and integrate low-resolution modality consistency to suppress noise and refine contrast fidelity (see Fig. 1).

Experiments on UK Biobank (healthy older cohorts), GBM (glioblastoma cohorts), and ABCD (healthy younger cohorts) across multiple modalities (T2-weighted, T2-FLAIR, DWI, ASL) and through-plane super-resolution factors ($\times 3, 5, 7$) demonstrate state-of-the-art reconstruction fidelity and cross-modality geometric consistency compared to subject-wise interpolation and contrast-agnostic INR/GS baselines, especially in clinical settings where slice thickness often exceeds 5 mm. We further show that our framework mitigates tumor hallucinations commonly observed in interpolation-based reconstructions, and that the learned Gaussian scaffold can also support in-plane super-resolution. These findings support our hypothesis that explicit geometry-guided representations provide a practical pathway toward retrospective multi-contrast MRI harmonization and the construction of reliable clinical reference spaces.

2 Methods

Let $I^{T1} \in \mathbb{R}^{H \times W \times D}$ denote an isotropic structural MRI (denoted as T1 in this section), and let I_{LR}^m denote a target modality acquired with thick through-plane sampling (in low resolution). Our objective is to learn a continuous high resolution representation

$$I_{HR}^m(s, t), \quad s = (x, y) \in \mathbb{R}^2, \quad t \in \mathbb{R}, \quad (1)$$

that:

- (i) preserves structural geometry defined by I^{T1} ,
- (ii) respects target acquisition physics, and
- (iii) enables arbitrary continuous re-slicing.

2.1 2D+t Gaussian Geometry with Latent Opacity and Appearance

We utilize a shared 2D+t Gaussian field from [23]

$$\mathcal{G} = \{\mathcal{G}_i\}_{i=1}^N, \quad (2)$$

where each primitive $\mathcal{G}_i$ represents a localized anatomical element and consists of:

- **geometric parameters:** spatial center $m_{s,i} \in \mathbb{R}^2$, through-plane center $m_{t,i} \in \mathbb{R}$, spatial covariance $\Sigma_{s,i} \in \mathbb{R}^{2 \times 2}$, through-plane variance $\sigma_{t,i}^2$,

- **modality-specific parameters:** opacity α_i , appearance c_i ,
- **deformation modeling parameters:** deformation function $f_i(\cdot)$, variance modulation function $a_i(t)$.

Starting from a separable Gaussian in (s, t) , we allow the spatial center to vary smoothly with t , thus the Folded-Gaussian density is defined as

$$FN_i(s, t) = \mathcal{N}(m_{s,i} + f_i(m_{t,i} - t), a_i(t)\Sigma_{s,i}) \cdot \mathcal{N}(m_{t,i}, \sigma_{t,i}^2). \quad (3)$$

This formulation models nonlinear anatomical continuity along the through-plane axis.

Conditioning of $\mathcal{G}$ at fixed t produces a standard 2D Gaussian in (x, y) , and rendering is formulated as:

$$I(s, t) = \sum_{i=1}^N \alpha_i FN_i(s, t) c_i. \quad (4)$$

2.2 Training Objectives

Training of $\mathcal{G}$ is performed in two stages:

- **Stage I (Structural learning):** to learn geometric parameters from I^{T1} .
- **Stage II (Modality fitting):** to fit modality-specific appearance using I_{LR}^m with frozen geometry.

In **Stage I**, all parameters are optimized towards I^{T1} , while in **Stage II**, only (α_i, c_i) are optimized.

Reconstruction. Reconstruction loss is formulated as

$$\mathcal{L}_{rec} = (1 - \lambda_{ssim}) \|\hat{I} - I_{gt}\|_1 + \lambda_{ssim} (1 - SSIM). \quad (5)$$

In **Stage I**, $\hat{I}$ is exactly I . While in **Stage II**, thick-slice supervision is modeled via slab integration along the through-plane coordinate t . For a slice centered at t_0 with thickness Δt , the predicted low-resolution slice is obtained by integrating the continuous Gaussian field within the slab:

$$\hat{I}(s, t_0) = \frac{1}{K} \sum_{k=1}^K I(s, t_0 + \delta_k), \quad \delta_k \in \left[-\frac{\Delta t}{2}, \frac{\Delta t}{2}\right], \quad (6)$$

where $\{\delta_k\}_{k=1}^K$ are uniformly spaced offsets inside the slice thickness.

Topology preservation. To suppress structural artifacts and preserve anatomical integrity, following Topology-GS [22], we impose a Persistence Homology (PH)-based constraint [7] $\mathcal{L}_{pers}$.

Latent appearance fitting. To ensure stable cross-modality fitting, appearance is factorized through a low-dimensional latent code:

$$c_i = f_\phi(\theta_i), \quad (7)$$

where θ_i is a compact per-Gaussian embedding and f_ϕ is a shared decoder.

Spatial smoothness is enforced by a kNN graph over Gaussian centers [3] and parameter regularization are applied:

$$\mathcal{L}_{reg} = \lambda_{smooth} \sum_{(i,j)} \|\theta_i - \theta_j\|_2^2 + \lambda_\theta \|\theta\|_2^2. \quad (8)$$

Total Loss.

$$\mathcal{L} = \mathcal{L}_{rec} + \lambda_{pers} \mathcal{L}_{pers} + \mathcal{L}_{reg}. \quad (9)$$

Two stages share the same loss. Stage I optimizes on I^{T1} , and Stage II optimizes on the target modality.

2.3 Low-Resolution (LR) Consistent Fusion

Structure-guided reconstruction may introduce anatomical artifacts under sparse thick-slice supervision. To balance structural prior and signal fidelity, we apply a LR consistent fusion at final rendering.

For high-resolution I_{HR}^{T1} and I_{HR}^m trained respectively from T1 and separately from target modality, and let $\mathcal{D}(\cdot)$ be the slab-aware degradation operator (blur + average pooling along t), we compute a voxel-wise LR confidence map via

$$w_{LR} = \text{softmax}_{\{T1, single\}} \left(-\frac{1}{\tau} [|\mathcal{D}(I_{HR}^{T1}) - I_{LR}^m|, |\mathcal{D}(I_{HR}^m) - I_{LR}^m|] \right)_{T1}, \quad (10)$$

where τ controls sensitivity and $(\cdot)_{T1}$ selects the T1-guided channel. After smoothing and clipping, w_{LR} is resampled as w and used to fuse:

$$I_{fused} = w I^{T1} + (1 - w) I^{single}. \quad (11)$$

3 Experiments and Results

Datasets We adopt multi-contrast through-plane super-resolution as an evaluation task of our representation. Experiments were done on three datasets: UK Biobank (UKBB) [24], UPenn-GBM (GBM) [2] and The Adolescent Brain Cognitive Development (ABCD) study [4]. ABCD was curated from FOMO300K dataset [5]. 4 target MRI sequences besides isotropic T1-MPRAGE are included, two of the them (T2-TSE and FLAIR) with consistent spacing, while the other 2 (diffusion-weighted imaging, DWI, and arterial spin labeling, ASL) have irregular slice spacing. UKBB and ABCD include respectively elder and younger healthy cohort, while GBM additionally includes patients with glioblastoma, where tumors alter brain anatomy (see Table 1 for details). All volumes underwent skull

Table 1. Dataset configuration, target modality resolution (in mm) and shape, and degradation factors.

Dataset	Train	Test	T1 Res.	T1 Shape	Target	Target Res.	Degrade Ratio
UKBB	100	100	1.0×1.0×1.0	(172,228,202)	Flair	1.0×1.0×1.0	×3 / ×5 / ×7
GBM	60	59	1.0×1.0×1.0	(240,240,140)	T2w	1.0×1.0×1.0	×7
					Flair	1.0×1.0×1.0	×7
ABCD	34	33	0.47×0.47×0.8	(512,512,256)	ASL	1.875×1.875×3.5	×3
					DWI	0.86×0.86×2.2	×3

stripping using HD-BET [10]. Low-resolution inputs are synthetically degraded along the slice direction with Gaussian blur averaging neighboring slices. Each method restores degraded frames from the input. We **don’t intentionally normalize** in order to evaluate our method’s ability to preserve absolute signal values, which is essential for downstream extension to quantitative MRI modalities (e.g. T1/T2 mapping).

Implementation For each subject, training proceeds in three stages: (1) Train T1 Gaussian scaffold; (2) Train single-modality Gaussian model on target modality; (3) Train and fit T1-guided model with frozen geometry on target modality. Each stages run for 10k iterations. Training takes an average of 20 minutes for one subject on an NVIDIA H100 GPU with 80 GB memory.

Optimization uses Adam [14]. Feature learning rate is 2.5×10^{-3} , opacity 2.5×10^{-2} , latent appearance 2.5×10^{-3} , and decoder 1×10^{-3} . Latent dimension is 4. Graph smoothness uses $k = 8$ neighbors. For loss regularization, $\lambda_{sim} = 0.2$, $\lambda_{pers} = 10$, $\lambda_{smooth} = 0.02$, $\lambda_{\theta} = 0.001$. See more hyper-parameters settings in GitHub repository.

Comparative Experiments We compared with through-plane super-resolution methods, include two interpolation methods: slab-wise linear interpolation (Linear), cubic B-spline interpolation (Cubic) [15]; two single-modality INR/GS based methods: MedGS [18], SA-INR [26]; and two multi-modal methods: T1-conditioned ALPINE-A2 (ALPINE) [25], Multi-contrast INR (MC-INR) [19]. SA-INR requires separate training set, while other methods don’t.

Quantitative Results In MAE, PSNR and SSIM, Our method achieves the best performance on UKBB across all degradation factors, with a widening margin under severe downsampling (e.g., ×7, typical in clinical protocols), indicating improved robustness under sparse thick-slice supervision (Table 2). Tumor preservation in the GBM cohort was assessed using nn-UNet Dice scores [11], where our model maintains stable lesion structure while modality-agnostic baselines often distort pathology and may even underperform linear interpolation. The method also generalizes to rare modalities with non-integral scaling ratios in the ABCD dataset (Table 3). Ablation Study shown in Table 4 supported

Table 2. Quantitative comparison results on UK Biobank (Flair) for through-plane super-resolution at $\times 3$, $\times 5$, and $\times 7$. Metrics are calculated inside brain mask. **Bold** indicates best results. Underline indicates second-best results. Numbers in brackets mean half-width Confidence Interval (CI).

Method	UKBB $\times 3$			UKBB $\times 5$			UKBB $\times 7$		
	MAE $\downarrow$	SSIM $\uparrow$	PSNR $\uparrow$	MAE $\downarrow$	SSIM $\uparrow$	PSNR $\uparrow$	MAE $\downarrow$	SSIM $\uparrow$	PSNR $\uparrow$
Interp	48.60 (1.33)	0.8697 (0.0014)	25.11 (0.20)	65.94 (1.81)	0.7866 (0.0021)	22.64 (0.20)	77.09 (2.26)	0.7373 (0.0026)	21.38 (0.20)
Cubic	46.78 (1.28)	<u>0.8793</u> (0.0013)	25.35 (0.20)	65.22 (1.78)	0.7919 (0.0021)	22.66 (0.20)	76.71 (2.25)	<u>0.7394</u> (0.0025)	21.34 (0.20)
MC-INR	53.44 (1.48)	0.8729 (0.0056)	24.69 (0.28)	72.92 (2.02)	0.7857 (0.0058)	22.07 (0.23)	86.25 (2.44)	0.7284 (0.0068)	20.71 (0.25)
SA-INR	46.95 (1.27)	0.8758 (0.0013)	25.17 (0.20)	67.23 (1.86)	0.7703 (0.0026)	22.13 (0.20)	80.15 (2.35)	0.7037 (0.0031)	20.69 (0.21)
ALPINE	239.35 (7.53)	0.4652 (0.0071)	13.17 (0.24)	236.06 (7.49)	0.4663 (0.0073)	13.25 (0.24)	232.46 (7.72)	0.4662 (0.0077)	13.31 (0.25)
MedGS	<u>46.18</u> (1.47)	0.8783 (0.0032)	<u>25.68</u> (0.24)	<u>62.57</u> (2.08)	<u>0.7924</u> (0.0053)	<u>23.17</u> (0.25)	<u>72.79</u> (2.47)	0.7362 (0.0055)	<u>21.91</u> (0.25)
Ours	40.16 (1.14)	0.9047 (0.0022)	27.01 (0.21)	52.21 (1.56)	0.8464 (0.0041)	24.90 (0.22)	60.90 (2.01)	0.8004 (0.0059)	23.69 (0.25)

Table 3. Quantitative Results on GBM (T2w, Flair) and ABCD (DWI, ASL). Metrics are calculated inside **Tumor Mask** for GBM and brain mask for ABCD. **DSC**: Dice score between groundtruth and predicted tumor mask using nn-UNet on reconstructions. **Bold** indicates best results. Underline indicates second-best results.

Method	GBM T2w $\times 7$			GBM Flair $\times 7$		ABCD DWI $\times 3$		ABCD ASL $\times 3$	
	MAE $\downarrow$	SSIM $\uparrow$	PSNR $\uparrow$	MAE $\downarrow$	DSC $\uparrow$	MAE $\downarrow$	SSIM $\uparrow$	MAE $\downarrow$	SSIM $\uparrow$
Interp	149.15	0.8273	20.26	44.12	0.6572	9.65	0.9792	7.11	0.9802
Cubic	<u>107.42</u>	0.9055	21.63	<u>29.27</u>	0.7247	5.78	0.9863	<u>4.63</u>	<u>0.9884</u>
SA-INR	108.97	<u>0.9075</u>	21.35	29.37	0.7243	<u>5.85</u>	<u>0.9863</u>	4.74	0.9877
MedGS	109.39	0.8782	<u>21.79</u>	31.38	0.7714	8.07	0.9824	4.89	0.9869
Ours	86.75	0.9133	23.63	27.36	<u>0.7458</u>	7.22	0.9859	4.60	0.9887

the value of (1) topology perservation; (2) latent appearance fitting; (3) LR-consistent fusion.

Qualitative Results High-resolution reconstructions on axial, coronal and sagittal view proved our model managed to learn a continous high-resolution representation (Fig. 2). Visualization on GBM dataset showed clear tumor boundary and eliminated halluciation introduced via interpolation, stressing clinical feasibility (Fig. 3 left). We also applied the Gaussian scaffold on T1 spacing and got ultra-high-resolution DWI and ASL scans in the ABCD dataset (Fig. 3 right).

4 Discussion and Conclusion

Our shared Gaussian framework achieves state-of-the-art through-plane reconstruction across multiple modalities and degradation levels without external

Table 4. Ablation Study Results. **T1-guided**: 2-stage fitting with free appearance fitting and only reconstruction loss; **+Topology**: T1-guided + Persist Loss; **+Latent**: T1-guided + Latent Appearance Fitting; **+LR-Fusion**: T1-guided + LR Consistent Fusion; **+All**: With all innovations.

Method	UKBB $\times 3$			UKBB $\times 5$			UKBB $\times 7$		
	MAE $\downarrow$	SSIM $\uparrow$	PSNR $\uparrow$	MAE $\downarrow$	SSIM $\uparrow$	PSNR $\uparrow$	MAE $\downarrow$	SSIM $\uparrow$	PSNR $\uparrow$
T1-guided	50.74	0.8646	24.87	63.12	0.8056	23.09	72.80	0.7565	21.98
+Topology	50.74	0.8646	24.87	63.06	0.8057	23.10	72.74	0.7566	21.99
+Latent	50.25	0.8687	24.92	62.71	0.8092	23.13	75.71	0.7585	21.92
+LR-Fusion	40.50	0.9028	26.96	52.53	0.8443	24.86	61.48	0.7972	23.62
+All	40.16	0.9047	27.01	52.21	0.8464	24.90	61.19	0.7998	23.66




Fig. 2. Axial, Coronal and Sagittal view of reconstructions results with comparative methods. 7x axial super-resolution from UKBB Flair. From left to right: Groundtruth, Linear, Cubic, MC-INR, SA-INR, MedGS, Ours.

training. The explicit shared geometric space restores in-plane structural detail and enables zero-shot adaptation of volume appearance.

Although geometry preservation is stable under anisotropic sampling, the framework has not been evaluated under motion corruption or anatomical deformation. Also the current formulation has not yet been extended to population-level atlas modeling. Expanding the shared scaffold toward atlas construction and large-scale harmonization is a natural next step.

Our geometry-aware design highlights both the value of structural priors as a lightweight yet powerful anchor, and the potential of explicit representation in medical image reconstruction. This paradigm enables both prospective protocol design and retrospective dataset reconciliation. Thanks to the flexibility and interpretability of gaussians, such explicit geometry-guided representations may provide a scalable foundation for cross-modal registration and mapping, longitudinal analysis, and image-guided intervention beyond super-resolution.




Fig. 3. Reconstruction results on GBM and ABCD datasets. GBM: green contour shows tumor regions, and error maps are given for reference. ABCD: UltraHigh denotes reconstructions generated at a sampling density matched to isotropic T1 resolution.

Acknowledgments. The authors acknowledge the computational resources and support provided by Cedars-Sinai Medical Center. The authors thank Yufeng Wang and Dr. JuDong Yang for providing support in access to UKBiobank brain imaging data.

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

References

1. Abbasi, S., et al.: Deep learning for the harmonization of structural MRI scans: a survey. *Biomed. Eng. Online* **23**, 90 (2024). <https://doi.org/10.1186/s12938-024-01280-6>
2. Bakas, S., et al.: The university of pennsylvania glioblastoma (UPenn-GBM) cohort: advanced MRI, clinical, genomics, & radiomics. *Scientific Data* **9**, 453 (2022). <https://doi.org/10.1038/s41597-022-01560-7>
3. Belkin, M., Niyogi, P.: Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.* **15**(6), 1373–1396 (2003). <https://doi.org/10.1162/089976603321780317>
4. Casey, B., et al.: The adolescent brain cognitive development (ABCD) study: Imaging acquisition across 21 sites. *Dev. Cogn. Neurosci.* **32**, 43–54 (2018). <https://doi.org/10.1016/j.dcn.2018.03.001>
5. Cerri, S., et al.: A large-scale heterogeneous 3D magnetic resonance brain imaging dataset for self-supervised learning. *arXiv:2506.14432* (2026), <https://arxiv.org/abs/2506.14432>
6. Choi, Y., et al.: TESLA: Test-time reference-free through-plane super-resolution for multi-contrast brain MRI. In: *MICCAI 2025. LNCS*, vol. 15972, pp. 584–593. Springer Nature Switzerland (2025). https://doi.org/10.1007/978-3-032-05169-1_56

7. Clough, J., et al.: A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE Trans. Pattern Anal. Mach. Intell.* **44**(12), 8766–8778 (2022). <https://doi.org/10.1109/TPAMI.2020.3013679>
8. Greenspan, H., Peled, S., Oz, G., Kiryati, N.: MRI inter-slice reconstruction using super-resolution. *Magn. Reson. Imaging* **20**, 437–446 (2002). [https://doi.org/10.1016/S0730-725X\(02\)00511-8](https://doi.org/10.1016/S0730-725X(02)00511-8)
9. Hu, F., et al.: Image harmonization: a review of statistical and deep learning methods for removing batch effects and evaluation metrics for effective harmonization. *NeuroImage* **274**, 120125 (2023). <https://doi.org/10.1016/j.neuroimage.2023.120125>
10. Isensee, F., et al.: Automated brain extraction of multisequence MRI using artificial neural networks. *Hum. Brain Mapp.* **40**(17), 4952–4964 (2019). <https://doi.org/10.1002/hbm.24750>
11. Isensee, F., et al.: nnu-net: self-configuring biomedical image segmentation. *Nat. Methods* **18**(2), 203–211 (2021). <https://doi.org/10.1038/s41592-020-01008-z>
12. Kerbl, B., et al.: 3D gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph. (SIGGRAPH)* (2023). <https://doi.org/10.1145/3592433>
13. Khateri, M., et al.: MRI super-resolution with deep learning: A comprehensive survey. *arXiv:2511.16854* (2025). <https://doi.org/10.48550/arXiv.2511.16854>
14. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. *arXiv:1412.6980* (2014). <https://doi.org/10.48550/arXiv.1412.6980>
15. Lehmann, T., Gönner, C., Spitzer, K.: Addendum: B-spline interpolation in medical image processing. *IEEE Trans. Med. Imaging* **20**(7), 660–665 (2001). <https://doi.org/10.1109/42.932749>
16. Liu, X., et al.: Deep unregistered multi-contrast MRI reconstruction. *Magn. Reson. Imaging* **81**, 33–41 (2021). <https://doi.org/10.1016/j.mri.2021.05.005>
17. Mahmoudzadeh, A., Kashou, N.: Interpolation-based super-resolution reconstruction: effects of slice thickness. *J. Med. Imaging (Bellingham)* **1**(3), 034007 (2014). <https://doi.org/10.1117/1.JMI.1.3.034007>
18. Marzol, K., et al.: MedGS: Gaussian splatting for multi-modal 3D medical imaging. *arXiv:2509.16806* (2025). <https://doi.org/10.48550/arXiv.2509.16806>
19. McGinnis, J., et al.: Single-subject multi-contrast MRI super-resolution via implicit neural representations. In: Greenspan, H., et al. (eds.) *MICCAI 2023*. LNCS, vol. 14229, pp. 173–183. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-43993-3_17
20. Molaei, A., et al.: Implicit neural representation in medical imaging: A comparative survey. In: *ICCVW*. pp. 2381–2391 (2023). <https://doi.org/10.1109/ICCVW60793.2023.00252>
21. Remedios, S., et al.: Self-supervised super-resolution for anisotropic MR images with and without slice gap. In: Wolterink, J., et al. (eds.) *SASHIMI 2023*. LNCS, vol. 14288, pp. 118–128. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-44689-4_12
22. Shen, T., Liu, S., Feng, J., Ma, Z., An, N.: Topology-aware 3D gaussian splatting: Leveraging persistent homology for optimized structural integrity. *arXiv:2412.16619* (2024). <https://doi.org/10.48550/arXiv.2412.16619>
23. Smolak-Dyżewska, W., et al.: VeGaS: Video gaussian splatting. *arXiv:2411.11024* (2024). <https://doi.org/10.48550/arXiv.2411.11024>
24. Sudlow, C., et al.: UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Med.* **12**(3), e1001779 (2015). <https://doi.org/10.1371/journal.pmed.1001779>

25. Vyas, K., et al.: Alpine: distributed PyTorch library for implicit neural representations. In: CVPR Workshop (2025), <https://github.com/kushalvyas/alpine>, last accessed 2026/02/25
26. Wang, X., et al.: SA-INR for arbitrary MRI slice-spacing reduction. *Med. Image Anal.* **94**, 103158 (2024). <https://doi.org/10.1016/j.media.2024.103158>
27. Yang, Q., Shomal-Zadeh, F., Gholipour, A.: MRI harmonization: a survey of acquisition, image-level, and feature-level methods. *arXiv:2507.16962* (2025). <https://doi.org/10.48550/arXiv.2507.16962>