



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.

Subject-Specific Low-Field MRI Synthesis via a Neural Operator

Ziqi Gao¹, Nicha C. Dvornek^{1,2}, Xiaoran Zhang¹, Gigi Galiana^{1,2}, Hemant D. Tagare^{1,2}, and R. Todd Constable^{1,2}

¹ Department of Biomedical Engineering, Yale University

² Department of Radiology & Biomedical Imaging, Yale University
ziqi.gao@yale.edu

Abstract. Low-field (LF) magnetic resonance imaging (MRI) improves accessibility and reduces costs but generally has lower signal-to-noise ratios and degraded contrast compared to high field (HF) MRI, limiting its clinical utility. Simulating LF MRI from HF MRI enables virtually evaluating novel imaging devices and developing LF algorithms. Existing low field simulators rely on noise injections and smoothing, which fails to capture the contrast degradation seen in LF acquisitions. To this end, we introduce an end-to-end LF-MRI synthesis framework that learns HF to LF image degradation directly from a small number of paired HF-LF MRIs. Specifically, we introduce a novel **HF** to **LF** coordinate-image decoupled neural Operator (H2LO) to model the underlying degradation process, and tailor it to capture high-frequency noise textures and dedicated image structure. Experimental results in T1w and T2w MRI demonstrate that H2LO produces more faithful simulated low-field images than existing parameterized noise synthesis model and popular image-to-image translation models. Furthermore, it improves performance in downstream image enhancement tasks, showcasing its potential to enhance LF MRI diagnostic capabilities.

Keywords: Neural Operator · Low-field MRI · Image Synthesis · Image Enhancement

1 Introduction

Low-field (LF) magnetic resonance imaging (MRI) offers improved accessibility and reduced cost compared to conventional high-field (HF) systems [2, 14, 17, 18, 24], yet generally produces images with lower signal-to-noise ratio (SNR), reduced spatial resolution, and diminished tissue contrast [1, 5, 8, 11, 13, 16], limiting its clinical utility. Establishing the clinical value of novel LF-MRI devices relies on prospective clinical studies [2, 14, 17, 24], which typically involve cohorts ranging from tens to over one hundred subjects. Assembling such validation cohorts often takes years [2, 14], creating a bottleneck for rapid device development and evaluation.

In-silico retrospective analysis [1] offers a complementary strategy to accelerate prospective clinical trials and device development by synthesizing subject-specific LF images from publicly available HF datasets. The HF-to-LF transformation can be learned from a small number of paired HF-LF scans, which are considerably easier to acquire than the large cohorts required for prospective validation. Arnold et al. proposed a parameterized noise-and-smoothing model that minimizes discrepancies in global histogram features using three subject pairs [1]. However, this approach assumes a monotonic intensity transformation and therefore cannot faithfully capture the non-linear, tissue-specific contrast shifts induced by field strength changes [16]. In particular, differential relaxation effects may alter the relative positions of tissue subpeaks (e.g., white matter and gray matter) within the intensity distribution, leading to inaccurate anatomical contrast in the synthesized LF images. Consequently, it has limited capacity to represent broader contrast variations across field strengths (Fig. 2(a), second column; (d), purple line).

To capture both SNR degradation and contrast changes, we formulate the HF-to-LF transformation as a field-strength-dependent transformation of the underlying MR signal and learn it in an end-to-end manner. Although this problem can be viewed from an image-to-image (I2I) translation perspective [7, 23], HF-to-LF translation differs from conventional pixel-level appearance mappings: it requires modeling how field strength alters the continuous spatial signal, including its frequency content and contrast behavior. In our experiments, popular pixel-wise I2I models [7, 23] are unable to simultaneously capture LF textures and accurately model contrast changes (3rd and 4th columns of Figure 2(a)). Therefore, we propose a novel high-frequency preserving 3D neural operator (H2LO), a lightweight coordinate-based network designed to model this transformation while preserving high-frequency structural information, and demonstrate its performance on both T1w and T2w MRI.

Besides LF MRI synthesis, we further demonstrate that H2LO supports LF algorithm development as a faithful data augmentation technique. We consider the LF image enhancement task [13, 8, 11, 4, 9], which aims to transform LF MRI into HF-like MRI to improve diagnostic utility. We augment an existing supervised LF image enhancement model [13], by simulating LF MRI from another public conventional MRI dataset, thereby creating additional LF-HF pairs.

In summary, our contributions are threefold. First, to support virtual clinical trials and algorithm development for novel MRI devices, we introduce the first end-to-end framework for LF MRI synthesis. Second, we formulate HF-to-LF translation as an operator learning problem. Building upon a theoretically bounded neural operator architecture [12], H2LO is tailored for preserving high-frequency in 3D volumetric data. Third, we compare H2LO with popular I2I translation networks and an LF simulator in terms of both synthesis quality and downstream utility, demonstrating superior performance on T1w and T2w MRI.




Fig. 1. Overview of the High-to-Low field Operator (H2LO) framework. The architecture maps HF MRI volumes to a continuous LF representation using a branch-trunk network optimized via pointwise intensity fidelity and gradient regularization.

2 Method

2.1 Problem Formulation

Let $I^{HF} \in \mathbb{R}^{H \times W \times D}$ and $I^{LF} \in \mathbb{R}^{H \times W \times D}$ denote a spatially aligned pair of HF and LF 3D MRI volumes acquired at different magnetic field strengths. From the perspective of MRI signal formation, image intensity depends on field-strength-dependent relaxation properties (e.g., T1 and T2) and SNR characteristics. Changing the magnetic field strength therefore alters the underlying signal response and its spatial frequency behavior, rather than merely modifying voxel-wise intensities.

We interpret each discretized MRI volume acquired at field strength E as samples from an underlying continuous spatial signal $f^E : \Omega \subset \mathbb{R}^3 \rightarrow \mathbb{R}$. Specifically, $I^{HF}[i, j, k] = f^{HF}(\mathbf{x}_{i,j,k})$, $I^{LF}[i, j, k] = f^{LF}(\mathbf{x}_{i,j,k})$, where $\{\mathbf{x}_{i,j,k}\}$ denotes the sampling grid.

Under this function-space formulation, HF-to-LF (H2L) translation amounts to learning a field-strength-dependent operator $\mathcal{T}_{HF \rightarrow LF} : f^{HF} \mapsto f^{LF}$. Let $u := f^{HF}$ for notation simplicity. We approximate $\mathcal{T}_{HF \rightarrow LF}$ with a neural operator $\mathcal{G}_\theta$ such that $\hat{f}^{LF}(\mathbf{x}) = \mathcal{G}_\theta(u)(\mathbf{x})$, $\mathbf{x} \in \Omega$.

This operator perspective models H2L translation as a mapping between function spaces rather than a fixed-grid voxel regression, aligning the formulation with the continuous nature of MR signal formation.

2.2 H2LO: High-to-Low Field Operator

We parameterize $\mathcal{G}_\theta$ utilizing branch-trunk factorization following DeepONet [12]. To let local tissue contrast modulate the operator response at each voxel, we extend the original formulation: while DeepONet’s branch coefficients are global summaries of the input function, our branch produces spatially varying coefficients by reading from a dense CNN feature field. The resulting operator takes the form

$$\mathcal{G}_\theta(u)(\mathbf{x}) = \sum_{k=1}^P b_k(u, \mathbf{x}) t_k(\mathbf{x}) + \beta. \quad (1)$$

where $b_k(u, \mathbf{x})$ are spatially varying field coefficients, $t_k(\mathbf{x})$ are coordinate-dependent basis functions, and β is a learnable bias term. As shown in Fig. 1, our framework comprises (1) a branch network and (2) a trunk network.

Branch network. The branch is a 3D image encoder that maps the input volume to a dense feature field $\mathbf{F}(u) \in \mathbb{R}^{P \times H \times W \times D}$. We adopt a lightweight 3D convolutional architecture from SRResNet [10]. For each query location $\mathbf{x}_{i,j,k}$, the coefficient vector $\mathbf{b}(u, \mathbf{x}_{i,j,k}) \in \mathbb{R}^P$ is obtained by reading the feature vector at the corresponding voxel. This produces spatially varying operator coefficients while requiring only a single encoder forward pass per volume.

Trunk network. The coordinate basis functions $t_k(\mathbf{x})$ are parameterized by a sinusoidal representation network (SIREN) [19] with four layers: one input sine layer, two hidden sine layers, and a linear output layer producing $P = 128$ outputs. Each sine layer computes $\sin(\omega_0 \mathbf{W}\mathbf{x} + \mathbf{b})$ with SIREN initialization. Sinusoidal activations enable accurate modeling of high-frequency spatial structures, preserving fine anatomical boundaries in the synthesized image.

2.3 Training and Inference

The total training loss combines an intensity fidelity term and a local gradient regularization term: $\mathcal{L} = \mathcal{L}_1 + \lambda_{grad}\mathcal{L}_{grad}$. The $\mathcal{L}_1$ term enforces zero-order fidelity over N randomly sampled voxel coordinates $\{\mathbf{x}_n\}_{n=1}^N$:

$$\mathcal{L}_1 = \frac{1}{N} \sum_{n=1}^N |\mathcal{G}_\theta(u)(\mathbf{x}_n) - f^{LF}(\mathbf{x}_n)|.$$

Reconstructing high-fidelity details requires further regularization of the local behavior of the learned operator. We therefore introduce a first-order consistency term by penalizing discrepancies between spatial derivatives of the predicted and target functions over randomly cropped sub-volumes:

$$\mathcal{L}_{grad} = \frac{1}{B} \sum_{b=1}^B \sum_{d \in \{x,y,z\}} \|\partial_d \mathcal{G}_\theta(u) - \partial_d f^{LF}\|_{L^2(V_b)}^2, \quad (2)$$

where $V_b \subset \Omega$ denotes the b -th randomly sampled sub-volume in a training iteration, B is the number of such sub-volumes, and ∂_d denotes differentiation along spatial dimension d . In practice, spatial derivatives are approximated using first-order derivatives of a 3D Gaussian kernel $G_\sigma(\mathbf{r}) = (2\pi\sigma^2)^{-3/2} \exp\left(-\frac{\|\mathbf{r}\|^2}{2\sigma^2}\right)$, with derivative filters $k_d(\mathbf{r}) = -\frac{r_d}{\sigma^2} G_\sigma(\mathbf{r})$. The discrete implementation becomes

$$\mathcal{L}_{grad} = \frac{1}{B} \sum_{b=1}^B \sum_{d \in \{x,y,z\}} \|\hat{V}_b * k_d - V_b * k_d\|_F^2, \quad (3)$$

where $\hat{V}_b$ and V_b denote the predicted and ground-truth sub-volumes, respectively, and $\|\cdot\|_F$ denotes the Frobenius norm over the discrete voxel grid.

Gaussian derivatives provide noise-robust and rotationally symmetric gradient estimates suited to low-SNR MRI. Matching spatial derivatives enforces first-order consistency of the operator output, complementing the zero-order intensity constraint.

Once trained, the synthesized LF volume is obtained by evaluating the learned operator on the discrete sampling grid: $\hat{I}^{LF}[i, j, k] = \hat{f}^{LF}(\mathbf{x}_{i,j,k})$.

3 Experiments

3.1 Low Field MRI Synthesis

A high-fidelity HF-to-LF transformation is essential for reliable in-silico retrospective clinical trials for novel LF device development, where only a small cohort of paired HF–LF scans—compared to the hundreds typically required in prospective studies [14]—is available to learn the mapping. In this study, we rigorously evaluate the synthesis performance of H2LO against related methods.

Dataset and Metrics. We train and evaluate on different partitions of the ds006557 dataset [21], which consists of 23 subjects scanned at both high-field (3T GE) and low-field (64mT Hyperfine) conditions. Both T1w and T2w contrasts are available, with one T1w subject excluded due to missing data. All volumes are skull-stripped using SynthStrip [6] and registered to a common space via ANTs [20], yielding paired volumes of size $208 \times 256 \times 256$. Each volume is normalized to $[0, 1]$ via max normalization. We perform 5-fold cross-validation, with each fold using 13-14 subjects for training, 2 for validation, and 7 for testing. Reconstruction quality is measured using peak signal-to-noise ratio (PSNR, in dB), structural similarity index (SSIM, in %), and normalized cross-correlation (NCC) between reconstructed and ground-truth images. Additionally, histogram-based metrics are used to assess intensity-distribution fidelity, including Wasserstein distance (WASS), histogram NCC (HNCC), Bhattacharyya distance (BHAT), and Jensen-Shannon (JS) divergence computed from intensity histograms. For histogram analysis, only foreground voxels are included using the ground-truth mask threshold ($I_{HF} > 0.01$).

Comparison Methods. We compare our method with three categories of approaches. (1) *Image-to-image (I2I) translation*, which learns a direct mapping between HF and LF MRI voxels. We include Pix2Pix [7] and DiffI2I [23] as representative GAN-based and diffusion-based models, respectively. (2) *Physics-inspired LF MRI simulation* [1], which approximates the LF acquisition process through noise modeling and spatial smoothing kernels. (3) *Conditional implicit neural representation (INR)*, which employs coordinate-based architectures similar to ours. Specifically, we adopt the 3D MRI implementation from [22], an extension of LIIF [3] designed for volumetric MRI data.

Implementation. The branch network adopts a lightweight 3D CNN encoder [10] with five convolutional layers (kernel size 3^3 , channels 32–32–64–64–128) with ReLU activations, which produces a spatially-resolved feature map of dimension $P = 128$. The trunk network is a 4-layer SIREN MLP (width 256, $\omega_0 = 30$) that

maps normalized 3D coordinates to P basis function values. For $\mathcal{L}_1$, N is set to 8000. For $\mathcal{L}_{grad}$, we set $\sigma = 1.0$ with a $5 \times 5 \times 5$ kernel based on hyperparameter search on validation sets and B is set to 1. We train for 500 epochs using Adam with an initial learning rate of 10^{-4} and cosine annealing to 10^{-6} .

3.2 Low Field MRI Enhancement

In this study, we examine the downstream utility of H2LO as a data engine for LF MRI enhancement, showcasing its ability to support LF algorithm development. We use an existing public HF dataset and simulate corresponding LF MRI using our trained HLFO. Generated HF-LF pairs thereby augment supervised LF MRI enhancement models.

Datasets and Metrics. The HF dataset is preprocessed 3T images from Open-neuro ds005752 [15], which contains 184 subjects with both T1w and T2w scans. These were split into 147/29 subjects for train/validation using a fixed random seed and an 80/20 partition. For each subject, LF counterparts for each HF structural MR volumes were generated (T1w and T2w available), with a typical voxel grid size of $208 \times 256 \times 256$. All images are intensity-normalized to $[0,1]$ and skull stripping is applied with SynthStrip. Real LF-HF pairs in ds006557 were retained for finetuning (using the training partition mentioned in Section 3.1) and evaluation (using the testing partition).

Comparison Methods. (1) *Real LF and Conventional Data Augmentations:* To increase data variation while preserving intensity consistency, we apply random paired in-plane augmentations (axial H,W flip and 90° rotations) to LF-HF training patches and keep the depth axis unchanged. (2) *Synthetic (Pre)training and Real LF Finetuning:* We generate synthetic LF inputs from pretrained HF-to-LF models including ours and comparison methods listed in Section 3.1.

Implementations. The PF-SR architecture [13] was adopted as the backbone network. Since the present task focuses on same-resolution enhancement rather than super-resolution, the 3D sub-pixel convolution upsampling module was removed. The final output is obtained via voxel-wise addition of the predicted residual and the original LF input. Pretraining was conducted on synthetic LF-HF image pairs with matched contrast, following the training protocol described in [13]. The model was subsequently fine-tuned on real LF-HF pairs using the same backbone and optimizer configuration. During fine-tuning, a cosine decay learning rate schedule was employed, with an initial learning rate of 2×10^{-6} , a 4-epoch warm-up phase starting from $0.5 \times$ the base learning rate, and a minimum learning rate of 5×10^{-7} . Experiments are conducted on NVIDIA H200, A100 and RTX A5000 GPUs.

4 Results

4.1 Low Field MRI Synthesis

We present the main synthesis results in Tab. 1 and Fig. 2(a). Since Arnold et al.’s physics simulator is explicitly optimized for matching histogram statistics,




Fig. 2. Visualization of generated LF MRIs and ablation studies. (a) Generated T1w (upper rows) and T2w (lower rows) LF images using multiple methods. Even rows visualize the absolute differences between simulated and real LF images. (b) Ablation of multiple components in our model. (c) Testing results of the 5-fold ablation study. (d) Comparison of histograms for LF images generated by Arnold et al. [1] and our method.

T1								
Method	PSNR(dB)↑	SSIM(%)↑	NCC↑	WASS↓	HNCC↑	BHAT↓	JS↓	Param.
HF Image	23.52±0.44	89.96±0.34	0.9222±0.0035	0.1211±0.0107	0.5130±0.0436	0.2148±0.0218	0.1495±0.0130	—
Arnold et. al. [1]	14.97±1.16	87.80±1.56	0.9466±0.0069	0.1099±0.0103	0.5759±0.0423	0.1009±0.0231	0.0823±0.0160	4
Pix2Pix [7]	26.46±0.33	92.25±0.27	0.9561±0.0032	0.0403±0.0046	0.8302±0.0530	0.0448±0.0106	0.0358±0.0080	4.86M
Diff21 [23]	27.16±0.54	92.50±0.83	0.9609±0.0054	0.0330±0.0046	0.9028±0.0145	0.0364±0.0065	0.0287±0.0047	16.39M
Conditional INR [22]	27.59±0.33	93.90±0.32	0.9650±0.0029	0.0327±0.0031	0.8642±0.0132	0.0426±0.0055	0.0342±0.0043	0.85M
Ours	28.98±0.27	94.76±0.22	0.9749±0.0017	0.0288±0.0049	0.9294±0.0141	0.0218±0.0068	0.0186±0.0057	0.58M
T2								
Method	PSNR(dB)↑	SSIM(%)↑	NCC↑	WASS↓	HNCC↑	BHAT↓	JS↓	Param.
HF Image	24.36±0.28	90.14±0.35	0.8059±0.0090	0.0583±0.0017	0.4818±0.0479	0.1759±0.0103	0.1442±0.0081	—
Arnold et. al. [1]	23.43±0.15	90.42±0.44	0.8297±0.0061	0.1094±0.0030	0.0415±0.0269	0.3626±0.0108	0.2709±0.0067	4
Pix2Pix [7]	25.45±0.28	91.42±0.46	0.8611±0.0095	0.0259±0.0031	0.9609±0.0124	0.0212±0.0043	0.0180±0.0036	4.86M
Diff21 [23]	26.88±0.23	92.33±0.36	0.8937±0.0074	0.0364±0.0037	0.9017±0.0203	0.0486±0.0056	0.0430±0.0049	16.39M
Conditional INR [22]	27.80±0.37	<u>93.78±0.43</u>	0.9147±0.0065	0.0260±0.0022	0.8729±0.0491	0.0504±0.0145	0.0439±0.0119	0.85M
Ours	28.37±0.38	94.29±0.48	0.9257±0.0064	0.0244±0.0050	0.9458±0.0094	0.0221±0.0027	0.0205±0.0024	0.58M

Table 1. Quantitative results of LF MRI synthesis. The mean $\pm$ standard deviation across five-fold cross-validation are reported. The **best** and second-best methods are highlighted in bold and underline, respectively.

we additionally visualize the intensity histograms in Fig. 2(d) to provide a more comprehensive evaluation. From Tab. 1, our method consistently achieves the best performance across almost all metrics for T1 synthesis. For T2 synthesis, Pix2Pix achieves the best results in part of the histogram-related statistics. However, the Pix2Pix column in Fig. 2(a) shows that its visual results exhibit unnatural grid-like artifacts (which also occur in Diff21), and such artifacts are not desirable for high-fidelity retrospective studies on novel devices. From the same figure, we can also observe that our method achieves the highest visual alignment with real LF images. From Fig. 2(a), Conditional INR produces overly smoothed LF images, while Arnold et al. simulate noise patterns but fail to properly model contrast changes, as reflected in the histogram comparison in Fig. 2(d). In contrast, our method maintains both strong quantitative performance and more visually coherent and anatomically faithful results.

Method	T1		T2		Subjects	Epoch
	PSNR(dB)	SSIM(%)	PSNR(dB)	SSIM(%)		
Real LF	30.12	92.93	29.26	92.61	13	200
Flip & Rotate	29.49	92.06	28.50	91.98	13 (*8)	400
Arnold et. al. [1]	30.37	93.48	29.39	92.75	13+147	150+50
Pix2Pix [7]	<u>30.39</u>	<u>93.77</u>	<u>29.59</u>	<u>93.17</u>	13+147	150+50
Diff21 [23]	30.27	93.33	29.52	93.01	13+147	150+50
Conditional INR [22]	30.10	92.97	29.34	92.97	13+147	150+50
Ours	30.55	93.71	30.02	93.72	13+147	150+50

Table 2. Quantitative result of LF MRI enhancement augmented with synthetic data. The **best** and second-best methods are highlighted in bold and underline, respectively.

4.2 Low Field MRI Enhancement

Table 2 summarizes the enhancement performance under different data augmentation strategies. Intensity-preserving conventional augmentation (flip & rotation) does not improve performance over real-only training, indicating that simple transformations are insufficient to compensate for limited LF data. In contrast, incorporating synthetic LF images consistently maintains or improves performance. Comparison across synthesis methods suggests that effective LF modeling requires both contrast and noise degradation modeling. Arnold et al. primarily simulate noise, providing texture cues but lacking accurate contrast modeling. Conditional INR captures contrast shifts but produces overly smooth results with limited noise characteristics. By modeling both contrast changes and noise degradation, our method achieves the best result for both T1w and T2w MRI.

4.3 Ablation Study

The ablation results for T1 synthesis are shown in Fig. 2(b)–(d). Visual comparisons in Fig. 2(b) indicate that the high-frequency gradient loss ($\mathcal{L}_{freq}$) improves local detail reconstruction and yields sharper LF textures. Histogram analysis in Fig. 2(b) and (d) further shows that incorporating SIREN enhances contrast modeling and contributes to finer structural representation. Quantitative results are reported in Fig. 2(c). Interestingly, ours w/o. SIREN achieves the highest PSNR and SSIM, while the full model ranks second. However, visual inspection shown in Fig. 2(b) suggests that the full model better preserves subtle anatomical details and produces more realistic LF characteristics.

5 Conclusion and Future Work

We present H2LO, a neural operator-based framework for subject-specific low-field MRI synthesis from high-field acquisitions. By formulating HF-to-LF translation as an operator learning problem, H2LO captures contrast shifts and high-frequency texture changes induced by field strength variation. Experiments on T1w and T2w MRI demonstrate improved synthesis fidelity over conventional

simulators and representative image-to-image translation models. Moreover, synthetic LF images generated by H2LO enhance downstream LF image enhancement performance, supporting its use for virtual clinical evaluation and algorithm development in emerging low-field MRI systems.

A limitation of this study is that it includes only healthy subjects. While this enables controlled evaluation, generalization to pathological cases remains to be established, as lesions or atrophy may alter the HF-to-LF mapping. Future work will extend the framework to clinical cohorts with diverse pathologies. In addition, the study is conducted on a limited number of paired subjects, reflecting current data availability and aligning with virtual clinical trial settings. To ensure methodological robustness, we mitigate the limited sample size through 5-fold cross-validation. Future work will validate the proposed method on larger and more diverse cohorts as such data become available.

Disclosure of Interests The authors have no competing interests in this work and other related research.

References

1. T Campbell Arnold, Steven N Baldassano, Brian Litt, and Joel M Stein. Simulated diagnostic performance of low-field mri: Harnessing open-access datasets to evaluate novel devices. *Magnetic resonance imaging*, 87:67–76, 2022.
2. Adrienne E Campbell-Washburn, Rajiv Ramasawmy, Matthew C Restivo, Ipshita Bhattacharya, Burcu Basar, Daniel A Herzka, Michael S Hansen, Toby Rogers, W Patricia Bandettini, Delaney R McGuirt, et al. Opportunities in interventional and diagnostic imaging by using high-performance low-field-strength mri. *Radiology*, 293(2):384–393, 2019.
3. Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8628–8638, 2021.
4. Zhuo-Xu Cui, Congcong Liu, Chentao Cao, Yuanyuan Liu, Jing Cheng, Qingyong Zhu, Yanjie Zhu, Haifeng Wang, and Dong Liang. Meta-learning enabled score-based generative model for 1.5 t-like image reconstruction from 0.5 t mri. *arXiv preprint arXiv:2305.02509*, 2023.
5. Matteo Figini, Hongxiang Lin, Godwin Ogbale, Felice D Arco, Stefano B Blumberg, David W Carmichael, Ryutaro Tanno, Enrico Kaden, Biobele J Brown, Ikeoluwa Lagunju, et al. Image quality transfer enhances contrast and resolution of low-field brain mri in african paediatric epilepsy patients. *arXiv preprint arXiv:2003.07216*, 2020.
6. Andrew Hoopes, Jocelyn S Mora, Adrian V Dalca, Bruce Fischl, and Malte Hoffmann. Synthstrip: skull-stripping for any brain image. *NeuroImage*, 260:119474, 2022.
7. Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *CVPR*, 2017.
8. Mohammad Javadi, Robert Griffin, Panagiotis Tsiamyrtzis, Ernst Leiss, Andrew G Webb, and Nikolaos V Tsekos. In-silico comparison of a diffusion model with conventionally trained deep networks for translating 64mt to 3t brain flair. *Scientific Reports*, 15(1):38052, 2025.

9. Seunghoi Kim, Henry FJ Tregidgo, Ahmed K Eldaly, Matteo Figini, and Daniel C Alexander. A 3d conditional diffusion model for image quality transfer—an application to low-field mri. *arXiv preprint arXiv:2311.06631*, 2023.
10. Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.
11. Xiyue Lin, Chenhe Du, Qing Wu, Xuanyu Tian, Jingyi Yu, Yuyao Zhang, and Hongjiang Wei. Zero-shot low-field mri enhancement via denoising diffusion driven neural representation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 775–785. Springer, 2024.
12. Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021.
13. Christopher Man, Vick Lau, Shi Su, Yujiao Zhao, Linfang Xiao, Ye Ding, Gilberto KK Leung, Alex TL Leong, and Ed X Wu. Deep learning enabled fast 3d brain mri at 0.055 tesla. *Science Advances*, 9(38):ead9327, 2023.
14. Mercy H Mazurek, Bradley A Cahn, Matthew M Yuen, Anjali M Prabhat, Isha R Chavva, Jill T Shah, Anna L Crawford, E Brian Welch, Jonathan Rothberg, Laura Sacolick, et al. Portable, bedside, low-field magnetic resonance imaging for evaluation of intracerebral hemorrhage. *Nature communications*, 12(1):5119, 2021.
15. Allison C. Nugent, Adam G Thomas, Margaret Mahoney, Alison Gibbons, Jarrod Smith, Antoinette Charles, Jacob S Shaw, Jeffrey D Stout, Anna M Namyst, Arshitha Basavaraj, Eric Earl, Dustin Moraczewski, Emily Guinee, Michael Liu, Travis Riddle, Joseph Snow, Shruti Japee, Morgan Andrews, Adriana Pavletic, Stephen Sinclair, Vinai Roopchansingh, Peter A Bandettini, and Joyce Chung. "the nimh healthy research volunteer dataset", 2025.
16. Thomas O'Reilly and Andrew G Webb. In vivo t1 and t2 relaxation time maps of brain tissue, skeletal muscle, and lipid measured in healthy volunteers at 50 mt. *Magnetic Resonance in Medicine*, 87(2):884–895, 2022.
17. Hemmen Sabir, Florian Kipfmüller, Soyhan Bagci, Till Dresbach, Tamara Grass, Patrizia Nitsch-Felsecker, Christos Pantazis, Joachim Schmitt, Lukas Schroeder, and Andreas Mueller. Feasibility of bedside portable mri in neonates and children during ecl. *Critical care*, 27(1):134, 2023.
18. Anja Samardzija, Kartiga Selvaganesan, Horace Z Zhang, Heng Sun, Chenhao Sun, Yonghyun Ha, Gigi Galiana, and R Todd Constable. Low-field, low-cost, point-of-care magnetic resonance imaging. *Annual Review of Biomedical Engineering*, 26, 2024.
19. Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.
20. Nicholas J Tustison, Philip A Cook, Andrew J Holbrook, Hans J Johnson, John Muschelli, Gabriel A Devenyi, Jeffrey T Duda, Sandhitsu R Das, Nicholas C Cullen, Daniel L Gillen, et al. The antsx ecosystem for quantitative biological and medical imaging. *Scientific reports*, 11(1):9068, 2021.
21. František Váša, Carly Bennallick, Niall J Bourke, Francesco Padormo, Levente Baljer, Ula Briski, Paul Cawley, Tomoki Arichi, Tobias C Wood, David J Lythgoe, et al. Ultra-low-field brain mri morphometry: test–retest reliability and correspondence to high-field mri. *Imaging Neuroscience*, 3:IMAG–a, 2025.

22. Qing Wu, Yuwei Li, Yawen Sun, Yan Zhou, Hongjiang Wei, Jingyi Yu, and Yuyao Zhang. 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.
23. Bin Xia, Yulun Zhang, Shiyin Wang, Yitong Wang, Xinglong Wu, Yapeng Tian, Wenming Yang, Radu Timotfe, and Luc Van Gool. Diffi2i: efficient diffusion model for image-to-image translation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(3):1578–1593, 2024.
24. Matthew M Yuen, Anjali M Prabhat, Mercy H Mazurek, Isha R Chavva, Anna Crawford, Bradley A Cahn, Rachel Beekman, Jennifer A Kim, Kevin T Gobeske, Nils H Petersen, et al. Portable, low-field magnetic resonance imaging enables highly accessible and dynamic bedside evaluation of ischemic stroke. *Science advances*, 8(16):eabm3952, 2022.