

Super-Field MRI Synthesis for Infant Brains Enhanced by Dual Channel Latent Diffusion

Austin Tapp¹, Can Zhao², Holger R. Roth², Jeffrey Tanedo³, Syed Muhammad Anwar^{1,7}, Niall J. Bourke⁴, Joseph Hajnal⁵, Victoria Nankabirwa⁶, Sean Deoni⁸, Natasha Lepore³, and Marius George Linguraru^{1,7}

¹ Sheikh Zayed Institute for Pediatric Surgical Innovation, Children's National Hospital, Washington, DC 20010

² NVIDIA Corporation, 2788 San Tomas Expressway Santa Clara, CA 95051, USA

³ CIBORG Lab, Department of Radiology, Children's Hospital Los Angeles and University of Southern California, Los Angeles, CA 90027, USA

⁴ Centre for Neuroimaging Sciences, King's College London, United Kingdom

⁵ Centre for the Developing Brain, School of Biomedical Engineering and Imaging Sciences, King's College London, London, United Kingdom

⁶ Makerere University School of Public Health, Kampala, Uganda

⁷ School of Medicine and Health Sciences, George Washington University, Washington, DC 20052, USA

⁸ Bill and Melinda Gates Foundation, PO Box 23350, Seattle, WA, 98102, USA
atapp@childrensnational.org
mlingura@childrensnational.org

Abstract. In resource-limited settings, portable ultra-low-field (uLF, i.e., 0.064T) magnetic resonance imaging (MRI) systems expand accessibility of radiological scanning, particularly for low-income areas as well as underserved populations like neonates and infants. However, compared to high-field (HF, e.g., ≥ 1.5 T) systems, inferior image quality in uLF scanning poses challenges for research and clinical use. To address this, we introduce Super-Field Network (SFNet), a custom swinUNETRv2 with generative adversarial network components that uses uLF MRIs to generate super-field (SF) images comparable to HF MRIs. We acquired a cohort of infant data ($n=30$, aged 0-2 years) with paired uLF-HF MRI data from a resource-limited setting with an underrepresented population in research. To enhance the small dataset, we present a novel use of latent diffusion to create dual-channel (uLF-HF) paired MRIs. We compare SFNet with state-of-the-art synthesis methods by HF-SF image similarity perceptual scores and by automated HF and SF segmentations of white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF). The best performance was achieved by SFNet trained on the latent diffusion enhanced dataset yielding state-of-the-art results in Fréchet inception distance at 9.08 ± 1.21 , perceptual similarity at 0.11 ± 0.01 , and PSNR at 22.64 ± 1.31 . True HF and SF segmentations had a strong overlap with Dice similarity coefficients of 0.71 ± 0.1 , 0.79 ± 0.2 , and 0.73 ± 0.08 for WM, GM, and CSF, respectively, in the developing infant brain with incomplete myelination, and displayed 166%, 107%, and 106% improvement over respective uLF-based segmentation metrics. SF MRI supports health equity by enhancing the clinical use of uLF imaging systems and improving the diagnostic capabilities of low-cost portable MRI systems in resource-limited settings and for underserved populations. Our code is made openly available at <https://github.com/AustinTapp/SFnet>.

Keywords: Magnetic Resonance Imaging, Ultra Low Field MRI, Synthesis, Latent Diffusion, Pediatric, Health Equity

1 Introduction

Clinically, neuroimaging of infants is crucial for assessing risks in brain development, detecting brain injuries, diagnosing neurological conditions, planning interventions, and monitoring treatment progress. In research, magnetic resonance imaging (MRI) has also enabled the characterization of early brain development, mostly in high-resource regions [1-4]. However, MRI is rarely available as a screening test in children despite the increasing demand for radiation-free pediatric imaging [5]. In fact, among medical imaging systems, high-field (HF, i.e., $\geq 1.5\text{T}$) MRI may be the most inequitably used due to high costs, complex operability with radiofrequency-shielded radiology suites, and the necessity for patients to travel to specialized clinics [6]. Thus, conventional HF MRI systems are rare or unavailable in resource-limited settings like the hospital infrastructure in Uganda, where this study’s data originates from, with only 10 MRI machines in a country of over 46 million inhabitants [6, 7].

Recent advancements in portable, ultra-low-field MRI (uLF MRI, i.e., 0.064T) technology has enabled imaging at the point of care and even around ferromagnetic materials [8, 9]. Compared to HF systems, uLF MRI offers additional advantages such as lower installation and maintenance costs, reduced power consumption, smaller space requirements, and elimination of cryogenic cooling [10, 11]. Moreover, uLF systems are great for children because of faster acquisition, lower acoustic noise, and open scanner designs, improving infant scanning success rates without sedation [12-16]. However, uLF MRI systems have weak sample signals, low signal-to-noise ratio (SNR), and poor contrast, leading to inferior image quality [9, 10].

Super-resolution (SR) techniques, which produce higher-resolution images from lower-resolution scans, offer potential solutions for improving image quality [17-22]. SR methods employing convolutional neural networks (CNNs), generative adversarial networks (GANs), and transformers have shown breakthrough performance among various SR benchmarks [23]. CNN methods like AUTOMAP have enhanced image quality and SNR gains in diverse scenarios, including uLF MRI acquired in adults [24], while methods such as SMORE use a self-supervised approach to restore image quality by leveraging the innate high and low-resolution image information [25]. Pure CNN-based methods are also tailored to include autoencoders, encoder-decoders, and GANs, which increase inference speed, expand the CNN’s receptive field, or capture rich texture details, respectively [23]. In the medical image domain, GAN SR approaches face challenges, such as hallucinations, and as in SOUP-GAN, are integrated with other methods for stabilization, thus outperforming previous work [26-28]. Further, transformer networks integrating reconstruction and super-resolution into two sub-branches can produce motion-artifact-free SR images from degenerated MRI [29].

In this study, we aim to use a single framework that enhances multiple aspects of uLF image quality metrics, including image contrast, SNR (Signal-to-Noise Ratio), and matched likeness to HF images; we refer to this as ‘super-field’ (SF) synthesis. Two approaches are closely related to our work: LoHiResGAN [19] and LF-SynthSR [21], but these methods focus on SR alone. LoHiResGAN integrates ResNet downsample and upsample blocks into a GAN architecture for low-resolution to high-resolution MRI translation that considers image structure and texture similarity. An extension of SynthSR, LF-SynthSR [21], applies 3-D U-net with a semantic segmentation loss in the architecture for predicting SR outputs from low-field MRI. LF-SynthSR incorporates a

synthetic data generator during training to present diverse image appearances. Although SR for uLF-MRI systems have been successful in adult cohorts [18-22], for infant and neonate populations, SR has only been explored by combining three orthogonal (coronal, sagittal, and axial) uLF scans to produce isotropic, high-resolution images; this recombination method is commonly known as SR reconstruction (SRR) [30-33]. While SR in [19] and [21] primarily focus on improving resolution, our aim for SF synthesis extends SR approaches by addressing image quality (SNR, perception, and structural representation). Image quality is especially critical for imaging infant brains, which exhibit poor contrast from incomplete myelination in healthy, developing brain structures. Therefore, SR differs from our SF approach, which yields resultant MRIs that exhibit enhanced quality and fidelity beyond SR and SRR outputs, thus making the original uLF images more suitable for clinical interpretation and analysis, as in Fig. 1.

Novelty and Contribution We present a novel SF MRI enhancement method that increases the image quality of emerging portable uLF MRI scans with the aim to rival the perceptual information offered by HF MRI at lower cost and at the point of care of underserved populations. By separating SF from SR, our proposed method, super-field network (SFNet), enhances uLF MRI quality by combining transformer attention mechanisms with CNN feature extraction in GAN-like framework. An additional novelty in our approach is the dual channel latent diffusion process that generates paired (uLF-HF) synthetic images to supplement our small dataset, a typical problem in infant imaging made trickier by the difficulty of acquiring uLF MRIs in limited-resource areas like Uganda. We compare our approach, SFNet, with three state-of-the-art methods, LoHiResGAN [19], LF-SynthSR [21], and SwinUNETRv2 [34], by quantitative perceptual metrics and anatomical segmentation analyses to demonstrate SFNet’s potential contribution to global health equity via empowering low-cost portable MRI to offer clinical insight similar to that of less accessible HF MRI.




Fig. 1. uLF (0.064T) images acquired in 3 directions orthogonally (axial, coronal, and sagittal); all 3 are used in traditional super resolution reconstruction (SRR), which is the image that serves as the input for our SFNet, which outputs a sharper, contrast enhanced, $1 \times 1 \times 1 \text{ mm}^3$ image synthesized through the super-field (SF) process.

2 Methods

2.1 Data Augmentation by Dual Channel Latent Diffusion

While latent diffusion has been used for MRI synthesis (e.g., T1 to T2), it has not been applied to a 2 channel (uLF-HF) paired data generation task [35]. As our infant data is small ($n = 30$), we utilize latent diffusion to generate images in a novel 2-channel approach stemming from methods described in [36]. A 3D variational autoencoder (VAE) maps a 2 channel (1 channel uLF and 1 channel HF) volume, $x \in \mathbb{R}^{2 \times H \times W \times D}$, with height (H), width (W), and depth (D), to a $6 \times 28 \times 28 \times 28 \text{ mm}^3$ uLF-HF latent representation, z , where $z \in \mathbb{R}^{6 \times H \times W \times D}$. The VAE decoder synthesizes a paired, 2 channel, uLF-HF image of size x from z , which is a probabilistic representation of the anatomical structure and features of the uLF and HF 2 channel MRIs. The VAE acts as a perceptual compression model trained through a combination of perceptual loss and a patch-based adversarial objective to yield images that look similar to input data [36]. After VAE training, latent noise of size $6 \times 28 \times 28 \times 28 \text{ mm}^3$ are used to train the diffusion model that learns the data distribution $p(x)$ by denoising a normally distributed variable. For synthesis, we use a reweighted variant of the variational lower bound on $p(x)$, mirroring denoising score-matching. During inference, a paired uLF-HF volume is produced from random, z -shaped noise. For latent diffusion network training, we use an Adam optimizer with a learning rate of 0.0001. A batch size of 1 is used for the VAE and a batch size of 6 is used for the diffusion model, which is a diffusion U-Net. The VAE is trained for 1,000 epochs and the diffusion U-Net is trained for 10,000 epochs. After dual channel image generation, a perceptual evaluation is conducted to assess the realism of synthesized images compared to the real images from the uLF-HF dataset. Evaluation metrics include Fréchet Inception Distance (FID) for gauging similarity and realism between real and synthetic images and Maximum Mean Discrepancy (MMD) to assess preservation of image characteristics. Synthetic image features were evaluated with a pretrained RadImageNet [37] and compared to image features from the input, x . Synthetic images were deemed adequate if they demonstrated FID of less than 16 and an MMD of less than 0.005 based on RadImageNet evaluation. The top 60 images meeting our quality metrics were subsequently added to the real uLF-HF paired image dataset to enhance SFNet training with a diverse range of realistic representations.




Fig. 2. uLF (0.064T) images are translated to SF MRI with methods prioritizing image perceptual quality via GAN components and supplement data with generated images.

2.2 Super-Field Synthesis

We propose SF enhancement on uLF images using SFNet, an enhanced swinUNETRv2 [34] that acts as the generator within a GAN framework. Fig. 2. shows an overview of the approach. Briefly, swinUNETRv2’s transformers and moving-window scheme support global image features. Further, depth feature extraction and multi-contrast feature fusion are contained within swin transformer residual blocks, each using multiple swin transformer layers (STL) for local attention and cross-window interactive learning. STL multi-head self-attention blocks and its multi-layer perception extract features used by multiple cascaded multi-scale aggregation blocks (MAB) [34]. MAB include a spatial adaptation block and joint residual feature aggregation block. Additionally, we utilize a discriminator in conjunction with the feature extraction performed natively by swinUNETRv2. Thus, output images are evaluated globally to yield more realistic images with limited anti-aliasing because of adversarial, SSIM, and perceptual component losses. The network’s overall loss is defined by Equation 1.

$$\mathcal{T} = \min_G \max_D (\lambda_1 \mathcal{L}_{ssim}(\hat{y}, y) + \lambda_2 \mathcal{L}_{Adv}(1 - D(\hat{y}, y)) + \lambda_3 \mathcal{L}_{Percept}(\hat{y}, y)) \quad (1)$$

Where y is the ground truth HF image, $\hat{y}$ is the SF image predicted from the uLF image. The λ_1 , λ_2 , and λ_3 terms are 2, 0.25 and 2, respectively. $\mathcal{L}_{Percept}$ represents the perceptual loss computed using a frozen-weight, pretrained deep neural network; in this case AlexNet. $\mathcal{L}_{Adv}$ represents the discriminator-based loss, a Patch-GAN discriminator based on Pix2PixHD, which learns in conjunction during training of the SFNet. $\mathcal{T}$ is the total function to optimize during training. All losses were computed volumetrically.

SFNet was curated through numerous ablation studies. In the first study, we pre-train SFNet in a self-supervised manner with 1,500 T2-weighted HF (1.5 or 3T) MRIs (mean age 16 months, 52% female) from developing Human Connectome projects’ (dHCP) open-source database [38]. Then, we fine-tune SFNet using either the original dataset (n=30) or the dataset supplemented by latent diffusion-generated data (n=90). In the second study, publicly available swinUNETR pre-trained weights from [39] were loaded into SFNet, which was then pre-trained using dHCP’s T2-weighted MRIs before training with the original (n=30) or supplemented datasets (n=90). In these prior studies, SFNet was also trained without GAN components (i.e., as swinUNETRv2). Finally, SFNet and SFNet without GAN components (i.e., swinUNETRv2) was trained from scratch with the original dataset (n=30) or latent diffusion-supplemented dataset (n=90). Our repository details training code for all models and their hyperparameters.

3 Experiments

3.1 Data

All images were acquired under IRB approval using a combination of uLF and HF MRI scanners from Makerere University in Kampala, Uganda. T2-weighted uLF images (spin echo TR/TE 1.5s/5ms) were obtained using the SWOOP (Hyperfine) portable 0.064T MR scanner as a trio of orthogonally acquired (axial, sagittal, and coronal orientation) anisotropic 3D T2-weighted MR images with in-plane resolution of 1.5mm x 1.5mm and 5 mm in slice selection direction. Paired T2-weighted HF images were

also acquired as a trio of orthogonal anisotropic 3D T2w MR images with in-plane resolution of 0.5mm x 0.5mm and 4.5 mm in the slice selection direction from 1.5T MAGNETOM Semptra (Siemens Healthineers) (TR/TE 5s/95ms) scanners. The healthy cohort of 22 male and 48 female subjects mean age was 381 ± 17 days old at scan time. Images were quality assessed for motion or banding artifacts. Images with severe distortions that failed quality checks, in both uLF and HF MRI, were removed from the study. After assessing image quality, 30 healthy patient pairs of HF and uLF T2 MRIs were included in the study; acquired orthogonal uLF images are shown in Fig. 1.

3.2 Super Resolution Reconstruction (SRR)

Each subject’s trio of orthogonally acquired uLF and HF images was reconstructed using Advanced Normalization Tools (ANTS) multivariate template construction tool [40, 41] as outlined in [17]. The tool solves a group-wise registration for the trio of images through iterative diffeomorphic transformations to generate a single, isotropic, super resolution reconstructed (SRR) volume with combined information from all three orthogonally acquired images. As shown in Fig. 1, after ANTS reconstruction, HF images have voxel grid spacing of $1.0 \times 1.0 \times 1.0 \text{ mm}^3$, while uLF images have a voxel grid spacing of $1.5 \times 1.5 \times 1.5 \text{ mm}^3$. Following SRR for both HF and uLF images, the uLF SRR volumes were rigidly registered to their SRR HF counterpart to ensure alignment and yield a uLF SRR image with voxel grid spacing of $1.0 \times 1.0 \times 1.0 \text{ mm}^3$. Throughout the manuscript, and for the remainder of the paper, HF SRR and uLF SRR volumes are referred to as HF and uLF, respectively; orthogonal images were termed explicitly.

3.3 SFNet Implementation and Results

We compare results from SFNet with the vanilla SwinUNETRv2 [34], LoHiResGAN [19] and LF-SynthSR [21]. LoHiResGAN and LF-SynthSR, are the only known open-source, pretrained models affirmed by authors as usable for uLF-HF synthesis ‘off-the-shelf’. Results are evaluated by the following perceptual metrics: FID, SSIM, PSNR, and learned perceptual image patch similarity (LPIPS). In addition, quantitative comparisons of white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF) segmented from the uLF and SF images were compared to segmentations from ground truth HF images obtained by iBeat, a state-of-the-art infant brain segmentation pipeline, using dice similarity coefficient (DSC) and relative volume error (RVE) metrics [42]. Significance was assessed at $p < 0.05$ with the Wilcoxon signed rank test. Training and inference were done on an NVIDIA A5000 (24 GB) GPU. Networks were implemented in the PyTorch-based framework MONAI. For each task, SFNet was trained for a total of 5,000 epochs. All inference times were under 1 minute; SF-Net requires 8.23 GB, SwinUNETRv2 requires 8.23 GB, LoHiResGAN requires 23.61 GB, and SynthSR requires 1.22 GB of GPU memory to generate MRIs of size (192, 192, 192).

Table 1 displays comparative results of our tested methods. The best SFNet used initial weights from [39], was pre-trained with dHCP’s T2 MRIs [38] and trained on the supplemented dataset ($n=90$), achieving state-of-the-art results in FID, LPIPS, and PSNR. Furthermore, when compared to HF segmentations calculated by iBeat, those from SF volumes demonstrated significant improvement over original uLF MRI. Specifically, DSC for WM, GM, and CSF were 0.71 ± 0.1 , 0.79 ± 0.2 , and 0.73 ± 0.08 , respectively. These coefficients represented improvements of 166%, 107%, and 106%

over the corresponding uLF-based segmentations. Figure 3 shows qualitatively the T2 images and their segmentations (WM=Blue, GM=Green, CSF=Red) output by iBeat.




Fig. 3. Segmentations obtained using iBeat; from left to right the SRR uLF MRI, the true HF MRI, our best SFNet MRI, SFNet MRI and a swinUNETRv2 SF MRI. † denotes pretraining.

Table 1. Model performance for generated images compared to ground truth high-field images. Only best performing SFNet models for multiple training configurations are shown (typically $n=90$). An * indicates significance over previous network performance; † denotes pretraining.

Model	FID ↓	SSIM ↑	LPIPS ↓	PSNR ↑	DSC ↑	RVE ↓
SFNet (n=90) [†]	9.08±1.21*	0.84±0.07*	0.11±0.01*	22.64 ±1.31*	0.74±0.03*	0.14±0.02*
SFNet (n=30) [†]	11.10±1.30*	0.74±0.17*	0.14±0.11*	19.64 ±1.21*	0.68±0.09*	0.24±0.04*
SFNet (n=90)	13.81±1.34	0.52±0.21	0.22±0.08	15.71 ±1.21	0.60±0.04	0.27±0.11
swinUv2 (90) [†]	14.12±1.30*	0.56±0.27*	0.24±0.11*	13.64 ±3.21	0.58±0.19	0.28±0.13
swinUv2 (90)	24.90±3.04	0.44±0.23	0.29±0.13	16.43 ±2.57	0.59±0.15*	0.27±0.09*
LHResGAN	26.43±3.52	0.41±0.25	0.31±0.11	16.66 ±2.35	0.34±0.05	0.38±0.13
LF-SynthSR	33.96±4.52	0.48±0.17	0.35±0.12	16.34 ±2.31	0.37±0.08	0.34±0.12
uLF SRR MRI	80.21±24.7	0.43±0.13	0.43±0.16	12.66 ±0.59	0.48±0.15	0.30±0.12

4 Discussion

In this study, we introduced SFNet, a novel SF MRI enhancement method designed to extend the benefits of MRI to underserved populations and resource-limited settings. In addition, SFNet can increase the clinical utility of emerging uLF imaging systems, particularly for low-cost portable MRI systems. Generally, in MRI, and especially in low-SNR uLF MRI, infant brains pose unique challenges due to incomplete myelination and low contrast between brain structures. Our SFNet focuses on elevating image quality, in addition to a more conventional SR reconstruction, to rival the perceptual information offered by HF MRI systems as shown by the significant boost

in perceptual metrics, as well as in the segmentation of infant brain structures. For this purpose, we designed a network that combines attention mechanisms with CNN feature extraction via a transformer-based network with GAN components. Additionally, our approach incorporates a dual-channel latent diffusion process to generate paired uLF-HF synthetic images and enhance small datasets, a typical problem in data collection from underrepresented populations in resource-constrained environments. By focusing on SF enhancement rather than classic SR, SFNet performed significantly better than prominent methods for image synthesis and SR reconstruction. Notably, SFNet reached state-of-the-art results for perceptual metrics e.g., average FID improved from 80.2 to 9.08 for uLF vs. SF MRI with an associated improvement in average DSC from 0.48 to 0.74, respectively. The use of uLF scanners supplemented with SF techniques to enhance scan quality presents a viable solution in such settings. We acknowledge our limitations regarding dataset size but obtaining paired HF data from resource-limited regions like Uganda, is difficult, and as uLF MRI technology is emerging, there are no public uLF datasets available. Future research will include more datasets and explore model applications in diverse cohorts, including patients with pathological conditions.

5 Conclusion

Our study highlights the transformative potential of super-field MRI enhancement to advance healthcare equity by bridging the gap in diagnostic capabilities between HF and uLF MRI systems. The latter have the advantage of low-cost, portability, and usability in ferromagnetic environments. However, they lack in image quality due to low SNR and contrast. By leveraging attentions mechanisms with GAN components, and dual channel latent diffusion-driven image augmentations, SFNet significantly improve the image quality and the quantitative information in uLF MRIs. As we demonstrated with data acquired from Ugandan infants, we aim to empower healthcare providers in resource-limited settings to deliver quality care and improve health outcomes for underserved populations worldwide.

Acknowledgments. Support for this work was provided by NIH NIDCR R01DE030286 and the Bill and Melinda Gates Foundation INV-005798.

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

References

1. Prayer, D., et al.: MRI of normal fetal brain development. *Eur. J. Radiol.* 57, 199-216 (2006).
2. Judas, M., et al.: Structural immunocytochemical and MR imaging properties of periventricular crossroads of growing cortical pathways in preterm infants. *AJNR Am. J. Neuroradiol.* 26, 2671-84 (2005).
3. Sadeghi, N., et al.: Regional characterization of longitudinal DTI to study white matter maturation of the early developing brain. *Neuroimage* 68, 236-47 (2013).
4. Gui, L., et al.: Morphology-driven automatic segmentation of MR images of the neonatal brain. *Med. Image Anal.* 16, 1565-79 (2012).

5. Huppi, P. S., Inder, T. E.: Magnetic resonance techniques in the evaluation of the perinatal brain: recent advances and future directions. *Semin. Neonatol.* 6, 195-210 (2001).
6. Mazurek, M. H., et al.: Portable bedside low-field magnetic resonance imaging for evaluation of intracerebral hemorrhage. *Nat. Commun.* 12, 5119 (2021).
7. Emerg Research & Consulting: HospeTrack Uganda (2020). [Online]. Available: <https://www.emerg-inc.com/wp-content/uploads/2020/01/HospeTrack-Uganda.pdf>
8. Lothar, S., et al.: Design of a mobile homogeneous and efficient electromagnet with a large field of view for neonatal low-field MRI. *MAGMA* 29, 691–698 (2016).
9. Campbell-Washburn, A. E., et al.: Opportunities in interventional and diagnostic imaging by using high-performance low-field-strength MRI. *Radiology* 293, 384–93 (2019).
10. Arnold, T. C., Freeman, C. W., Litt, B., Stein, J. M.: Low-field MRI: Clinical promise and challenges. *J. Magn. Reson. Imaging* 57(1), 25–44 (2023).
11. Kimberly, W. T., et al.: Brain imaging with portable low-field MRI. *Nat. Rev. Bioeng.* 1(9) (2023). [Online]. Available: <https://doi.org/10.1038/s44222-023-00086-w>
12. Mazurek, M. H., et al.: Portable bedside low-field magnetic resonance imaging for evaluation of intracerebral hemorrhage. *Nat. Commun.* 12(1), 5119 (2021).
13. Rusche, T., et al.: More Space Less Noise-New-generation Low-Field Magnetic Resonance Imaging Systems Can Improve Patient Comfort: A Prospective 0.55T-1.5T-Scanner Comparison. *J. Clin. Med.* 11(22), 6705 (2022).
14. Sheth, K. N., et al.: Assessment of Brain Injury Using Portable Low-Field Magnetic Resonance Imaging at the Bedside of Critically Ill Patients. *JAMA Neurol.* (2020).
15. Yuen, M. M., et al.: Portable low-field magnetic resonance imaging enables highly accessible and dynamic bedside evaluation of ischemic stroke. *Sci. Adv.* 8(16) (2022).
16. Chetcuti, K., Chilingulo, C., et al.: Implementation of a Low-Field Portable MRI Scanner in a Resource-Constrained Environment: Our Experience in Malawi. *AJNR Am. J. Neuroradiol.* 43(5) (2022). [Online]. Available: <https://doi.org/10.3174/ajnr.A7494>
17. Deoni, S. C. L., et al.: Accessible pediatric neuroimaging using a low field strength MRI scanner. *NeuroImage* 238, 118273 (2021).
18. Deoni, S. C. L., O'Muircheartaigh, J., Ljungberg, E., Huentelman, M., Williams, S. C. R.: Simultaneous high-resolution T2-weighted imaging and quantitative T2 mapping at low magnetic field strengths using a multiple TE and multi-orientation acquisition approach. *Magn. Reson. Med.* 88(3), 1273-1281 (2022).
19. Islam, K. T., et al.: Improving portable low-field MRI image quality through image-to-image translation using paired low- and high-field images. *Sci. Rep.* 13(1) (2023).
20. Cooper, R. E., et al.: Bridging the gap: improving correspondence between low-field and high-field magnetic resonance images in young people. *Front. Neurol.* 15 (2024). [Online]. Available: <https://doi.org/10.3389/fneur.2024.1339223>
21. Iglesias, J. E., et al.: Quantitative Brain Morphometry of Portable Low-FieldStrength MRI Using Super-Resolution Machine Learning. *Radiology* 306(3), e220522 (2023). [Online]. Available: <https://doi.org/10.1148/radiol.220522>
22. Iglesias, J. E., et al.: SynthSR: A public AI tool to turn heterogeneous clinical brain scans into high-resolution T1-weighted images for 3D morphometry. *Sci. Adv.* 9, eadd3607 (2023). [Online]. Available: <https://doi.org/10.1126/sciadv.add3607>
23. Yang, H., et al.: Deep learning in medical image super resolution: a review. *Appl. Intell.* 53, 20891–20916 (2023). [Online]. Available: doi.org/10.1007/s10489-023-04566-9
24. Koonjoo, N., et al.: Boosting the signal-to-noise of low-field MRI with deep learning image reconstruction. *Sci. Rep.* 11, 8248 (2021).
25. 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 Trans. Med. Imaging* 40(3), 805-17 (2021).

26. Zhang, K., Hu, H., Philbrick, K., Conte, G. M., Sobek, J. D., Rouzrokh, P., Erickson, B. J.: SOUP-GAN: Super-Resolution MRI Using Generative Adversarial Networks. *Tomography* 8(2) (2022).
27. Ledig, C., Theis, L.: Photo-Realistic Single image Super-Resolution using a generative adversarial network. In: *Computer Vision and Pattern Recognition (CVPR)*, pp. 105–114 (2017).
28. Wang, X., et al.: ESRGAN: Enhanced super-resolution generative adversarial networks. In: *European Conference on Computer Vision Workshops (ECCVW)*, pp. 63–79 (2018).
29. Feng, C. M., Yan, Y., Fu, H., Chen, L., Xu, Y.: Task Transformer Network for Joint MRI Reconstruction and Super-Resolution. In: de Bruijne, M., et al. (eds.) *MICCAI 2021. LNCS*, vol. 12906, Springer, Cham, pp. 1–16 (2021).
30. Askin Incebacak, N., et al.: Super-resolution reconstruction of T2-weighted thick-slice neonatal brain MRI scans. *J. Neuroimaging* 32(1) (2022).
31. Sui, Y., Afacan, O., Gholipour, A., Warfield, S. K.: Fast & High-Resolution Neonatal Brain MRI Through Super-Resolution Reconstruction From Acquisitions With Variable Slice Selection Direction. *Front. Neurosci.* 15, 636268 (2021).
32. Baljer, L., et al.: Multi-orientation U-Net for Super-Resolution of Ultra-Low-Field Paediatric MRI. [Online]. Available: <https://doi.org/10.1101/2024.02.16.580639>
33. de Leeuw den Bouter, M. L., et al.: Deep learning-based single image super-resolution for low-field MR brain images. *Sci. Rep.* 12, 6362 (2022).
34. He, Y., Nath, V., Yang, D., Tang, Y., Myronenko, A., Xu, D.: SwinUNETR-V2: Stronger Swin Transformers with Stagewise Convolutions for 3D Medical Image Segmentation. In: *LNCS (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 14223 LNCS, pp. 416–426 (2023).
35. Kim, J., Park, H.: Adaptive Latent Diffusion Model for 3D Medical Image to Image Translation: Multi-modal Magnetic Resonance Imaging Study. [Online]. Available: <https://arxiv.org/pdf/2311.00265.pdf>
36. Tapp, A., et al.: MR to CT Synthesis Using 3D Latent Diffusion. In: *2024 IEEE 21st International Symposium on Biomedical Imaging (ISBI)*, Athens, Greece. May 2024.
37. Mei, X., et al.: RadImageNet: An Open Radiologic Deep Learning Research Dataset for Effective Transfer Learning. *Radiol. Artif. Intell.* 4(5) (2022). [Online]. Available: <https://doi.org/10.1148/ryai.210315>
38. Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., Ugurbil, K. for the WU-Minn HCP Consortium: The WU-Minn Human Connectome Project: An overview. *NeuroImage* 80, 62–79 (2013).
39. Valanarasu, J. M., et al.: Disruptive Autoencoders: Leveraging Low-level features for 3D Medical Image Pre-training. *arXiv preprint arXiv:2307.16896* (2023). [Online]. Available: <https://arxiv.org/abs/2307.16896>
40. Avants, B. B., Tustison, N. J., Song, G., Cook, P. A., Klein, A., Gee, J. C.: A reproducible evaluation of ANTs similarity metric performance in brain image registration. *Neuroimage* 54(3), 2033–2044 (2011).
41. Fonov, V., Evans, A. C., Botteron, K., Almli, C. R., McKinstry, R. C., Collins, D. L.; Brain Development Cooperative Group: Unbiased average age-appropriate atlases for pediatric studies. *Neuroimage* 54(1), 313–27 (2011).
42. Wang, L., Wu, Z., Chen, L., Sun, Y., Lin, W., Li, G.: iBEAT V2.0: a multisite-applicable deep learning-based pipeline for infant cerebral cortical surface reconstruction. *Nat. Protocols* 18, 1488–1509 (2023).