Multi-orientation U-Net for Super-Resolution of Ultra-Low-Field Paediatric MRI

Levente Baljer[a], Yiqi Zhang[a], Niall J Bourke[a], Kirsten A Donald[b], Layla E Bradford[b], Jessica E Ringshaw[a,b], Simone R Williams[b], Sean CL Deoni[c], Steven CR Williams[a], Khula SA Study Team[b,1], František Váša[a,2], Rosalyn J Moran[a,2]

[a] Department of Neuroimaging, Kings College London, UK
[b] Department of Paediatrics and Child Health, Red Cross War Memorial Children's Hospital, University of Cape Town, South Africa
[c] Maternal Newborn and Child Nutrition and Health (MNCNH), Bill & Melinda Gates Foundation, Seattle, Washington, USA

[1]A complete list of the Khula SA Study Team can be found in the Supplementary Information
[2]Joint senior authors

Abstract

Purpose

Owing to the high cost of modern MRI systems, their use in clinical care and neurodevelopmental research is limited to hospitals and universities in high income countries. Ultra-low-field systems with significantly lower scanning costs bear the potential for global adoption, however their reduced SNR compared to 1.5 or 3T systems limits their applicability for research and clinical use.

Methods

In this paper, we describe a deep-learning based super-resolution approach to generate high-resolution isotropic T_2-weighted scans from low-resolution inputs. We train a 'multi-orientation U-Net', which uses multiple low-resolution anisotropic images acquired in orthogonal orientations to construct a super-resolved output.

Results

Our approach exhibits improved quality of outputs compared to current state-of-the-art methods for super-resolution of ultra-low-field scans in paediatric populations. The average correlation value between volume estimates from high-field scans and super-resolved outputs rises to 0.77 using our method, compared to 0.71 using earlier techniques.

Conclusion

Our research serves as proof-of-principle of the viability of training deep-learning based super-resolution models for use in neurodevelopmental research and presents the first U-Net trained exclusively on paired ultra-low-field and high-field data from infants.

Introduction

Neuroimaging studies on infant and child populations have become increasingly vital in establishing the relationship between neurodevelopment and resultant cognitive functioning. Magnetic resonance imaging (MRI) has proven to be particularly essential in this endeavour, owing to its ability to provide insight into structural, functional, and metabolic brain development, in addition to recognising pathologies pertaining to neurobiological disorders (Nolen-Hoeksema, 2014). Such powerful capabilities are inherently coupled with major financial costs, however: on top of the cost of a scanner itself (roughly $1,000,000 per Tesla (Arnold, 2023)), the high field strength of existing systems requires facilities with sufficient electromagnetic shielding and specialised staffing. Furthermore, since most modern MRI systems use superconducting magnets, they require the use of cryogens, which themselves come with storage, transportation and maintenance costs (Sarracanie et al., 2015). Altogether, these expenses establish strict financial boundaries on neuroimaging studies and clinical work. Even in hospitals and universities within high-income countries (HICs), scanning costs ($500-$1000/hr per research scan) place a limit on the number of participants scanned and the duration of longitudinal research.

More concerningly from a global health perspective, infrastructural costs severely limit MRI accessibility in low- and middle-income countries (LMICs). Evaluating MRI accessibility based on ratio of MRI units per million population, HICs such as Japan and the US boast 51.67 and 38.96 units/million, respectively. Comparing these values with those of Nigeria and Ghana (which own 0.30 and 0.48 units/million, respectively) reflects a hundred-fold difference between HICs and LMICs (Ogbole et al., 2018). Consequently, our current understanding of neurodevelopment in such regions, where adversities pertaining to nutrition, sanitation and higher rates of infectious risk play a crucial role, is primarily based on psychometric measures or low-cost, functional imaging methods such as EEG or fNIRS (Jensen et al., 2021) (Perdue et al., 2019).

Ultra-low-field (ULF) imaging systems with magnetic field strengths ranging from 50 to 100mT (e.g., the 64mT Hyperfine Swoop) present a potential solution to issues of access inequality (Sarracanie et al., 2015). Such systems rely on the use of large, permanent magnets instead of superconducting electromagnets, and as such have significantly reduced component prices, reduced room requirements, and reduced costs for power, cooling and maintenance

(Campbell-Washburn et al., 2019). Unfortunately, these benefits are paired with a significantly diminished signal-to-noise ratio (SNR), rendering image outputs sub-optimal for visual reading by radiologists or for a large portion of automated processing methods in currently available neuroimaging toolkits. Such processing methods usually require high-resolution MRI scans (1mm isotropic), whereas product sequences on the Hyperfine system yield a default spatial resolution of 1.5×1.5×5 mm. Higher resolution images may be acquired with increased acquisition times (∼12–15 min for a single 2×2×2mm T_2-weighted image), however the former approach is considered more clinically feasible in sensitive populations - such as infants or elderly participants - as acquisition time is reduced to approximately 3-6 minutes per image.

One technique that has found success in enhancing the SNR of ULF outputs is super-resolution (SR) reconstruction of a single image from multiple lower-resolution images, acquired in three orthogonal orientations (i.e., axial, sagittal, and coronal) (Deoni et al., 2022). This is carried out through repeated multi-resolution registration (MRR) of the scans using the Advanced Normalization Tools (ANTs) multivariate template construction tool, outputting a final image with effective dimensions of 1.5×1.5×1.5mm (Figure 1).

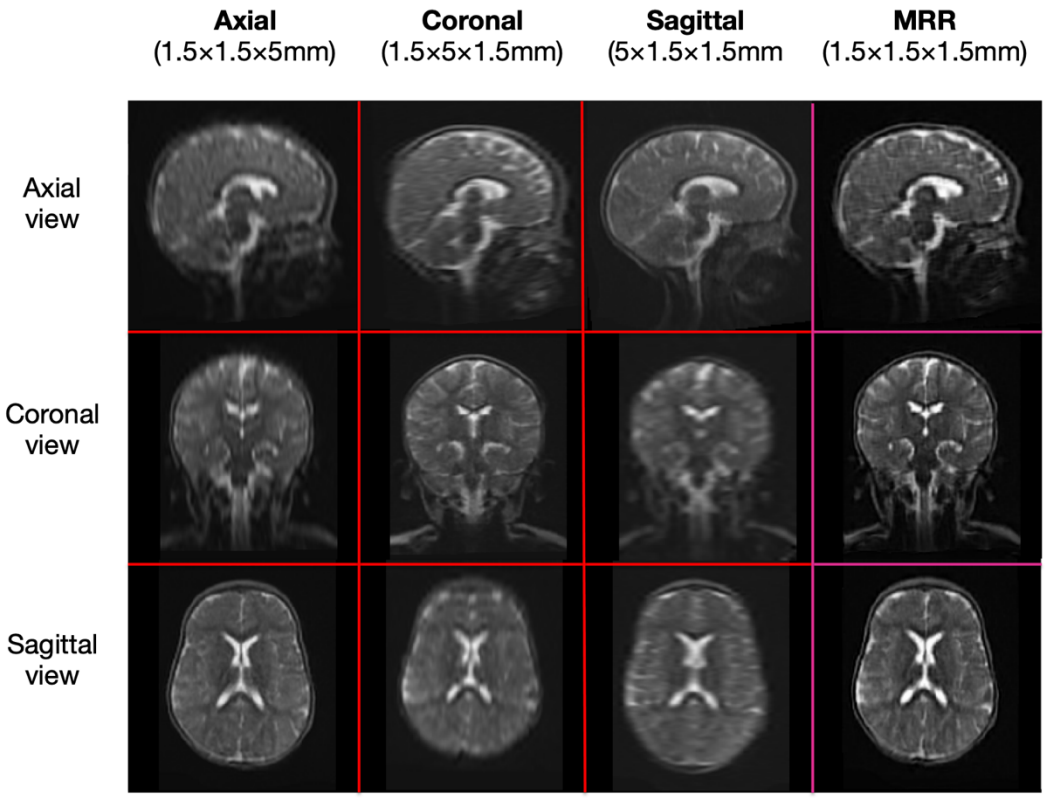

Figure 1) Multi-Resolution Registration (MRR) - Columns 1-3: non-isotropic Hyperfine scans of a 6-month-old subject (1.5 x 1.5 x 5 mm, with high-resolution plane along each of 3 orthogonal orientations); Column 4: isotropic output of MRR (1.5 x 1.5 x 1.5 mm)

An alternative approach involves the use of modern image synthesis methods based on convolutional neural networks (CNNs), which learn a transformation between a source image and a target image. Several architectures have emerged over the years to tackle this problem, one of the earliest being the super resolution convolutional neural network (SRCNN), a three-layer network performing patch extraction, nonlinear mapping of patches, and image reconstruction through each consecutive layer (Dong et al., 2014). More complex architectures were proposed in subsequent years, such as the very deep residual network for Super Resolution (VDSR) (Kim et al., 2015), and the Efficient Sub-Pixel Convolutional Neural Network (ESPCN) (Shi et al., 2016), which was one of the first models to integrate the upsampling operation within the network, making interpolation unnecessary. Several of these networks later found success in the context of super-resolution of 3D brain MRI scans, such as a 3D-SRCNN (Pham et al., 2019) and a modified VDSR (Oktay, 2016). It is important to note, however, that these pioneering architectures were originally developed for single-image super-resolution, not necessarily within a medical domain.

In contrast, the U-Net was originally developed for semantic segmentation of biomedical image data (Ronneberger, 2015) and since its inception in 2015 it has remained one of the most popular architectures for use in this field. The architecture comprises of a contracting path to capture context and a symmetric expanding path that enables precise localization of biological structures. For segmentation, both high-level contextual information and fine-grained spatial information are vital, as the former captures key semantic features of the image (i.e. large-scale anatomical data), whereas the latter allows the output to accurately represent the contents of the original input (i.e. determine whether an individual pixel/voxel is part of a specific brain region). Through the contracting path, this spatial information is normally lost, however skip connections between the two paths ensure feature reusability. This innovative approach is the key behind the U-Net's success in medical image segmentation and is also responsible for its recent success in use for super-resolution of medical images (Iglesias et al., 2021, Iglesias et al., 2023). Here, we investigate whether a U-Net trained on paediatric paired ULF-HF data can surpass MRR in quality of SR outputs.

Methods

Imaging data

MRI data used in this paper stems from a study investigating the neurodevelopment of executive function over the first 3 years of life in participants based in South Africa and Malawi. Images used here included subjects from South Africa exclusively, at ages of either 3- or 6-months. A total of 82 subjects attended two scanning sessions (ULF and HF) at either 3-months or 6-months of age. All subjects had HF scans acquired using a Siemens 3T scanner and had ULF scans acquired using a Hyperfine Swoop, with high in-plane resolution along three orthogonal planes (axial, sagittal, coronal). Of these subjects, 63 successfully completed all four scanning protocols (one HF and three ULF). Pre-processing involved dense-rigid registration between all HF scans and an age-specific template, and each subject's ULF with the corresponding, pre-registered HF scan (Ashburner, 2007). Following this, we skull-stripped all HF and ULF scans separately using HD-BET, a deep-learning tool for MRI brain extraction (Isensee et al., 2019). As seven subjects had poor image outputs following pre-processing steps, our final dataset included 56 paired sets of ULF (64mT T_2; 1.5×1.5×5 mm) and HF (3T T_2; 1×1×1×mm) scans. To maximise the data available from a single subject, we built networks that use three ULF inputs (Hyperfine scans acquired in three, orthogonal orientations) instead

of just one. Consequently, each matched 'pair' of ULF-HF scans included three orthogonal ULF scans and one HF scan. The full pre-processing pipeline and exclusion of images is shown in Figure 2.

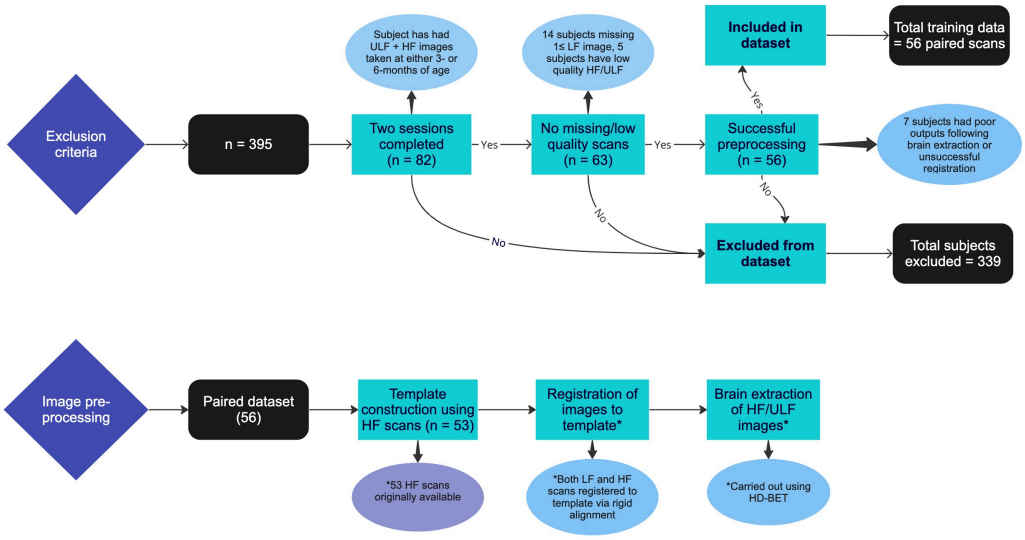

Figure 2) Top row: Participant exclusion; Bottom row: Pre-processing of images prior to being fed into the model.

CNN architecture

Our CNN (Figure 3) builds on the architecture of the 3D U-net (Çiçek et al., 2016). It has three input channels which lead into three encoding levels, each consisting of a convolution with a 3×3×3 kernel (allowing the model to capture fine-grained details while reducing parameter requirements), group normalisation (speeding up convergence by preventing internal covariate shift), a rectified-linear unit (ReLU) activation and maxpooling. The first layer has 64 features, with the number of features doubling after each maxpooling and halving after each upsampling. The final layer utilises a linear activation to produce the final image output. The model is implemented on a GPU using PyTorch.

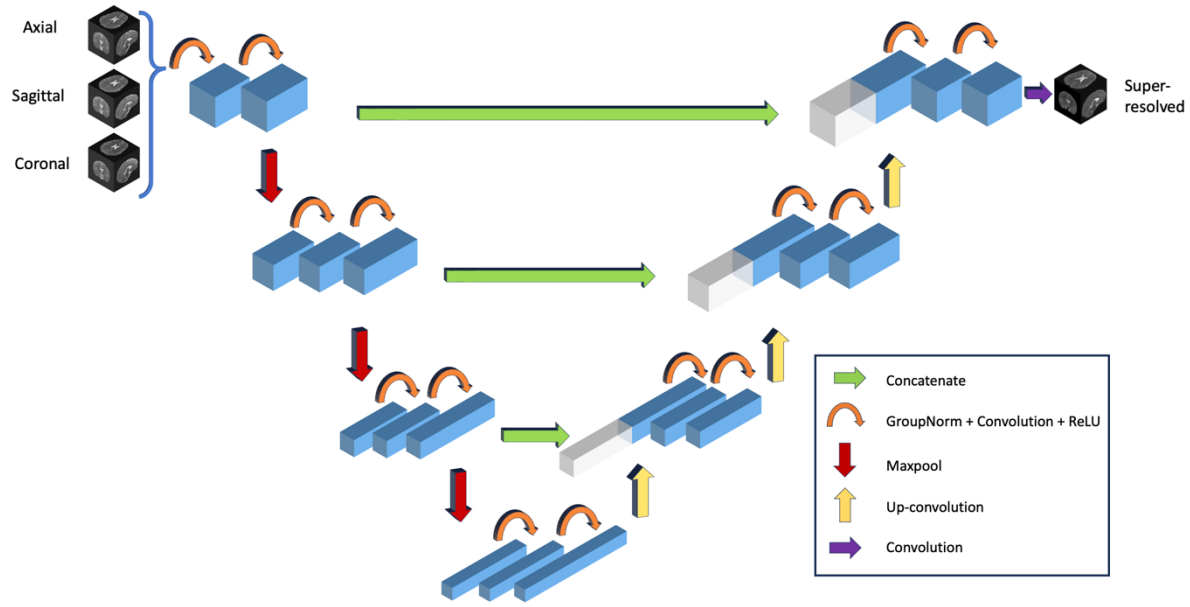

Figure 3) Model architecture: 3D-U-NET with three input channels and one output channel, allowing a transformation from three ULF scans to one SR image. Using this architecture, we aim to maximise the data available from a single subject.

Learning and inference

The training of our network minimises the peak-signal-to-noise ratio with respect to the neural network parameters θ:

$$\mathcal{L}_{PSNR}(\theta; \hat{Y}, Y) = 10 * \log_{10}\left(\frac{MAX_I^2}{MSE}\right)$$

Where $\boldsymbol{MAX_I}$ denotes the maximum image intensity (a pre-defined input parameter) and $\boldsymbol{MSE}$ is the mean-squared error, or:

$$\mathcal{L}_{MSE}(\theta; \hat{Y}, Y) = \sum_{i=1}^{n} (y - \hat{y})^2$$

Where $\hat{\boldsymbol{y}}$ denotes the model output and $\boldsymbol{y}$ denotes the target image.

To assess the dependence of network outputs on the size of the training dataset (i.e. number of matched pairs of HF and ULF scans), we trained 4 models on 10, 20, 30, and 40 paired images, respectively, with the remaining 16 pairs being used to test model outputs. Each U-Net was trained using the Adam optimizer (Kingma, 2014) with a learning rate of 10^{-4} for 1000 epochs (~9 hours for 10-scan model, ~30 hours for 40-scan model using an Nvidia RTX 3090 GPU).

ULF images were normalised and scaled to 1-mm isotropic resolution, ensuring that input and output resolutions match. Augmentation of images was carried out on the fly, with random affine transform or random elastic deformation applied to both ULF and HF data with pre-set probability ($p = 0.5$). After training, weights of the model were frozen and inference was carried out (~10 seconds on an RTX 3090 GPU and ~1 minute on a modern CPU).

Evaluation of model outputs

The quality of model outputs from each of the 4 models relative to "ground-truth" HF images was evaluated via Dice overlap of segmented brain regions within participants and tissue volume correlations across participants, with both segmentation outputs and volume estimations obtained using SynthSeg (Billot, 2020). In addition to segmenting cortical grey matter (GM), cortical white matter (WM) and cerebrospinal fluid (CSF), we segmented the following deep brain structures: accumbens, amygdala, pallidum, hippocampus, caudate, putamen, thalamus, and lateral ventricles. Considering that one of the primary applications of successful SR techniques would be use in ULF-based neurodevelopmental research in LMIC settings, selecting regions that play an important role in neurodevelopmental outcomes is particularly beneficial to underscore the value of our technique. In the case of conditions such as deep nuclear grey matter injury in neonatal hypoxic-ischemic encephalopathy (HIE) (Hassett et al., 2022) or perinatal stroke (Ilves et al., 2022), all regions described are affected in a manner such that damage is associated with worse neurodevelopmental outcomes. To compare how our model fares in capturing these regions compared to current state-of-the-art methods for super-resolution of paediatric ULF scans, we additionally computed the same metrics on images obtained using MRR.

Using our pre-trained models, we additionally assessed performance with varying numbers of input scans. More specifically, we assessed model outputs when the U-Nets only received an axial scan or axial and sagittal scans (as opposed to all three orientations, with which it was originally trained). Since our models require three input volumes, inference with missing scans was achieved by cloning the axial scan either once or twice, depending on whether one or two distinct inputs were provided. As such, we ran inference with each subject a total of three times, with the following combination of scans: 1) axial, axial, axial; 2) axial, axial, sagittal; 3) axial, sagittal, coronal. The choice of prioritising scans in the order of axial >

sagittal > coronal was done based on the scanning protocol used, which involved scanning participants in this exact order.

Statistical methods

Given the small test sample (n=16) we opted for nonparametric statistics to compare populations by using the Wilcoxon signed-rank tests to compare the medians (Wilcoxon, 1945), as opposed to t-tests.

Results

Outputs from our U-Nets yielded visually superior images compared to original ULF scans. Model outputs appeared less noisy, but also smoother than raw ULF scans or MRR outputs (Figure 4).

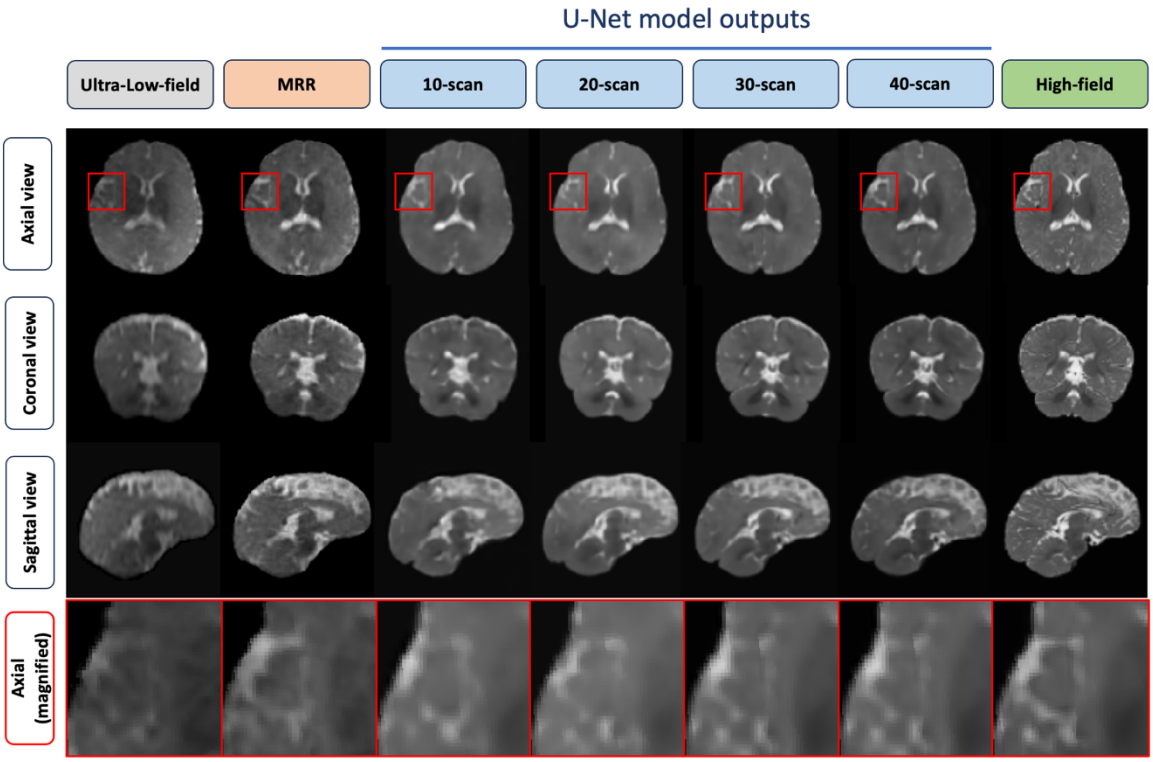

Figure 4) Model outputs from a single out-of-sample subject. Left to right: raw Axial ULF scan, output of multi-resolution registration, outputs of U-Nets as a function of increasing training set size, and the "ground-truth" high-field scan.

We observed a progressive increase in Dice overlap of segmented model predictions with segmented HF scans as the training set increased in size, across segmentation labels (Figure S1). Comparing Dice overlap of model predictions and HF scans to multi-resolution

registration outputs and HF scans, we saw an improvement relative to MRR starting with the 20-scan model, with average Dice coefficients of deep brain structures rising from 0.73 to 0.77 (Figure 5). This value peaked with the 40-scan model at 0.81, reflecting an average increase of 11% relative to MRR.

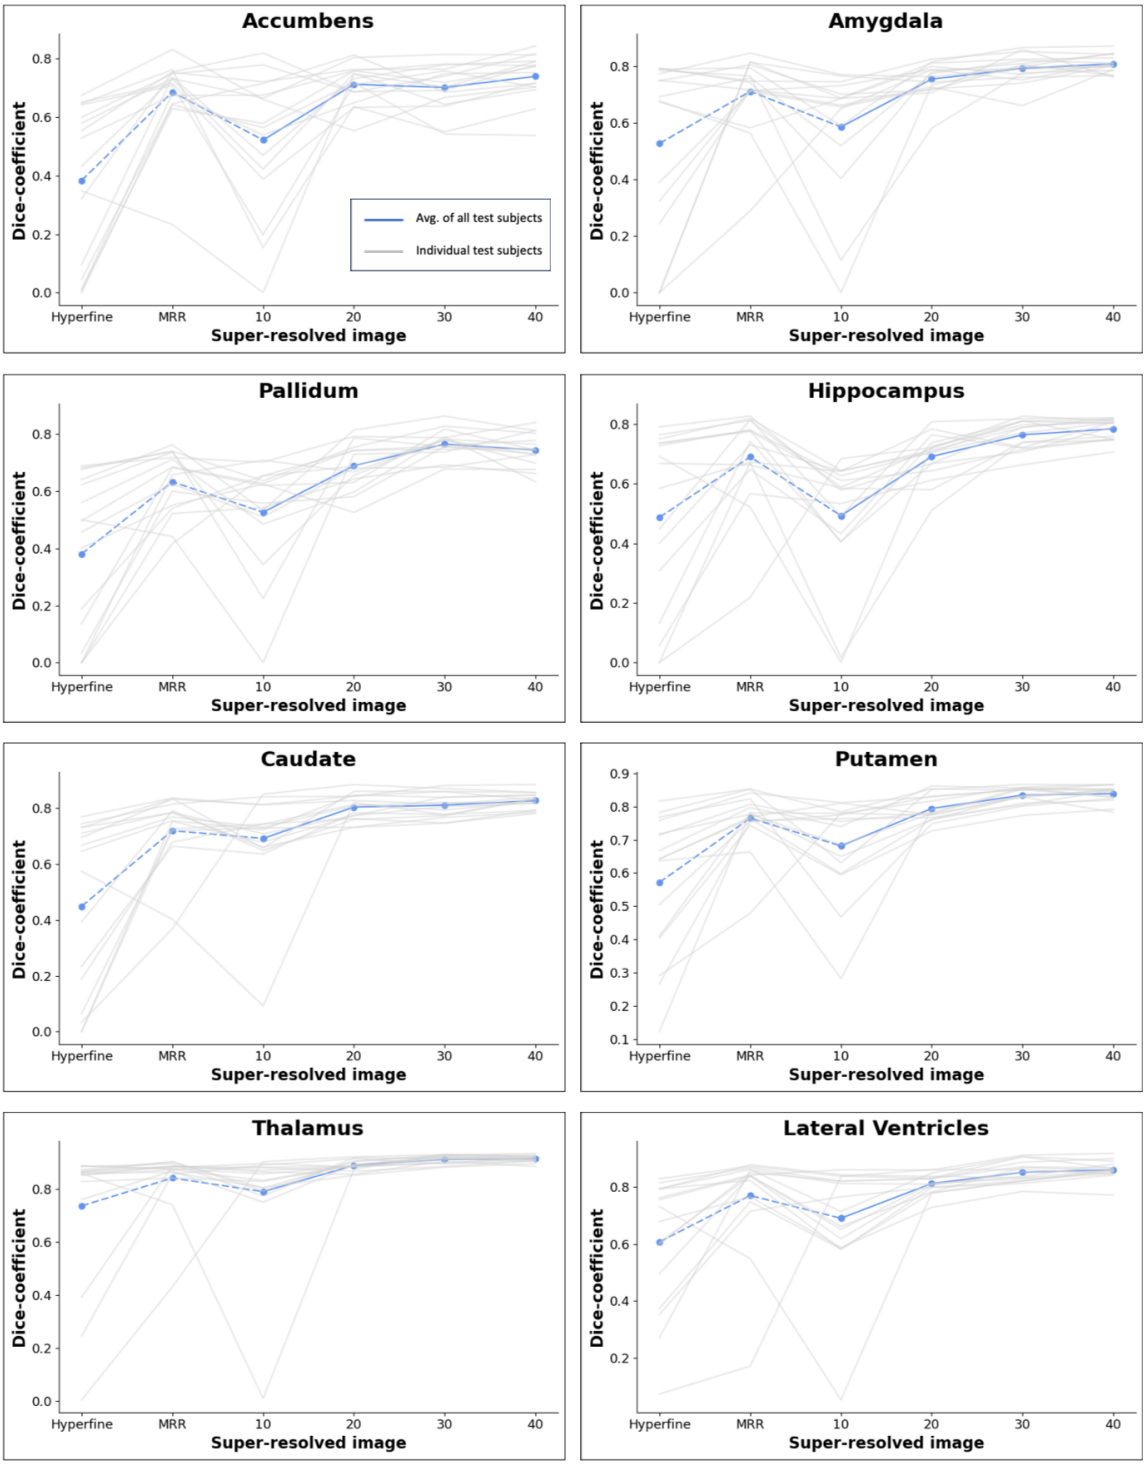

Figure 5) Region-specific Dice coefficients between HF scan and the original Hyperfine scan, MRR output, and outputs of U-Nets as a function of increasing training set. Dice score of original Hyperfine scan is obtained via averaging Dice of axial, coronal, and sagittal scans

A similar pattern was observed when quantifying super-resolution quality by comparing regional volumes between super-resolved outputs and HF scans, across participants (Figure S2). The average Spearman's correlation coefficient of deep brain structures across both hemispheres rose from 0.71 in MRR outputs to 0.77 in predictions of the 40-scan model, reflecting an 8.5% increase. Peak r=0.96 was seen in the thalamus and lateral ventricles (Figure 6 & 7).

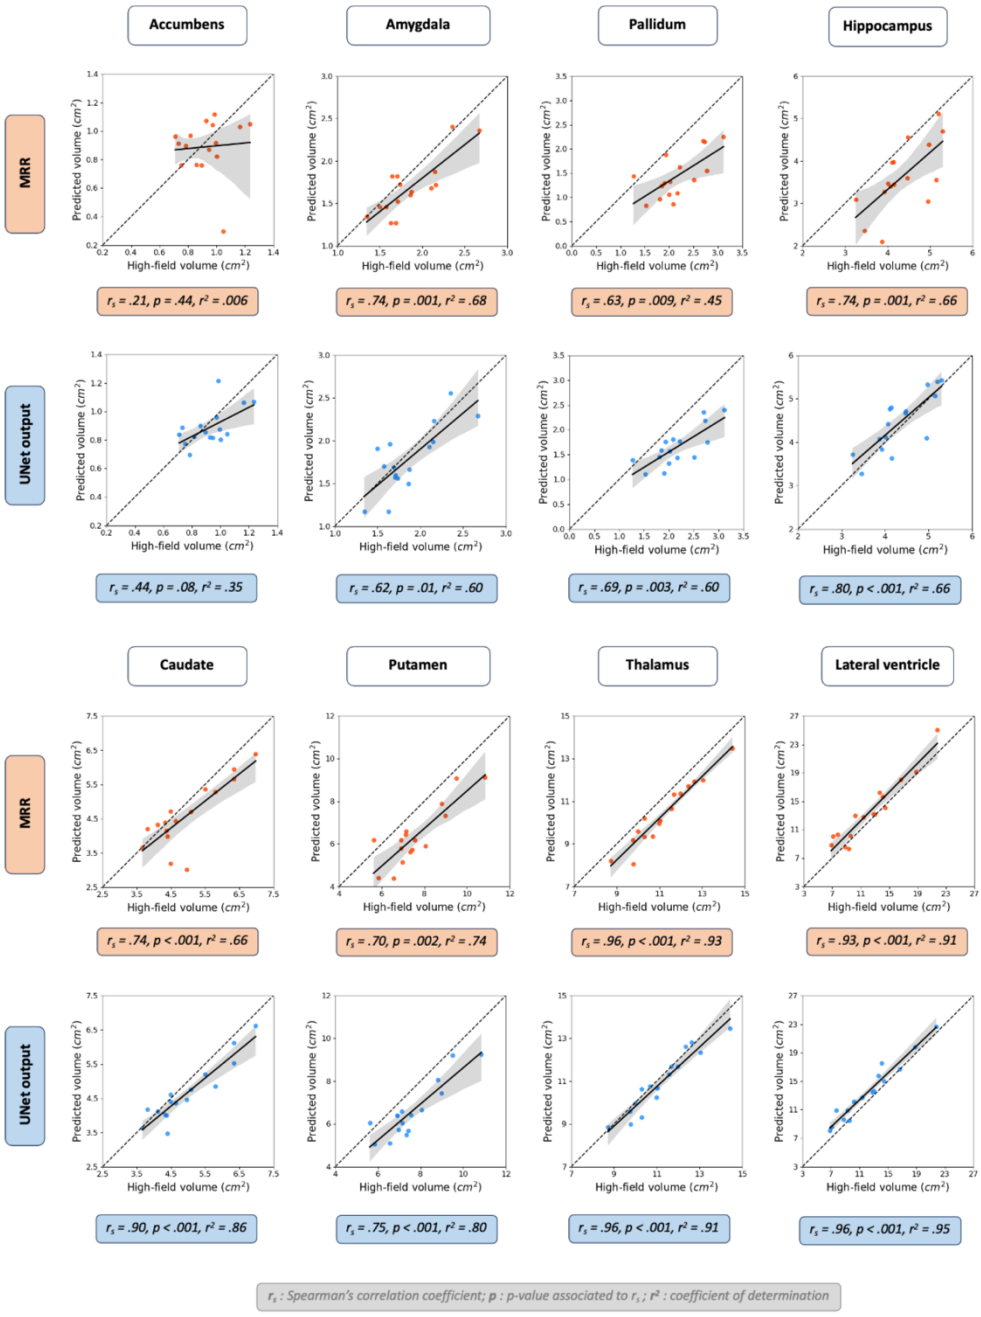

Figure 6) Region-specific volume correlations between MRR outputs and ground-truth HF scans (orange box) vs volume correlations between 40-scan U-Net output and ground-truth HF scan (blue box)

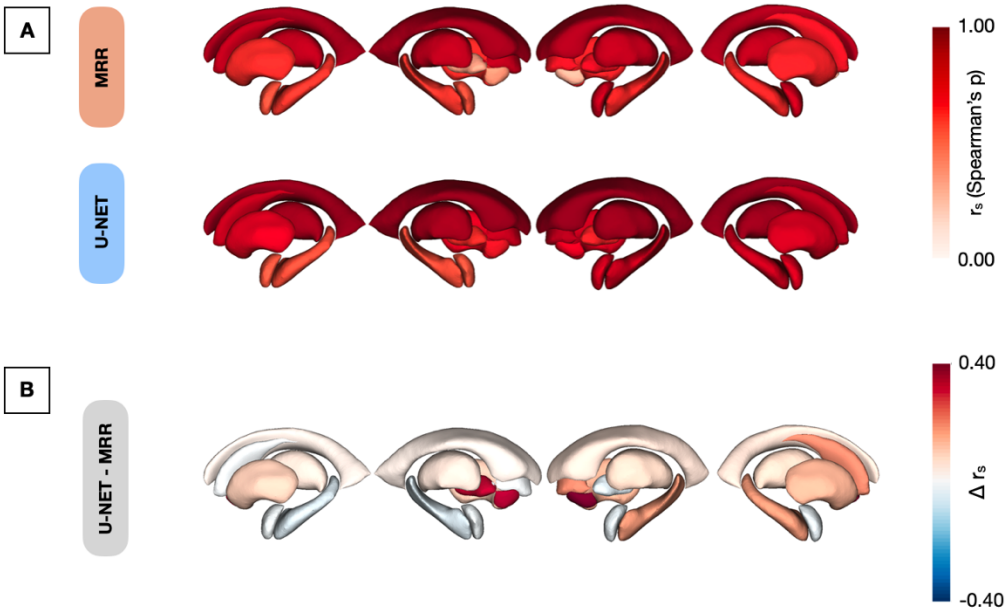

Figure 7) 3D-plot of deep brain structures (left & right accumbens, amygdala, caudate, hippocampus, pallidum, putamen, thalamus and lateral ventricle). A) Volume correlations with HF scans; Top row = MRR outputs, Bottom row = U-Ne outputs. B) Difference in volume correlation values between U-Net outputs and MRR outputs (U-Net - MRR). Images constructed using the ENIGMA toolbox (Larivière et al, 2022)

Significance testing using the Wilcoxon signed-rank test was applied to Dice scores obtained for all sub-cortical regions described above, and additionally to Dice scores for cerebrospinal fluid (CSF), cortical grey matter (GM) and cortical white matter (WM). Given a family-wise error rate (FWER) of 0.0045, significant differences in all subcortical regions were observed when comparing Dice scores from U-Net outputs to those of ULF scans. In contrast, six out of eight regions displayed a significant difference in Dice scores between MRR outputs and ULF scans. Furthermore, when comparing the two SR techniques (our proposed U-Net vs MRR), a significantly greater Dice score was seen in five out of eight regions when using the U-Net. With regards to GM, WM and CSF, Dice scores obtained from MRR outputs did not exhibit any significant difference to those obtained from ULF scans. In contrast, all three of these regions saw significant improvements in Dice scores when using predictions from our U-Net (Table 1).

Table 1) Wilcoxon signed-rank test applied to Dice scores of 16 test subjects, comparing outputs of MRR to Hyperfine, UNet to Hyperfine, and UNet to MRR. For each comparison, the rank biserial correlation (RBC) and significance are displayed.

Region	**MRR > Hyperfine**		**UNet > Hyperfine**		**UNet > MRR**	
Subcortical	*RBC*	*Sig.*	*RBC*	*Sig.*	*RBC*	*Sig.*
Accumbens	0.96	0.000076*	0.99	0.000031*	0.34	0.13
Amygdala	0.51	0.037	0.99	0.000031*	0.82	0.0011*
Pallidum	0.96	0.000076*	1.0	0.000015*	0.87	0.00050*
Hippocampus	0.85	0.00066*	1.0	0.000015*	0.72	0.0046
Caudate	0.87	0.00050*	1.0	0.000015*	0.94	0.00011*
Putamen	1.0	0.000015*	1.0	0.000015*	1.0	0.000015*
Thalamus	0.72	0.0046	1.0	0.000015*	0.99	0.000031*
Lateral ventricles	0.85	0.00066*	1.0	0.000015*	0.86	0.00066*
Cortical	*RBC*	*Sig.*	*RBC*	*Sig.*	*RBC*	*Sig.*
Grey Matter	0.71	0.0055	0.99	0.00031*	0.23	0.23
White Matter	0.60	0.016	1.0	0.000015*	0.46	0.06
CSF	0.60	0.017	0.85	0.00066*	0.0	0.51

When assessing how our models performed using a reduced number of inputs, we observed minimal difference in Dice scores given axial only or axial + sagittal inputs compared to all three orientations in all models trained on >10 paired sets ULF-HF scans. In instances when only one or two input scans were given, a training-set of 30 paired scans was sufficient for the U-Net to supersede MRR in Dice scores for all subcortical and cortical regions (Table 2; Figure S3).

Table 2) Region-specific Dice scores between segmentations obtained from MRR vs. model outputs as a function of increasing training set size. Size of training data is further stratified according to how many unique inputs each model received (A = Axial, AS = Axial + Sagittal, ACS = Axial + Sagittal + Coronal). Grey colouring is used to indicate points at which U-Net outputs exceed MRR outputs in terms of Dice scores. Abbreviations: Subcort. = subcortical, Accumb. = Accumbens, Hip. = hippocampus, Lat. ven. = lateral ventricles.

Region		**10-scan model**			**20-scan model**			**30-scan model**			**40-scan model**		
Subcort.	*MRR*	*A*	*AS*	*ACS*	*A*	*AS*	*ACS*	*A*	*AS*	*ACS*	*A*	*AS*	*ACS*
Accumb.	.69	.37	.34	.52	.67	.71	.71	.69	.69	.70	.71	.72	.74
Amygdala	.71	.34	.34	.59	.76	.76	.75	.78	.77	.79	.78	.80	.81
Pallidum	.63	.44	.45	.53	.66	.67	.69	.76	.77	.77	.72	.73	.74
Hip.	.69	.34	.35	.49	.69	.67	.69	.74	.75	.76	.77	.77	.78
Caudate	.72	.55	.50	.69	.79	.79	.80	.82	.81	.81	.82	.82	.83
Putamen	.76	.57	.56	.68	.79	.80	.79	.84	.84	.83	.82	.83	.84
Thalamus	.84	.60	.67	.79	.89	.89	.89	.91	.91	.91	.91	.91	.91
Lat. ven.	.77	.52	.49	.69	.81	.81	.81	.86	.85	.85	.86	.85	.86
Cortical	*MRR*	*A*	*AS*	*ACS*	*A*	*AS*	*ACS*	*A*	*AS*	*ACS*	*A*	*AS*	*ACS*
GM	.70	.61	.60	.66	.71	.71	.70	.74	.74	.74	.74	.75	.74
WM	.74	.63	.63	.69	.74	.74	.74	.77	.77	.77	.77	.77	.77
CSF	.47	.31	.31	.39	.45	.46	.44	.51	.53	.50	.50	.51	.50

Discussion

Obtaining accurate, high-resolution MRI scans in paediatric populations is essential for studies of neurodevelopment and for diagnosis of neurological conditions. In LMICs where MRI accessibility is significantly reduced, this may be achieved via the use of ULF scanners, supplemented with techniques to improve scan quality. Here we demonstrate improved results compared to previous efforts (Deoni et al., 2022) using a U-Net trained for SR of ULF paediatric images. More specifically, we present greater Dice overlap measures of segmented deep-brain structures and improved volume correlations in these regions. We additionally exhibit the feasibility of training an SR model with a limited number of paired scans, as with 40 paired ULF-HF scans we show success in surpassing MRR on test subjects from the same sample population.

The design of our model assumes that feeding in three separate ULF scans maximises the amount of anatomical information available, allowing the U-Net to more accurately reconstruct a given subject's brain. When assessing Dice scores of deep brain structures from model outputs with varying numbers of unique input scans, however, we see marginal improvement in reconstructions from one scan vs. three unique scans. Tolerance to long scanning times is low in infants (Barkovich et al., 2019), thus successfully completing three separate 4-minute-scanning sessions is often infeasible. As such, having a model that can reconstruct an image from just one ULF scan has major practical benefits. On the other hand, the relatively small improvement in image quality going from one to three input scans suggests that encoding layers may not be extracting features at the single-image level, and instead the model treats these scans in unison. To encourage image-specific encodings, future work may involve modifying the architecture to have one full encoding branch for each input image, where these separate encodings are then combined before they enter the expanding path (Lau et al 2022).

A limitation of our analyses is that we only assess model performance on unseen subjects of the same age from the same scanning site. To fully assess the limits of our model, we would run inference on ULF scans from subjects of varying ages and alternative sites, investigating the magnitude of deviation from the training set that results in a sharp decline in model performance. Additionally, including test subjects with pre-specified neurobiological conditions/lesions would allow us to investigate whether using deep learning-based SR runs

the risk of failing to capture pathologies. This would provide insight as to whether a model trained on normative images can be applied to subjects with potential pathologies, or if disease-specific models need to be trained to ensure reliable SR outputs.

An additional limitation stems from the use of SynthSeg for image segmentation (Billot et al., 2023). Seeing as SynthSeg was trained on adult data, its application to paediatric images runs the risk of producing erroneous volume estimations. Owing to differing water and fat content in the developing brain compared to adults, alterations in signal intensities in newborns and infants are seen with both T_1- and T_2-weighted anatomical MRI. In particular, the first 6 months of life are characterised by a reversal of the normal adult contrast (T_1-weighted: lower WM intensity than GM; T_2-weighted: higher WM intensity than GM) (Dubois et al., 2021). These neurodevelopmental changes amplify the drastic difference already present in T_1 and T_2 contrast between 3T and ULF, and as such a more accurate approach to testing model performance would involve the use of a segmentation model trained specifically on paediatric scans.

Finally, although the U-Net has proven to be the gold-standard for many deep-learning applications to medical imaging, the use of novel architectures such as diffusion models or visual transformers may yield further improvements in quality of model outputs. These methods have the potential to address critical limitations of previous techniques, enhancing overall realism in super-resolved images. Within the domain of single-image super-resolution (SISR) of medical modalities, both diffusion models and visual transformers have been shown to outperform CNN and GAN-based models in terms of Structural Similarity Index Measure (SSIM) and Peak Signal-to-Noise Ratio (PSNR) (Kazerouni et al., 2023) (Zhu et al., 2023).

Conclusions

Ultra-low-field (ULF) imaging presents a potential paradigm-shift in neuroimaging, however its ingress into widespread research and clinical use is impeded by limitations on resolution and the range of available contrasts. Here we demonstrate how the use of CNNs can aid in deriving high-resolution T_2-weighted scans of paediatric patients from ULF scanners. More specifically, we demonstrate the feasibility of training U-Nets with a limited number of

paired ULF-HF scans, such that their outputs outperform current non-deep-learning based SOTA methods for super-resolving paediatric scans.

Funding information

Artificial Intelligence Methods for Low Field MRI Enhancement, Bill and Melinda Gates Foundation (INV- 032788). Wellcome Leap 1kD programme (The First 1000 Days) [222076/Z/20/Z].

Conflict of interest

The authors report no significant financial conflicts of interest with respect to the subject matter of this manuscript.

References

Arnold, T.C., Freeman, C.W., Litt, B. and Stein, J.M. (2023), Low-field MRI: Clinical promise and challenges. *J Magn Reson Imaging, 57*: 25-44. https://doi.org/10.1002/jmri.28408

Ashburner J. (2007). A fast diffeomorphic image registration algorithm. *Neuroimage* 38, 95–113 10.1016/j.neuroimage.2007.07.007

Barkovich MJ, Li Y, Desikan RS, Barkovich AJ, Xu D. Challenges in pediatric neuroimaging. *Neuroimage*. 2019 Jan 15;185:793-801. doi: 10.1016/j.neuroimage.2018.04.044.

Billot, B., et al. (2020). A Learning Strategy for Contrast-agnostic MRI Segmentation. Proceedings of the Third Conference on Medical Imaging with Deep Learning. *Proceedings of Machine Learning Research, PMLR*. 121: 75--93

Campbell-Washburn, A. E., et al. (2019). "Opportunities in Interventional and Diagnostic Imaging by Using High-Performance Low-Field-Strength MRI." *Radiology* 293(2): 384-393.

Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O. (2016). 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016. *Springer, Cham*. https://doi.org/10.1007/978-3-319-46723-8_49

Deoni SCL, O'Muircheartaigh J, Ljungberg E, Huentelman M, Williams SCR. 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.* 2022; 88(3): 1273-1281. doi:10.1002/mrm.29273

Dong, C., Loy, C.C., He, K., Tang, X. (2014). Learning a Deep Convolutional Network for Image Super-Resolution. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds) Computer Vision – ECCV 2014. ECCV 2014. Lecture Notes in Computer Science, vol 8692. *Springer, Cham*. https://doi.org/10.1007/978-3-319-10593-2_13

Dubois, J., Alison, M., Counsell, S.J., Hertz-Pannier, L., Hüppi, P.S. and Benders, M.J.N.L. (2021), MRI of the Neonatal Brain: A Review of Methodological Challenges and Neuroscientific Advances. *J Magn Reson Imaging*, 53: 1318-1343. https://doi.org/10.1002/jmri.27192

Hassett, J., Carlson, H., Babwani, A., & Kirton, A. (2022). Bihemispheric developmental alterations in basal ganglia volumes following unilateral perinatal stroke. *NeuroImage: Clinical, 35*, 103143. https://doi.org/10.1016/j.nicl.2022.103143

Iglesias, J. E., Billot, B., Balbastre, Y., Tabari, A., Conklin, J., Gilberto González, R., . . . Fischl, B. (2021). Joint super-resolution and synthesis of 1 mm isotropic MP-RAGE volumes from clinical MRI exams with scans of different orientation, resolution and contrast. *NeuroImage*, 237, 118206. https://doi.org/10.1016/j.neuroimage.2021.118206

Iglesias, J. E., et al. (2023). "Quantitative Brain Morphometry of Portable Low-Field-Strength MRI Using Super-Resolution Machine Learning." *Radiology* 306(3): e220522.

Ilves, N., Männamaa, M., Laugesaar, R., Ilves, N., Loorits, D., Vaher, U., . . . Ilves, P. (2022). Language lateralization and outcome in perinatal stroke patients with different vascular types. *Brain Lang, 228,* 105108. doi:10.1016/j.bandl.2022.105108

Isensee F, Schell M, Tursunova I, Brugnara G, Bonekamp D, Neuberger U, Wick A, Schlemmer HP, Heiland S, Wick W, Bendszus M, Maier-Hein KH, Kickingereder P. Automated brain extraction of multi-sequence MRI using artificial neural networks. *Hum Brain Mapp*. 2019; 1–13. https://doi.org/10.1002/hbm.24750

Jensen, S. K. G., Xie, W., Kumar, S., Haque, R., Petri, W. A., & Nelson, C. A., 3rd. (2021). Associations of socioeconomic and other environmental factors with early brain development in Bangladeshi infants and children. *Dev Cogn Neurosci, 50,* 100981. doi:10.1016/j.dcn.2021.100981

Kazerouni, A., et al. (2023). "Diffusion models in medical imaging: A comprehensive survey." *Medical Image Analysis* 88: 102846.

Kim, J., Lee J. K., Lee, K. M. "Accurate Image Super-Resolution Using Very Deep Convolutional Networks," *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR),* Las Vegas, NV, USA, 2016, pp. 1646-1654, doi: 10.1109/CVPR.2016.182.

Kingma, D.P., Ba, J., Adam: A method for stochastic optimization. *arXiv*:1412.6980 [cs.LG] (22 December 2014).

Larivière, S., Paquola, C., Park, By. et al. The ENIGMA Toolbox: multiscale neural contextualization of multisite neuroimaging datasets. *Nat Methods* 18, 698–700 (2021). https://doi.org/10.1038/s41592-021-01186-4

Lau V, Xiao L, Zhao Y, et al. Pushing the limits of low-cost ultra-low-field MRI by dual-acquisition deep learning 3D superresolution. *Magn Reson Med*. 2023; 90: 400-416. doi: 10.1002/mrm.29642

Nolen-Hoeksema, S. (2014). *Abnormal psychology*. New York, NY, McGraw-Hill Education.

Ogbole GI, Adeyomoye AO, Badu-Peprah A, Mensah Y, Nzeh DA. Survey of magnetic resonance imaging availability in West Africa. *Pan Afr Med J*. 2018 Jul 31;30:240. doi: 10.11604/pamj.2018.30.240.14000

Oktay, O. et al. (2016). Multi-input Cardiac Image Super-Resolution Using Convolutional Neural Networks. Medical Image Computing and Computer-Assisted Intervention - MICCAI 2016, *Springer, Cham*. https://doi.org/10.1007/978-3-319-46726-9_29

Perdue, K. L., Jensen, S. K. G., Kumar, S., Richards, J. E., Kakon, S. H., Haque, R., . . . Nelson, C. A. (2019). Using functional near-infrared spectroscopy to assess social information processing in poor urban Bangladeshi infants and toddlers. *Dev Sci*, 22(5), e12839. doi:10.1111/desc.12839

Pham, C.-H., Tor-Díez, C., Meunier, H., Bednarek, N., Fablet, R., Passat, N., & Rousseau, F. (2019). Multiscale brain MRI super-resolution using deep 3D convolutional networks. *Computerized Medical Imaging and Graphics, 77,* 101647. https://doi.org/10.1016/j.compmedimag.2019.101647

Ronneberger, O., Fischer, P., Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, *Springer, Cham*. https://doi.org/10.1007/978-3-319-24574-4_28

Sarracanie, M., LaPierre, C. D., Salameh, N., Waddington, D. E. J., Witzel, T., & Rosen, M. S. (2015). Low-Cost High-Performance MRI. *Scientific Reports*, 5(1), 15177. doi:10.1038/srep15177

Shi, W. et al., "Real-Time Single Image and Video Super-Resolution Using an Efficient Sub-Pixel Convolutional Neural Network," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR),* Las Vegas, NV, USA, 2016 pp. 1874-1883. doi: 10.1109/CVPR.2016.207

Wilcoxon, Frank (Dec 1945). "Individual comparisons by ranking methods". *Biometrics Bulletin. 1* (6): 80–83. doi:10.2307/3001968

Zhu, J., et al. (2023) A residual dense vision transformer for medical image super-resolution with segmentation-based perceptual loss fine-tuning. *arXiv*:2302.11184 DOI: 10.48550/arXiv.2302.11184

Supplementary Information

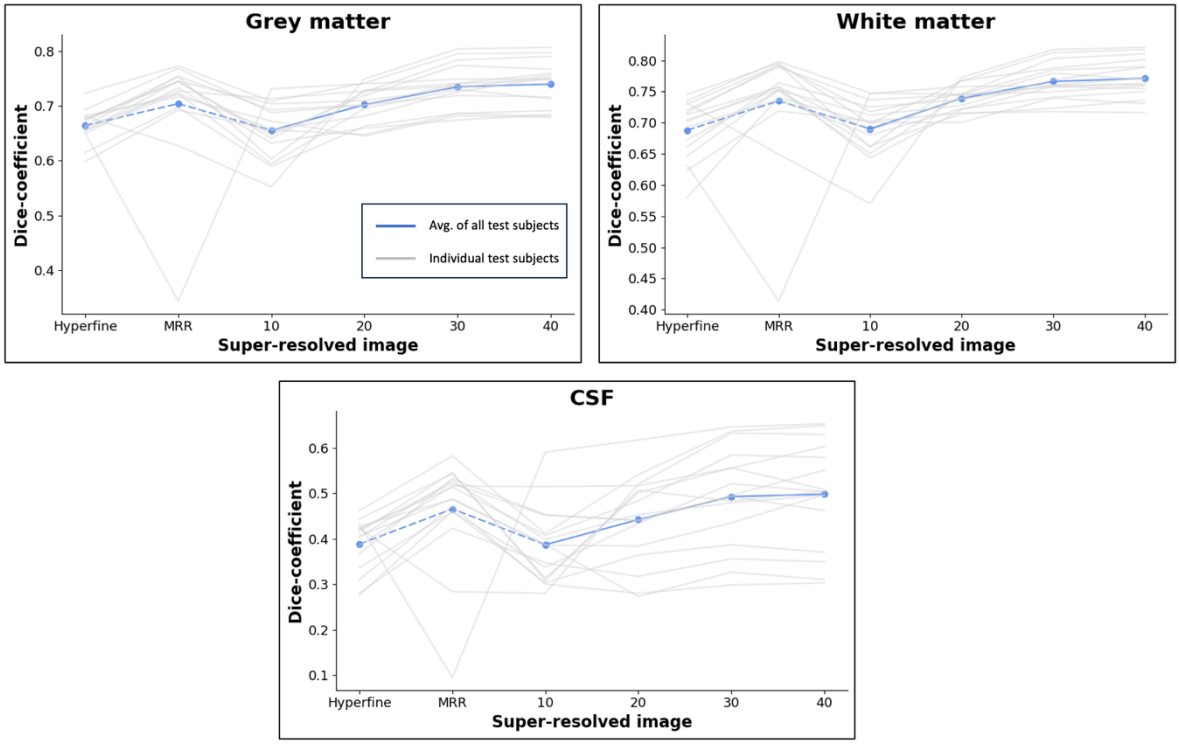

Figure S1) Region-specific Dice coefficients between HF scan and the original Hyperfine scan, MRR output, and outputs of U-Nets as a function of increasing training set. Regions depicted: Grey Matter, White Matter and Cerebrospinal fluid (CSF)

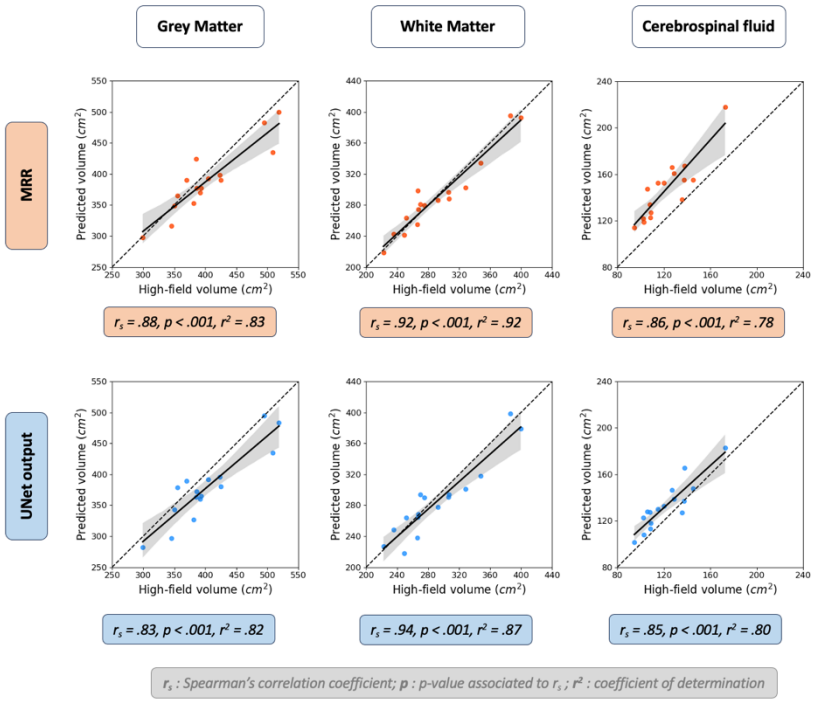

Figure S2) Region-specific volume correlations between MRR outputs and ground-truth HF scans (orange box) vs volume correlations between 40-scan U-Net output and ground-truth HF scan (blue box). Regions depicted: Grey Matter, White Matter, Cerebrospinal fluid

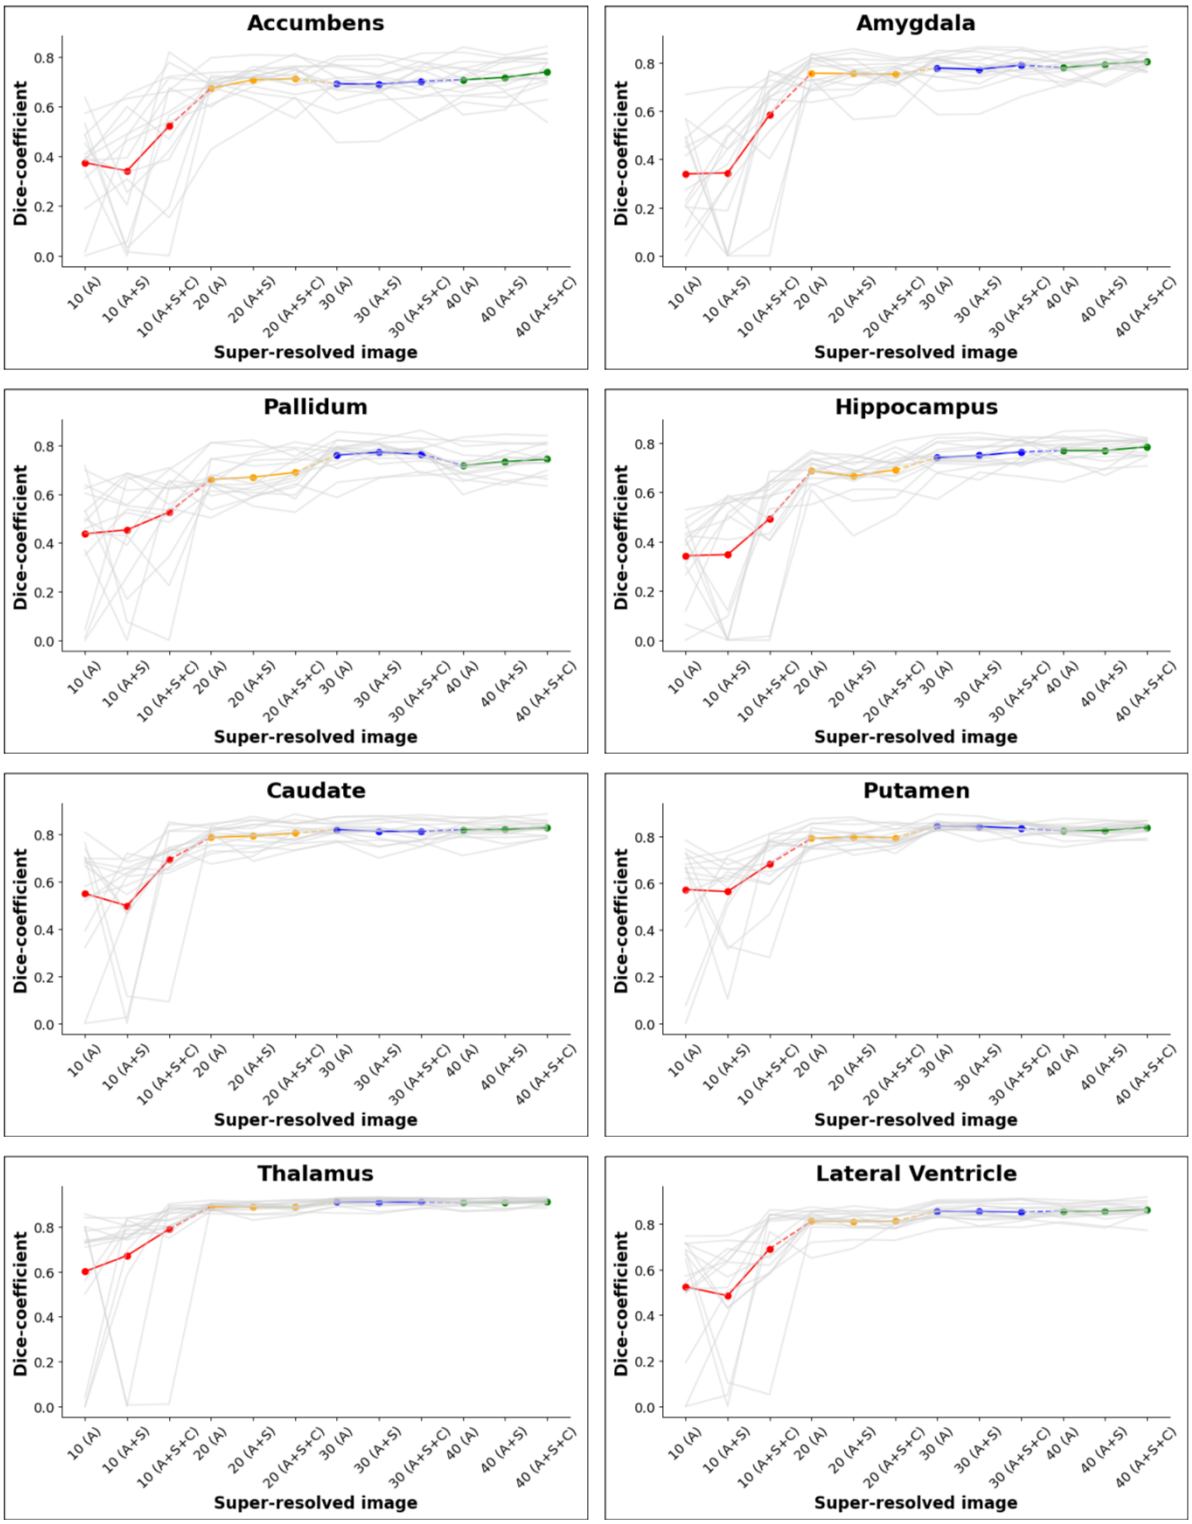

Figure S3) Region-specific Dice coefficients between HF scan and outputs of U-Nets as a function of increasing training set and increasing number of inputs given to the models. A = axial, A+S = axial and sagittal, A+S+C = axial, sagittal and coronal.

Khula SA Study Team

Michal R. Zieff
Donna Herr
Chloë A. Jacobs
Sadeeka Williams
Zamazimba Madi
Nwabisa Mlandu
Tembeka Mhlakwaphalwa
Lauren Davel
Reese Samuels
Zayaan Goolam
Thandeka Mazubane
Bokang Methola
Khanyisa Nkubungu
Candice Knipe