Boosting the Signal-to-Noise of Low-Field MRI With Deep Learning Image Reconstruction

Neha Koonjoo (✉ nkoonjoo@mgh.harvard.edu)

Massachusetts General Hospital

Bo Zhu

Massachusetts General Hospital

Cody Bagnall

Texas A&M University

Matthew Rosen

Massachusetts General Hospital

Research Article

Keywords: Signal to noise, Image Reconstruction, AUTOMAP

Posted Date: December 17th, 2020

DOI: <https://doi.org/10.21203/rs.3.rs-126917/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Version of Record: A version of this preprint was published at Scientific Reports on April 15th, 2021. See the published version at <https://doi.org/10.1038/s41598-021-87482-7>.

Abstract

Recent years have seen a resurgence of interest in inexpensive low-field (<0.3 T) MRI systems mainly due to advances in magnet, coil and gradient set designs. However, most of these advances are focused on hardware development and signal acquisition while far less attention has been given to how advanced image reconstruction can improve image quality at low field. We describe here the use of our end-to-end deep neural network approach (AUTOMAP) to improve the image quality of highly noise-corrupted low-field MRI data. We compare the performance of this approach to two additional state-of-the-art denoising pipelines. We find that AUTOMAP improves image reconstruction of data acquired on two very different low-field MRI systems: human brain data acquired at 6.5 mT, and plant root data acquired at 47 mT, demonstrating SNR gains above Fourier reconstruction by factors of 1.5- to 4.5-fold, and 3-fold, respectively. In these applications, AUTOMAP outperformed both contemporary denoising algorithms and suppressed noise-like spike artifacts in reconstructed images. The impact of domain-specific training corpora on the reconstruction performance is discussed. The AUTOMAP approach to image reconstruction will enable significant image quality improvements at low-field, especially in highly noise-corrupted environments.

Introduction

Introduction

MRI scanners that operate in the low magnetic field regime (i.e., < 0.3 Tesla) as a rule suffer from reduced image quality that arises from the low Boltzmann polarization at these field strengths resulting in weak NMR signals¹. Accordingly, images obtained at low-field suffer from low signal-to-noise ratio (SNR) which can be mitigated in part by increased acquisition times². The situation is far worse in the militesla regime including efforts at ultra-low field (ULF) (<10 mT) where, in addition to extremely small NMR signals, spatial encoding gradient field strengths are usually weak (<10 mT/m), limiting attainable spatial resolutions on these systems compared to those at high field³. Despite these challenges, there has been significant recent interest in ULF MRI as a low-cost strategy for increasing the availability and worldwide accessibility of MRI scanners⁴. The ability to acquire images at ULF has been bolstered by contemporary hardware developments including high-performance magnets, RF coils, and gradient sets, and well as improved strategies for data acquisition and new image processing algorithms^{2,5,6}. However, as part of this evolving arsenal of tools, the introduction of recently developed machine learning (ML) approaches has not yet impacted low-field MRI, namely noise-robust machine learning-based image reconstruction.

At higher MR field strengths, the application of ML to the problem of transforming the acquired (raw) data to final image (known as image reconstruction) has led to a host of technical improvements including rapid (of order milliseconds) reconstruction of high-dimensional raw data, improvement in image quality and image contrast, including for highly-undersampled data sets, and dynamic and multi-channel image reconstruction⁷⁻¹⁴. To date, none of these SNR-boosting ML approaches have been applied to real noisy MRI data acquired at very low magnetic field, a regime with unique challenges in acquisition and

reconstruction. MR imaging coils must maximize coverage over the volume of interest while minimizing losses, and at low field (i.e., low Larmor frequency), these losses are mainly due to resistive losses in the coil (the so-called Johnson noise regime) as body/sample noise is negligible. RF coil designs using parallel imaging technology are limited in their ability to accelerate imaging at low magnetic field due to a general need to signal average; trading acceleration for a loss in SNR is a zero sum game in terms of acquisition time in this regime. While minimizing losses during signal acquisition will improve image quality and boost SNR, image reconstruction with noise-robust machine-learning approaches can also contribute to low field imaging.

In this paper, we investigate the capability of deep learning-based image reconstruction approaches to address the low-SNR challenges of low-field MRI. We focus in particular on the use of the AUTOMAP (Automated transform by Manifold Approximation) neural network framework¹⁰ and describe the SNR boost and image artifact removal that AUTOMAP provides in comparison to other widely-used image denoising techniques. Various deep learning denoising algorithms have substantially gained attention due to their flexible neural network architectures, less need to manually set parameters, and more generalizable denoising problem solving^{15–19}. However, these neural networks have primarily been used for denoising in the spatial- and not the signal-domain, and are focused mainly on Gaussian and other idealized noise distributions due to a general dearth of real-world low-field imaging data.

The AUTOMAP method recasts image reconstruction as a supervised learning task, whereby the relationship between two data domains, sensor (or k -space) and image, emerges during supervised learning from a training corpus generated using the known forward encoding model. Hence, AUTOMAP directly learns the spatial decoding transform between the k -space and the image space. Once trained, the feed-forward network operates between sparse manifolds making the reconstruction highly robust to noise and other artifacts. As described in¹⁰, the neural network architecture consists of two fully connected layers followed by sparse convolutional layers as shown in **Fig. 1a**.

Raw data was acquired from two different MRI systems operating at low magnetic field: a 6.5 mT human brain scanner operating in our laboratory³ and a 47 mT plant root imager²⁰ designed to operate outdoors underground in clay soil. The 6.5 mT ULF scanner was used to acquire image data of both phantoms and human heads. The 47 mT scanner was purpose built to image sorghum root architecture as it grows in natural soils²⁰.

In this paper, we assess the reconstruction performance of AUTOMAP on these different low field MRI datasets. For the 6.5 mT scanner, image SNR was evaluated for AUTOMAP reconstructing 2D phantom data and 3D *in vivo* brain data, as compared to the standard inverse Fast Fourier Transform (IFFT) reconstruction of the same data. We then compared the SNR boost from using AUTOMAP to jointly denoise and reconstruct the raw data with a compartmentalized approach that first reconstructed with IFFT and then used the two state-of-the-art image-domain-only denoising pipelines – 1) a deep convolutional neural network gaussian noise denoiser¹⁸ and 2) the Block-matching and 3D filtering (BM3D) denoising algorithm²¹. The ability of AUTOMAP to remove system imperfections such as noise

intrusion from the raw data was also assessed and compared to a denoising pipeline operating in image-domain.

The role of training corpus in the reconstruction performance was then explicitly studied by reconstructing raw data acquired on the 47 mT scanner of sorghum roots with AUTOMAP trained on each of three different training sets: 1. brain images, 2. synthetic vascular tree-like images, and 3. synthetic plant root images. Reconstruction performance was evaluated in each case compared with the standard IFFT reconstruction method.

Results

2D structured phantom MR Imaging at 6.5 mT

In order to characterize image SNR improvement in a well-controlled environment, a phantom experiment was carried out, where a single slice image was acquired with different number of averages (NA): 40, 60, 80, 100, 200, 300, 400, 500, 600, 700, and 800 respectively.

The 2D water-filled phantom raw datasets were reconstructed with AUTOMAP trained on T_1 -weighted MRI brain images. The SNR of the reconstructions were compared with the conventional IFFT reconstruction method. As observed in **Fig. 1 (b and c)**, all the reconstructed slices with AUTOMAP exhibit a lower noise and overall image quality is improved such that small features become observable even at low number of averages. The overall mean SNR was evaluated and plotted in **Fig. 1d** - left axis. The SNR of AUTOMAP reconstructed images were higher than the SNR of IFFT reconstructed images. The mean gain in SNR was of 62% with NA=40 (**Fig. 1d** - right axis) and plateaus at NA=800 with less than 5% gain compared to the IFFT. Taking the reference image as that acquired with the maximum NA, the graph in **Fig. 1e** shows a significant reduction in the RMSE at lower NA compared to higher NA. The decrease in RMSE for NA less than 200 is around 35% to 40% and around 10% to 20% for higher NA. The error maps (Fig. S1) also depicted a significant decrease in error in the AUTOMAP reconstructed 2D images with respect to the reference IFFT image at NA=800, compared to the IFFT reconstructed 2D images.

3D human brain MR Imaging at 6.5 mT

Healthy volunteer human subjects were scanned under IRB in our 6.5 mT scanner with a 3D-bSSFP sequence. The raw k -space data was reconstructed with either AUTOMAP or IFFT. The AUTOMAP reconstructed images of a 11-minute scan with 50 averages of one of the subjects are shown in **Fig. 1f**. In comparison with the IFFT reconstruction method (**Fig. 1g**), the AUTOMAP reconstructed images have lower noise floor and also more clarity in brain structures like the ventricles. The reference slices acquired with 100 averages in 22 minutes are shown in **Fig. 1h**. Mean SNR improvements in the brain region of 1.5 to 4.5 times higher were obtained across slices.

As observed in both the phantom and brain datasets, the AUTOMAP approach significantly reduces the noise floor in the reconstructed images as a consequence of it operating as transformation between

learned sparse manifolds¹⁰.

In this light, we wanted to compare the end-to-end reconstruction and denoising performance of AUTOMAP with 2 denoising algorithms. The first denoiser is a state-of-the-art image-only deep learning denoising approach (DnCNN)^{18,19} (recently incorporated as a built-in MATLAB function `denoiseImage`), which utilizes a deep single-scale convolutional neural network in a residual learning context to perform the denoising. The second denoiser is the well-established BM3D algorithm with collaborative filtering²¹ (http://www.cs.tut.fi/~foi/GCF-BM3D/bm3d_matlab_package_3.0.7.zip). DnCNN and BM3D were both applied to the (magnitude) images obtained from reconstruction using IFFT and AUTOMAP. A single illustrative slice from the NA=50 reconstructed brain dataset from **Fig. 1g** is displayed in **Fig. 2 (a-d)**, reconstructed with either AUTOMAP (**Fig. 2a**) or IFFT reconstruction (**Fig. 2b**). To improve the display of the background noise, the images are shown with the window-level adjusted to more conveniently evaluate the low pixel values. The additional DnCNN or BM3D image denoising step was applied to the image data reconstructed with IFFT (**Fig. 2c and d**). The denoised IFFT image in **Fig. 2 (c and d)** seem to indicate that both DnCNN and BM3D act more on the background and not the signal.

The denoised images were evaluated using image quality metrics and compared to the those reconstructed without the additional DnCNN/BM3D denoising stage. Results are shown in **Fig. 2 (e-h)**. The mean overall SNR within the head region is plotted in **Fig. 2e**. The mean SNR for AUTOMAP is higher by ~2 folds than both the IFFT + DnCNN or IFFT + BM3D. The PSNR (**Fig. 2f**) is also ~ 2.0 dB higher for the AUTOMAP as compared to the IFFT alone. The PSNR of IFFT + BM3D is improved by 1.0 dB over IFFT alone; 1.0 dB of the improvement in the PSNR is due to AUTOMAP. The RMSE is lower for AUTOMAP compared to the IFFT + DnCNN or IFFT + BM3D (**Fig. 2g**). The same conclusion applies for SSIM analysis, where the similarity index is higher for images reconstructed with AUTOMAP alone than for those reconstructed via IFFT + DnCNN or IFFT + BM3D images (**Fig. 2h**). This imaging and analysis were repeated for a second human subject and similar results were obtained (data not shown). The table in **Fig. 2i** summarizes the mean metric values across all the 15 slices and shows that AUTOMAP alone provides a mean SNR gain of 3.1 with a maximum at 4.5.

In addition to the improvement in SNR and robustness to noise, the use of AUTOMAP has the potential to reduce or eliminate certain imaging artifacts (**see Fig. 3, a-d**) which might arise from a low-field scanner operating in an unshielded environment or other real-world non-ideal imaging scenarios. A human head scan acquired at 6.5 mT in 22 min (NA = 50) in our laboratory (**Fig. 3a**) inadvertently contained a spiking artifact caused by noisy gradient amplifier switching that can be seen as structured noise in the phase-encode direction across all slices when reconstructed with IFFT. A reference scan (NA=100) was acquired with the scanner operating normally (i.e, with no artifact), and error maps between the reference and the NA=50 scan also depict clearly the spike artifact in **Fig. 3 (a and b)**. This artifact was eliminated when the same dataset was reconstructed with AUTOMAP as seen in **Fig. 3(c and d)**.

As a way to evaluate the ability of AUTOMAP to reduce *k*-space spike artifacts, we randomly generated high-intensity *k*-space samples, which were added to the hybrid *k*-space data to simulate spike corrupted

multi-slice k -space data (**Fig. 3h**). This data was reconstructed with the same trained model used earlier (see Methods for details). As described in ¹⁰, the AUTOMAP reconstruction approach may be seen as a transformation between sparse manifolds that are conditioned during training. In this case, since the training corpus was based on clean, artifact-free brain MRI data, spike artifacts would be detected as non-signal noise-like features outside the training corpus and be suppressed during the reconstruction. The AUTOMAP reconstruction of the spike corrupted k -space (**Fig. 3i**) is mostly artifact-free compared to the highly-corrupted images as can be seen in the IFFT reconstruction (**Fig. 3k**), although there are still some residual artifact-like structures as quantified in the table described in **Fig. 3n**. We note that the non-corrupted images reconstructed with IFFT are shown in **Fig. 3f and 3g**.

To further improve the immunity of the learned approach to spiking artifacts, we trained AUTOMAP on a new training corpus that included random spike corruption as described in Methods. The spike corrupted k -space data (**Fig. 3h**) was reconstructed with this new trained AUTOMAP and as seen in **Fig. 3j**, and the reconstructed images are artifact-free, appearing identical to the uncorrupted k -space ground truth reconstructed with AUTOMAP (**Fig. 3e**). We found that by training AUTOMAP on a corpus that includes forward encoding of classes of corruption that are representative of those likely to be found in the experimental data, reconstruction performance is greatly improved, reducing or even eliminating artifacts in the final image.

A sparse transform that is conditioned on properties of its training corpus operates very differently from a purely image-based denoising method. As shown in **Fig. 3l and 3m**, the DnCNN denoising and the BM3D denoising applied following the transform from k -space to image-space are not able to significantly reduce spike artifacts beyond blurring them. Quantitative image quality metrics measurements (in **Fig. 3n**) corroborate this. In particular, the SSIM values of the spike-corrupted dataset reconstructed using the standard gaussian noise-trained AUTOMAP are slightly lower than the SSIM values of the clean dataset reconstructed using the same model. The SSIM metrics of the spike-corrupted dataset reconstructed using the spike- and gaussian-noise trained AUTOMAP model are identical to the spike-free dataset reconstructed with the standard gaussian noise-trained AUTOMAP. In summary, AUTOMAP was able to eliminate most spiking artifacts when trained on a standard corpus without prior knowledge of artifacts, and once trained on a corpus that included spike-like signals, AUTOMAP completely removed spike artifacts, with a reconstruction performance equal to that of the uncorrupted data set.

Plant root MR Imaging at 45 mT

Recent prior work has demonstrated the feasibility of studying the root phenotype of sorghum growing in natural clay soils with low field MRI ²⁰. With the aim of further improving root images acquired from the 47 mT cylindrical electromagnet-based scanner described therein, raw k -space from several 2D root projection imaging experiments were reconstructed with AUTOMAP and the images were compared with the conventional IFFT reconstruction. In addition to evaluating the performance of AUTOMAP to reconstruct root images compared to IFFT, we also investigate use of different training corpora (**Fig. 4, a-c**) and the impact of each training corpus on the reconstructed images. **Fig. 4d** summarizes the list of the

different resolution of the acquired root datasets and the corresponding training set used for reconstruction.

First, 48 x 48 root k -space datasets were reconstructed using the “brain trained” model, i.e., where the training set was composed of T_1 -weighted MRI brain images that were downsampled to match the 48x48 matrix size. The results of the reconstructed 2D data in **Fig. 5a** show consistent improvement in SNR with AUTOMAP across all the datasets. AUTOMAP was able to reconstruct the images accurately, without any visible distortion or information loss as the standard IFFT images as seen in **Fig. 5d**, and with better overall SNR of the roots. Enhancements of 1.37 to 1.53 were observed in the different images as stated in the **Fig. 5e**. The noise floor in all AUTOMAP reconstructed datasets are considerably reduced.

We note here that by visual inspection, the sulcus-like structures in the MRI brain image data (see **Fig. 4a**) are not present in the linear and branching features of the root images. As a first attempt to generate a more domain-appropriate training corpus for roots that would enhance feature extraction and improve the image quality of the root reconstruction, we naïvely drew visual comparison to the natural branched structure of blood vessels, and leveraged existing synthetic vascular tree models to generate training images (**Fig. 4b**). This vascular tree-based trained model is described in the Methods and was evaluated by measuring the mean overall SNR of the images and comparing them with the IFFT method. The results are shown in **Fig. 5b**, where a noticeable reduction in noise is seen in all the six root datasets. The SNR enhancement evaluated in those images are reported in the **Fig. 5e**. A further improvement in SNR was obtained with SNR gains of 2.3- to 2.6-fold higher than the IFFT reconstruction. Along with the decrease in noise floor, we also observed that the small signal arising from soil water is attenuated. This filtering of soil signal can eventually help in extracting specific root measurements using existing root data analysis toolbox.

Our third training corpus for root reconstruction uses a specialized root architecture toolbox – *RootBox* – to generate synthetic root structures, and this was used to train AUTOMAP (<https://github.com/Plant-Root-Soil-Interactions-Modelling/CRootBox>). Some representative root reconstructions are shown in **Fig. 4c**, and these reconstructed images in **Fig. 5c** appear to be nearly indistinguishable from those reconstructed with vascular tree-based model. **Fig. 5e** summarizes the results as SNR and SNR enhancements for all three trained models and clearly shows that the domain specificity of the synthetic vascular tree and the synthetic roots promotes the reduction in the noise floor in the images and takes advantage of the sparse features of those two training sets.

Roots datasets acquired with a higher spatial resolution of 0.83 mm were also reconstructed using synthetic root trained AUTOMAP model since the roots appeared to be more densely packed root crown structures. All the eight projections from a root dataset were reconstructed with AUTOMAP and a significant reduction in noise floor can be observed in **Fig. 6a** as compared with IFFT in **Fig. 6b**. SNR improvements of 2-fold or higher than IFFT were evaluated. **Fig. 6c** summarizes the SNR analysis of both the AUTOMAP reconstruction and the IFFT reconstruction over the eight projections.

Discussion

In this work, we have shown that our end-to-end deep learning-based image reconstruction approach improves reconstruction of SNR-starved MRI images acquired at low magnetic field. From the structured phantom images and human brain images acquired at 6.5 mT, to plant root images acquired at 47 mT, we observed significantly improved image SNR in all cases with AUTOMAP reconstruction. The well-established relationship between SNR and scan time allows us to utilize the gain in SNR from AUTOMAP to enable a reduction in scan time of at least 2-fold. Indeed, the results from the 2D structured phantom shows that as the number of averages decreases, the AUTOMAP reconstruction approach is able to extract more features from the input k -space as compared to the conventional IFFT method. Also, we should note that the reconstruction of the phantom data was done using a model that was trained on brain images. Even though brain structures exhibit different visual features than the phantom, the model learned the relationship between the k -space and the image and reconstructed the data with improved SNR and also shows robustness to noise¹⁰.

Furthermore, in the *in vivo* human experiment, training the neural network without the addition of white gaussian noise in the k -space shows mean SNR improvements of 1.3 to 1.54-fold corresponding to 1.0 to 1.5 dB increase in the PSNR (data not shown). However, the addition of random white gaussian noise (10 to 35 dB) in each k -space data in the training database shows a further improvement in the output images with significant noise reduction. As per the SSIM values, we did not observe any loss of information compared to the reference image dataset.

State-of-the-art denoising methods like BM3D¹⁷, WNNM²² or TNRD²³ which are based on nonlocal self-similarity image prior models are popular and effective denoisers that generally act on the reconstructed magnitude images. However, these approaches typically employ time-consuming iterative optimizations that suffer from generalization to multiple noise levels. In 2016, Zhang K. et al proposed a deep convolutional neural network denoising (DnCNN) approach that uses a residual network architecture to learn the noise and subtract it from the image. This work showed that the state-of-the BM3D with collaborative filtering as well as the DnCNN approach can attenuate the additive white Gaussian noise operating purely on the image (magnitude) data for a range of noise levels, and that DnCNN has been evaluated as one of the standard machine learning methods to denoising magnitude MR images¹⁹. We compared DnCNN-denoised low field images and BM3D-denoised low field images to AUTOMAP reconstructed images, in particular on the *in vivo* brain dataset. Because BM3D and DnCNN are both limited to reducing noise only in the image domain, its reconstructed images generally exhibited over-blurred features. We also note that the statistical properties of the noise profile changed after denoising. On the contrary, AUTOMAP, which operates directly on the raw data to solve the inverse problem operating between learned sparse manifolds, is more robust to noise at a fundamental level. From the quantitative analysis, we see that AUTOMAP reconstructed images have a higher PSNR and SSIM values than denoised IFFT images.

Inverting the raw k -space data using AUTOMAP has another interesting advantage in cases involving high frequency spike-corrupted datasets. The reconstruction of spike-corrupted data demonstrated significant suppression of artifacts, which were outside the training corpus and thus treated as non- k -space data to be reduced, or as shown above, essentially eliminated from the reconstructed image. In addition, we observed significant improvement in image accuracy with improved SSIM values as compared with the denoising pipeline.

As already mentioned above, both image-based denoising pipelines reduce gaussian noise in the image domain, but mainly blurs the entire image²⁴. On the contrary, AUTOMAP learns that spike artifacts in the sensor domain do not contain any 'signal-like' or "image-like" information and hence tries to discard them. Furthermore, if training is performed on an artifact-based database, where the neural network learns the relationship between a high frequency corrupted k -space and a clean image, we show that the SSIM, the PSNR and the RMSE values of the reconstructed images with corrupted k -space match those images with clean k -space data. We also note that the artifact-based trained model works better for high spatial frequency spikes. Besides threshold-based spike correction techniques, the Fourier-Transform based approach developed by Yi-Hsuan Kao and his team is an effective spike removal algorithm even for low spatial frequencies²⁵. However, some conditions need to be fulfilled for optimal results, for example, 2 spikes have to be at least 7 pixels apart. Another spike removal technique, the Principal Component Analysis algorithm-based approach that acts directly on the k -space data, separating the sparse low-rank spike information, is one that successfully identifies high frequency RF spikes efficiently. However, the sparsity penalty parameter needs to be adjusted for every dataset for optimal results^{26,27}. Using residual learning can be one of the solutions where the neural network learns the relationship between the corrupted k -space and the subtracted image between the clean and corrupted one^{28,29}.

Besides training on T_1 -weighted brain images, we investigated domain-specific training on plant root MRI datasets acquired on a 47 mT magnet²⁰. Reconstruction of the plant roots datasets is feasible with a T_1 -weighted brain trained model; however, the low-dimensional representation of plant root images is different than that of brain images. The use of a model trained on either vascular tree structures or synthetic roots images provided better SNR improvements. The performance with both training sets were similar and qualitatively the AUTOMAP reconstructed images look the same as the IFFT reconstructed images. Unfortunately, due to lack of ground truth images, SSIM and PSNR were not evaluated. Scaling up AUTOMAP reconstruction to 96 × 96 roots dataset clearly shows the robustness to noise even in conditions where we have higher SNR. The significant improvement in SNR will eventually help in feature extraction of the roots. The difference between roots structures and noise is enhanced and thereby postprocessing analyses such as segmentation would work more effectively.

In conclusion, we have applied the end-to-end AUTOMAP ML-based image reconstruction approach to different ULF datasets and showed the significant improvement in SNR and image quality without any loss of information. We have demonstrated its application to spike-corrupted data and have shown the impact of the training set on the reconstruction. As low field MRI is rapidly becoming a modality of

interest for developing countries for more sustainable and affordable medical care³⁰, we believe that employing noise-robust machine learning image reconstruction approaches will be an important component in shortening scan times and improving image quality to accelerate the practical use and adoption of low-field MRI.

Materials And Methods

2D Phantom - Data Acquisition at 6.5 mT

2D data at 6.5 mT (276 kHz ¹H frequency) was acquired with a balanced Steady State Free Precession (b-SSFP) sequence of a hemispheric water-filled structured resolution phantom placed inside a single-channel spiral volume head coil described previously³. The sequence parameters were TR = 31ms, matrix size = 64 x 64, spatial resolution = 4.2mm x 3.6mm, and slice thickness = 12mm. Several imaging sequences were run, each with a different number of averages (NA): 40, 60, 80, 100, 200, 300, 400, 500, 600, 700, and 800 respectively. The total acquisition time of each dataset is shown in **Fig. 1, b-c**.

3D *In vivo* brain - Data Acquisition at 6.5 mT

The same single-channel spiral volume head coil described above was used to acquired 3D human brain data. A 3D b-SSFP sequence was used: TR = 31 ms, matrix size = 64 x 75 x 15, spatial resolution = 2.5 mm x 3.5 mm x 8 mm, and 50% under-sampled was performed along both along the phase-encode and the slice direction. Two *in vivo* datasets were collected: 1. a 11-min scan with NA=50 and 2. a 22-min scan with NA=100. The 22-min scan was used as the reference scan.

To simulate spike/herringbone corrupted *k*-space data, we altered the raw *k*-space data of the 11-min scan by multiplying the signal with a scaling factor. The scaling factor varied from 2 to 20 and the number of spikes, which were randomly spread over the *k*-space ranges from 2 to 10.

Sorghum root imaging at 47 mT

2D projection images were acquired from rhizotron cores²⁰. The 2D projection images were acquired using a 2D spin-warp pulse sequence. Sixteen spin echoes with an echo spacing of 7-ms were generated and the image acquisition time for each 2D projection in this experiment was 15 minutes. The sequence parameters are as follows: TR = 500 ms, Field-of-View = 80 mm, matrix size = 48x48, and spatial in-plane resolution = 1.64 mm. Different roots data with matrix size 72 x 72 and spatial resolution of 1.11 mm or matrix size of 96 x 96 and spatial resolution of 0.83 mm were acquired for visualizing finer root structures.

AUTOMAP training dataset for phantom and brain imaging at 6.5 mT:

Two training corpora were assembled from 50,000 2D T₁-weighted brain MR images selected from the MGH-USC Human Connectome Project (HCP) public database³¹. The images were cropped and were

subsampled to either 64 x 64 (for the phantom dataset) or 75 x 64 (for the brain dataset), symmetrically tiled to create translational invariance. Random additive white gaussian noise (ranging from 10 dB - 35 dB) was applied to each image in the training set. To produce the corresponding *k*-space for training, each image was Fourier Transformed with MATLAB's native 2D FFT function. The neural network was trained from the noise corrupted *k*-space encodings and target 'noise-free' images to learn an optimal feed-forward reconstruction of *k*-space domain into the image domain.

Spike/herringbone artifact AUTOMAP training dataset for brain imaging at 6.5 mT

An artifact-based 75 x 64 training set was generated using the same HCP dataset. The *k*-space inputs were corrupted with randomly distributed spikes outside of the central 17% of *k*-space. The number of spikes varied from 1 to 25 and its magnitude was multiplied randomly by a factor between 20 to 300. Additive white gaussian noise was also added in the *k*-space data (ranging from 15 dB-35 dB). The neural network was hence trained from the noise corrupted spiky *k*-space encodings and target 'noise-free' images using a neural network architecture as described below.

AUTOMAP training dataset for sorghum imaging at 47 mT

AUTOMAP was trained on the Fourier forward-encoding model using three different training corpora and each model was used to reconstruct the low-field plant root raw data. The performance of AUTOMAP on those three training sets was evaluated.

Brain training set - The first training set used was built from the same HCP database as mentioned above, using 51,000 2D T_1 -weighted MRI brain images (**shown in Fig. 4a**) downsampled to matrix size 48 x 48. The images in each training set were symmetrically tiled to create translational invariance and finally normalized to the maximum intensity of the data. To produce the corresponding *k*-space representations for training, each image was Fourier Transformed with MATLAB's native 2D FFT function.

Vascular training set - The second training set was assembled from 55,000 2D synthetic vasculature images using the Vascular Tree Synthesis Software, VascuSynth³². The first step was to generate 3D volumes of the random synthetic vessels that were 48 x 48 x 48 with random branching sizes, numbers and flow rates. Then 2D images of the vasculature were obtained by randomly slicing the volumes and summing in either the first, the second or the third dimension, resulting into 55000 images in total. Examples from the training set are illustrated in **Fig. 4b**. Random additive white gaussian noise (ranging from 10 dB-35 dB) was applied to each image in the training set. Each image was symmetrically tiled to create translational invariance and finally normalized to the maximum intensity of the data. To produce the corresponding *k*-space representations for training, each image was Fourier Transformed with MATLAB's native 2D FFT function. The neural network was trained from the noise corrupted *k*-space encodings and target 'noise-free' images to learn an optimal feed-forward reconstruction of *k*-space domain into the image domain.

RootBox training set - The third corpus was assembled from 75,000 2D synthetic 48 x 48 root images. The roots images were generated using a 3D root system growth model implemented with MATLAB - called RootBox³³ (see Fig. 4c). Random additive white gaussian noise (ranging from 10dB-35dB) was applied to each image in the training set. The images were symmetrically tiled to create translational invariance and finally normalized to the maximum intensity of the data. To produce the corresponding k -space representations for training, each noise-corrupted image was Fourier Transformed with MATLAB's native 2-D FFT function. The neural network was trained from the noise corrupted k -space encodings and target 'noise-free' images to learn an optimal feed-forward reconstruction of k -space domain into the image domain. RootBox training sets were also generated for matrix size 96 x 96.

Architecture of Neural Network

The neural network was trained from the noise corrupted k -space encodings and target 'noise-free' images to learn an optimal feed-forward reconstruction of k -space domain into the image domain. The same network architecture and hyperparameters were used for our experiments as previously described in¹⁰. The real and the imaginary part of datasets were trained separately. The network was composed of 2 fully connected layers (input layer and 1 hidden layer) of dimension $n^2 \times 1$ and activated by the hyperbolic tangent function. The 2nd layer was reshaped to $n \times n$ for convolutional processing. Two convolutional layers convolved 64 filters of 3x3 with stride 1 followed each by a rectifier nonlinearity. The last convolution layer was finally de-convolved into the output layer with 64 filters of 3x3 with stride 1. The output layer resulted into either the reconstructed real or imaginary component of the image.

For the spike- and noise- corrupted k -space database, the neural network architecture was composed of 2 fully connected layers of dimension $n^2 \times 1$ and activated by the hyperbolic tangent function. The 2nd layer was reshaped to $n \times n$ for convolutional processing. Three convolutional layers convolved 128 filters of 3x3 with stride 1 followed each by a rectifier nonlinearity. Then the last convolution layer was finally de-convolved into the output layer with 128 filters of 3x3 with stride 1. As stated above the real and imaginary components were trained separately and hence the output layer resulted into either the reconstructed real or imaginary component of the image.

Training details

Multiplicative noise at 1% was also applied to the input to learn robust representations from corrupted inputs^{10,34}. The RMSProp algorithm was used with following parameters: minibatches size = 100, learning rate = 0.0001, momentum = 0.0, and decay = 0.9. The loss function and the L1 norm penalty applied on the convolutional layer during training were kept the same as described in the previous work¹⁰. The network was trained for 100 epochs on the Tensorflow deep learning framework using 2 NVIDIA Tesla GV100 GPUs with 16GB memory capacity each.

3D Image reconstruction with AUTOMAP

The *in vivo* raw datasets of each slice were reconstructed with either AUTOMAP or IFFT. Due to memory limitation of the network architecture of AUTOMAP, we explicitly applied a 1D FFT along the partition direction of the 3D k -space, and then applied AUTOMAP to the resultant hybrid-space data to reconstruct images slice-by-slice.

2D Image reconstruction with AUTOMAP

The raw 2D k -space datasets from all samples were stacked and multiplied by a scalar so the range of signal intensities lies within that of the corresponding training models above. The stacked k -space were then reconstructed with the corresponding trained model.

Image Analysis

The signal magnitude of each dataset was normalized to unity to enable fair comparison between both reconstruction methods. SNR was then computed by dividing the signal magnitude by the standard deviation of the noise (obtained from a background ROI of the image-space data). For the *in vivo* brain datasets, error maps were computed using the 22-min scan as the reference image. Image quality metrics were evaluated using RMSE (root mean square error), as well as in-built MATLAB functions - PSNR (Peak SNR), and SSIM (Structure Similarity Index for Measuring image quality).

Declarations

Acknowledgments

We acknowledge support from the Advanced Research Projects Agency-Energy (ARPA-E), U.S. Department of Energy, under Award Number DE-AR0000823 for contributing to fund the information, data, or work presented herein. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. We also acknowledge support for this work from GE Healthcare.

Author contributions:

N.K., B.Z. and M.S.R. performed the phantom and human experiments and conceived the initial research study. G.C.B. acquired the plant root datasets. All authors interpreted results. N.K. wrote the manuscript, with significant input and critical revisions from all authors.

Competing interests:

All authors in this article declare to have no competing interests.

MSR is a co-founder of Hyperfine Research, Inc.

Materials and correspondence:

References

1. Waddington, D. E. J., Boele, T., Maschmeyer, R., Kuncic, Z. & Rosen, M. S. High-sensitivity in vivo contrast for ultra-low field magnetic resonance imaging using superparamagnetic iron oxide nanoparticles. *Sci. Adv.***6**, eabb0998 (2020).
2. Marques, J. P., Simonis, F. F. J. & Webb, A. G. Low-field MRI: An MR physics perspective. *J. Magn. Reson. Imaging***49**, 1528–1542 (2019).
3. Sarraçanie, M. *et al.* Low-Cost High-Performance MRI. *Sci. Rep.***5**, 15177 (2015).
4. 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) doi:10.1001/jamaneurol.2020.3263.
5. Ginde, A. A., Foianini, A., Renner, D. M., Valley, M. & Camargo, Jr, C. A. Availability and Quality of Computed Tomography and Magnetic Resonance Imaging Equipment in U.S. Emergency Departments. *Acad. Emerg. Med.***15**, 780–783 (2008).
6. Cooley, C. Z. *et al.* Two-dimensional imaging in a lightweight portable MRI scanner without gradient coils. *Magn. Reson. Med.* (2015) doi:10.1002/mrm.25147.
7. Qin, C. *et al.* Convolutional recurrent neural networks for dynamic MR image reconstruction. *IEEE Trans. Med. Imaging* (2019) doi:10.1109/TMI.2018.2863670.
8. Hyun, C. M., Kim, H. P., Lee, S. M., Lee, S. & Seo, J. K. Deep learning for undersampled MRI reconstruction. *Phys. Med. Biol.***63**, (2018).
9. Hammernik, K. *et al.* Learning a variational network for reconstruction of accelerated MRI data. *Magn. Reson. Med.* (2018) doi:10.1002/mrm.26977.
10. Zhu, B., Liu, J. Z., Cauley, S. F., Rosen, B. R. & Rosen, M. S. Image reconstruction by domain-transform manifold learning. *Nature***555**, 487–492 (2018).
11. Qin, C. *et al.* k-t NEXT: Dynamic MR Image Reconstruction Exploiting Spatio-Temporal Correlations. in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 505–513 (2019). doi:10.1007/978-3-030-32245-8_56.
12. Chen, F. *et al.* Variable-density single-shot fast spin-echo MRI with deep learning reconstruction by using variational networks. *Radiology* (2018) doi:10.1148/radiol.2018180445.
13. Ramzi, Z., Ciuciu, P. & Starck, J. L. Benchmarking MRI reconstruction neural networks on large public datasets. *Appl. Sci.* (2020) doi:10.3390/app10051816.
14. Ravishankar, S., Ye, J. C. & Fessler, J. A. Image Reconstruction: From Sparsity to Data-Adaptive Methods and Machine Learning. *Proc. IEEE* (2020) doi:10.1109/JPROC.2019.2936204.
15. Tian, C. *et al.* Deep learning on image denoising: An overview. *Neural Networks***131**, 251–275 (2020).
16. Lucas, A., Iliadis, M., Molina, R. & Katsaggelos, A. K. Using Deep Neural Networks for Inverse Problems in Imaging. *IEEE Signal Process. Mag.* (2018).

17. Dabov, K., Foi, A., Katkovnik, V. & Egiazarian, K. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Trans. Image Process.* (2007) doi:10.1109/TIP.2007.901238.
18. Zhang, K., Zuo, W., Chen, Y., Meng, D. & Zhang, L. Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising. *IEEE Trans. Image Process.***26**, 3142–3155 (2017).
19. Fan, L., Zhang, F., Fan, H. & Zhang, C. Brief review of image denoising techniques. *Vis. Comput. Ind. Biomed. Art***2**, 7 (2019).
20. Bagnall, G. C. *et al.* Low-field magnetic resonance imaging of roots in intact clayey and silty soils. *Geoderma***370**, (2020).
21. Makinen, Y., Azzari, L. & Foi, A. Collaborative Filtering of Correlated Noise: Exact Transform-Domain Variance for Improved Shrinkage and Patch Matching. *IEEE Trans. Image Process.* (2020) doi:10.1109/tip.2020.3014721.
22. Gu, S., Zhang, L., Zuo, W. & Feng, X. Weighted nuclear norm minimization with application to image denoising. in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2014). doi:10.1109/CVPR.2014.366.
23. Chen, Y. & Pock, T. TNRD. *IEEE Trans. Pattern Anal. Mach. Intell.* (2017) doi:10.1109/TPAMI.2016.2596743.
24. Jiang, D. *et al.* Denoising of 3D magnetic resonance images with multi-channel residual learning of convolutional neural network. *Jpn. J. Radiol.* (2018) doi:10.1007/s11604-018-0758-8.
25. Kao, Y. H. & MacFall, J. R. Correction of MR fc-space data corrupted by spike noise. *IEEE Trans. Med. Imaging* (2000) doi:10.1109/42.875184.
26. Campbell-Washburn, A. E. *et al.* Using the robust principal component analysis algorithm to remove RF spike artifacts from MR images. *Magn. Reson. Med.* (2016) doi:10.1002/mrm.25851.
27. Jin, K. H. *et al.* MRI artifact correction using sparse + low-rank decomposition of annihilating filter-based hankel matrix. *Magn. Reson. Med.* (2017) doi:10.1002/mrm.26330.
28. Jin, K. H., McCann, M. T., Froustey, E. & Unser, M. Deep Convolutional Neural Network for Inverse Problems in Imaging. *IEEE Trans. Image Process.* (2017) doi:10.1109/TIP.2017.2713099.
29. Eo, T. *et al.* KIKI-net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images. *Magn. Reson. Med.* (2018) doi:10.1002/mrm.27201.
30. Obungoloch, J. *et al.* Design of a sustainable prepolarizing magnetic resonance imaging system for infant hydrocephalus. *Magn. Reson. Mater. Physics, Biol. Med.* (2018) doi:10.1007/s10334-018-0683-y.
31. Fan, Q. *et al.* MGH-USC Human Connectome Project datasets with ultra-high b-value diffusion MRI. *Neuroimage* (2016) doi:10.1016/j.neuroimage.2015.08.075.
32. Hamarneh, G. & Jassi, P. VascuSynth: Simulating vascular trees for generating volumetric image data with ground-truth segmentation and tree analysis. *Comput. Med. Imaging Graph.* (2010) doi:10.1016/j.compmedimag.2010.06.002.

33. Dunbabin, V. M. *et al.* Modelling root-soil interactions using three-dimensional models of root growth, architecture and function. *Plant Soil* (2013) doi:10.1007/s11104-013-1769-y.
34. Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y. & Manzagol, P. A. Stacked denoising autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion. *J. Mach. Learn. Res.* (2010).

Figures

Figure 1

AUTOMAP neural network architecture and AUTOMAP Image Reconstruction at 6.5 mT – (a) The neural network architecture schema - The input is the complex k-space data and the output is either the real component or the imaginary component of the image. The architecture is comprised of 2 fully connected layers that learns the transform between the k-space and the image space by example from the forward model, followed by convolutional layers that takes advantage of the fact that there is inherent sparsity in

the k-space as in the image space. (b and c) Comparison of AUTOMAP Reconstruction with IFFT Reconstruction for 2D imaging in a water-filled structured phantom - Images were acquired with a bSSFP sequence (matrix size = 64×64 , TR=31 ms at 6.5 mT. The NA increases from left to right with their respective scan time below. (b) Upper panel shows AUTOMAP-reconstructed images and (c) the lower panel shows same image reconstructed with IFFT. The window level of reconstructed images (b and c for each NA) is identical. (d) SNR Analysis of the 2D phantom dataset across the different number of averages with the overall mean SNR versus the NA on the left axis with the AUTOMAP (■) and IFFT (□). The right axis plots the SNR Gain across the NA with the ratios of SNRs of AUTOMAP over IFFT. (e) The Root Mean Square Error (RMSE) of either the AUTOMAP reconstructed images (●) or the IFFT reconstructed images (●) were evaluated with respect to the 800-average IFFT reconstructed image as reference. (f, g and h) Reconstruction of 3D human head dataset - An 11-min (NA=50) 3D acquisition dataset was reconstructed with AUTOMAP (f) and IFFT (g). Shown here are 10 slices from the full 15 slice dataset. For comparison, a 22-minute (NA=100) acquisition reconstructed with IFFT is shown in (h). The window level is unchanged in all images.

Figure 1

AUTOMAP neural network architecture and AUTOMAP Image Reconstruction at 6.5 mT – (a) The neural network architecture schema - The input is the complex k-space data and the output is either the real component or the imaginary component of the image. The architecture is comprised of 2 fully connected layers that learns the transform between the k-space and the image space by example from the forward model, followed by convolutional layers that takes advantage of the fact that there is inherent sparsity in

the k-space as in the image space. (b and c) Comparison of AUTOMAP Reconstruction with IFFT Reconstruction for 2D imaging in a water-filled structured phantom - Images were acquired with a bSSFP sequence (matrix size = 64×64 , TR=31 ms at 6.5 mT. The NA increases from left to right with their respective scan time below. (b) Upper panel shows AUTOMAP-reconstructed images and (c) the lower panel shows same image reconstructed with IFFT. The window level of reconstructed images (b and c for each NA) is identical. (d) SNR Analysis of the 2D phantom dataset across the different number of averages with the overall mean SNR versus the NA on the left axis with the AUTOMAP (■) and IFFT (□). The right axis plots the SNR Gain across the NA with the ratios of SNRs of AUTOMAP over IFFT. (e) The Root Mean Square Error (RMSE) of either the AUTOMAP reconstructed images (●) or the IFFT reconstructed images (●) were evaluated with respect to the 800-average IFFT reconstructed image as reference. (f, g and h) Reconstruction of 3D human head dataset - An 11-min (NA=50) 3D acquisition dataset was reconstructed with AUTOMAP (f) and IFFT (g). Shown here are 10 slices from the full 15 slice dataset. For comparison, a 22-minute (NA=100) acquisition reconstructed with IFFT is shown in (h). The window level is unchanged in all images.

Fig 2: Noise floor comparison and Image metric analysis on the 3D human brain dataset. (a – d) Noise floor comparison - Slice 4 from the NA=50 reconstructed brain dataset shown above in Fig. 1f and 1g is displayed here with two different window levels: a normalized image on the left and a window level chosen to highlight the noise on the right. AUTOMAP is shown in (a), and IFFT is shown in (b). An additional DnCNN or BM3D image denoising step was applied to the image data reconstructed with IFFT (c and d respectively). (e – h) Image metric analysis of AUTOMAP and IFFT reconstruction with- and without the DnCNN step or the BM3D denoiser following transformation of the raw k -space data - The mean overall SNR in the whole-head ROI across all the 15 slices is shown in (g) for IFFT (●), denoised IFFT with BM3D (■) with DnCNN (p), AUTOMAP (○). Three additional metrics are computed: PSNR (f), RMSE (g), and SSIM (h). (i) The table summarizes the mean PSNR, SSIM, RMSE, SNR and SNR Gain values across all the slices. The SNR gain was calculated with respect to the conventional IFFT.

Figure 2

Noise floor comparison and Image metric analysis on the 3D human brain dataset. (a – d) Noise floor comparison - Slice 4 from the NA=50 reconstructed brain dataset shown above in Fig. 1f and 1g is displayed here with two different window levels: a normalized image on the left and a window level chosen to highlight the noise on the right. AUTOMAP is shown in (a), and IFFT is shown in (b). An additional DnCNN or BM3D image denoising step was applied to the image data reconstructed with IFFT

(c and d respectively). (e – h) Image metric analysis of AUTOMAP and IFFT reconstruction with- and without the DnCNN step or the BM3D denoiser following transformation of the raw k-space data - The mean overall SNR in the whole-head ROI across all the 15 slices is shown in (g) for IFFT (●), denoised IFFT with BM3D(■) with DnCNN (▲), AUTOMAP (○). Three additional metrics are computed: PSNR (f), RMSE (g), and SSIM (h). (i) The table summarizes the mean PSNR, SSIM, RMSE, SNR and SNR Gain values across all the slices. The SNR gain was calculated with respect to the conventional IFFT.

Fig 3: **Artifacts:** (a–d) Elimination of inadvertent hardware artifacts at 6.5 mT - Two slices from a 3D bSSFP (NA=50) are shown. When reconstructed with IFFT (a, b), a vertical artifact (red arrows) is present across slices. When the same raw data was reconstructed with AUTOMAP (c, d), the artifacts are eliminated. The error maps of each slice with respect to a reference scan (NA=100) is shown for both IFFT and AUTOMAP reconstruction. (e–g) Uncorrupted k-space (NA=50) was reconstructed with AUTOMAP (e) and IFFT (f). The reference NA=100 scan is shown in (g). (h–m) AUTOMAP reconstruction of simulated k-space artifacts. Two slices of the hybrid k-space from the 11-minute (NA=50) brain scan was corrupted with simulated spikes (h). In (i), the data was reconstructed with AUTOMAP trained on the standard corpus of white gaussian noise corrupted brain MRI images. In (j), the k-space data was reconstructed with AUTOMAP with a training corpus of k-space data including variable number of random spikes. IFFT reconstructed images are shown in (k), where the spiking artifacts are clearly seen. Denoised IFFT with DnCNN reconstructed images are shown in (l) and denoised IFFT with BM3D reconstructed images are shown in (m). (n) The table summarizes image quality metrics for the reconstruction task of the 3 slices, both with- and without spike corruption. PSNR, SSIM and RMSE were evaluated for reconstruction using IFFT, denoised IFFT with either DnCNN or

Artifacts: (a–d) Elimination of inadvertent hardware artifacts at 6.5 mT - Two slices from a 3D bSSFP (NA=50) are shown. When reconstructed with IFFT (a, b), a vertical artifact (red arrows) is present across slices. When the same raw data was reconstructed with AUTOMAP (c, d), the artifacts are eliminated. The error maps of each slice with respect to a reference scan (NA=100) is shown for both IFFT and AUTOMAP reconstruction. (e–g) Uncorrupted k-space (NA=50) was reconstructed with AUTOMAP (e) and IFFT (f). The reference NA=100 scan is shown in (g). (h–m) AUTOMAP reconstruction of simulated k-space artifacts. Two slices of the hybrid k-space from the 11-minute (NA=50) brain scan was corrupted with simulated spikes (h). In (i), the data was reconstructed with AUTOMAP trained on the standard corpus of white gaussian noise corrupted brain MRI images. In (j), the k-space data was reconstructed with AUTOMAP with a training corpus of k-space data including variable number of random spikes. IFFT reconstructed images are shown in (k), where the spiking artifacts are clearly seen. Denoised IFFT with DnCNN reconstructed images are shown in (l) and denoised IFFT with BM3D reconstructed images are shown in (m). (n) The table summarizes image quality metrics for the reconstruction task of the 3 slices, both with- and without spike corruption. PSNR, SSIM and RMSE were evaluated for reconstruction using IFFT, denoised IFFT with either DnCNN or BM3D, and AUTOMAP trained on either the standard gaussian noise-corrupted corpus or on a spike- and gaussian noise-corrupted corpus.

Fig 3: Artifacts: (a–d) Elimination of inadvertent hardware artifacts at 6.5 mT - Two slices from a 3D bSSFP (NA=50) are shown. When reconstructed with IFFT (a, b), a vertical artifact (red arrows) is present across slices. When the same raw data was reconstructed with AUTOMAP (c, d), the artifacts are eliminated. The error maps of each slice with respect to a reference scan (NA=100) is shown for both IFFT and AUTOMAP reconstruction. (e–g) Uncorrupted k -space (NA=50) was reconstructed with AUTOMAP (e) and IFFT (f). The reference NA=100 scan is shown in (g). (h–m) AUTOMAP reconstruction of simulated k -space artifacts. Two slices of the hybrid k -space from the 11-minute (NA=50) brain scan was corrupted with simulated spikes (h). In (i), the data was reconstructed with AUTOMAP trained on the standard corpus of white gaussian noise corrupted brain MRI images. In (j), the k -space data was reconstructed with AUTOMAP with a training corpus of k -space data including variable number of random spikes. IFFT reconstructed images are shown in (k), where the spiking artifacts are clearly seen. Denoised IFFT with DnCNN reconstructed images are shown in (l) and denoised IFFT with BM3D reconstructed images are shown in (m). (n) The table summarizes image quality metrics for the reconstruction task of the 3 slices, both with- and without spike corruption. PSNR, SSIM and RMSE were evaluated for reconstruction using IFFT, denoised IFFT with either DnCNN or

Figure 3

Artifacts: (a–d) Elimination of inadvertent hardware artifacts at 6.5 mT - Two slices from a 3D bSSFP (NA=50) are shown. When reconstructed with IFFT (a, b), a vertical artifact (red arrows) is present across slices. When the same raw data was reconstructed with AUTOMAP (c, d), the artifacts are eliminated. The error maps of each slice with respect to a reference scan (NA=100) is shown for both IFFT and AUTOMAP reconstruction. (e–g) Uncorrupted k -space (NA=50) was reconstructed with AUTOMAP (e) and IFFT (f).

The reference NA=100 scan is shown in (g). (h-m) AUTOMAP reconstruction of simulated k-space artifacts. Two slices of the hybrid k-space from the 11-minute (NA=50) brain scan was corrupted with simulated spikes (h). In (i), the data was reconstructed with AUTOMAP trained on the standard corpus of white gaussian noise corrupted brain MRI images. In (j), the k-space data was reconstructed with AUTOMAP with a training corpus of k-space data including variable number of random spikes. IFFT reconstructed images are shown in (k), where the spiking artifacts are clearly seen. Denoised IFFT with DnCNN reconstructed images are shown in (l) and denoised IFFT with BM3D reconstructed images are shown in (m). (n) The table summarizes image quality metrics for the reconstruction task of the 3 slices, both with- and without spike corruption. PSNR, SSIM and RMSE were evaluated for reconstruction using IFFT, denoised IFFT with either DnCNN or BM3D, and AUTOMAP trained on either the standard gaussian noise-corrupted corpus or on a spike- and gaussian noise-corrupted corpus.

Figure 4

Domain-specific training corpora used on plant roots dataset – (a, b and c) Representative images from three training sets for root MRI reconstruction. (a) 2D images from the Human Connectome Project database. (b) 2D images from the training set based on synthetic vascular trees. Each of the 2D images was obtained by summing up the 3D synthetic vascular tree volumes in all 3 dimensions. (c) Images of realistic simulated root system from the RootBox toolbox. (d) The list of matrix sizes of the acquired root datasets and their corresponding training set used for image reconstruction.

Figure 4

Domain-specific training corpora used on plant roots dataset – (a, b and c) Representative images from three training sets for root MRI reconstruction. (a) 2D images from the Human Connectome Project database. (b) 2D images from the training set based on synthetic vascular trees. Each of the 2D images was obtained by summing up the 3D synthetic vascular tree volumes in all 3 dimensions. (c) Images of realistic simulated root system from the RootBox toolbox. (d) The list of matrix sizes of the acquired root datasets and their corresponding training set used for image reconstruction.

Fig 5: Reconstruction of sorghum dataset acquired at 47 mT using AUTOMAP trained either on T1-weighted MRI brain images, synthetic vascular tree Images, or synthetic root images. Six 2D projections of roots images (labelled dataset 1 to 6) extracted from six different roots' samples are shown. All the 48 x 48 datasets were acquired at 47mT. In the upper panel (a), the raw data were reconstructed with AUTOMAP trained on T1-weighted MRI brain images, the panel (b) are the same datasets reconstructed with AUTOMAP trained on synthetic Vascular Tree Images and finally the same datasets were reconstructed with AUTOMAP trained on synthetic Root Images in panel (c). The lower panel (d) shows the images reconstructed with the standard IFFT method. All the images were windowed to the same level for comparison. The scale of 1cm represented one of the images is the same for all the 2D projections. (e) The table summarizes the mean SNR analysis of the 2D projections of the roots 48 x 48 datasets acquired at 47mT.

Figure 5

Reconstruction of sorghum dataset acquired at 47 mT using AUTOMAP trained either on T1-weighted MRI brain images, synthetic vascular tree Images, or synthetic root images. Six 2D projections of roots images (labelled dataset 1 to 6) extracted from six different roots' samples are shown. All the 48 x 48 datasets were acquired at 47mT. In the upper panel (a), the raw data were reconstructed with AUTOMAP trained on T1-weighted MRI brain images, the panel (b) are the same datasets reconstructed with

AUTOMAP trained on synthetic Vascular Tree Images and finally the same datasets were reconstructed with AUTOMAP trained on synthetic Root Images in panel (c). The lower panel (d) shows the images reconstructed with the standard IFFT method. All the images were windowed to the same level for comparison. The scale of 1cm represented one of the images is the same for all the 2D projections. (e) The table summarizes the mean SNR analysis of the 2D projections of the roots 48 x 48 datasets acquired at 47mT.

Fig 5: Reconstruction of sorghum dataset acquired at 47 mT using AUTOMAP trained either on T1-weighted MRI brain images, synthetic vascular tree Images, or synthetic root images. Six 2D projections of roots images (labelled dataset 1 to 6) extracted from six different roots' samples are shown. All the 48 x 48 datasets were acquired at 47mT. In the upper panel (a), the raw data were reconstructed with AUTOMAP trained on T1-weighted MRI brain images, the panel (b) are the same datasets reconstructed with AUTOMAP trained on synthetic Vascular Tree Images and finally the same datasets were reconstructed with AUTOMAP trained on synthetic Root Images in panel (c). The lower panel (d) shows the images reconstructed with the standard IFFT method. All the images were windowed to the same level for comparison. The scale of 1cm represented one of the images is the same for all the 2D projections. (e) The table summarizes the mean SNR analysis of the 2D projections of the roots 48 x 48 datasets acquired at 47mT.

Figure 5

Reconstruction of sorghum dataset acquired at 47 mT using AUTOMAP trained either on T1-weighted MRI brain images, synthetic vascular tree Images, or synthetic root images. Six 2D projections of roots images (labelled dataset 1 to 6) extracted from six different roots' samples are shown. All the 48 x 48 datasets were acquired at 47mT. In the upper panel (a), the raw data were reconstructed with AUTOMAP trained on T1-weighted MRI brain images, the panel (b) are the same datasets reconstructed with AUTOMAP trained on synthetic Vascular Tree Images and finally the same datasets were reconstructed with AUTOMAP trained on synthetic Root Images in panel (c). The lower panel (d) shows the images reconstructed with the standard IFFT method. All the images were windowed to the same level for comparison. The scale of 1cm represented one of the images is the same for all the 2D projections. (e) The table summarizes the mean SNR analysis of the 2D projections of the roots 48 x 48 datasets acquired at 47mT.

Figure 6

AUTOMAP reconstruction using the RootBox synthetic roots database versus IFFT reconstruction of a 96 x 96 root dataset. All eight 2D projections reconstructed with AUTOMAP are shown in the upper panel (a) and the lower panel (b) is the same dataset reconstructed with IFFT. The window level for projections 1 to 7 were set to the same value except for projection 8, where the threshold was lowered on both panels

revealing the noise floor differences. In (c) the SNR was evaluated for AUTOMAP reconstruction using the RootBox training and compared to IFFT reconstruction. (Left axis) The plotted graph shows the SNR of each of the 8 projections reconstructed with AUTOMAP (▲) and IFFT (△). (Right Axis) The SNR enhancement was evaluated and plotted for each projection in red (*).

Figure 7

AUTOMAP reconstruction using the RootBox synthetic roots database versus IFFT reconstruction of a 96 x 96 root dataset. All eight 2D projections reconstructed with AUTOMAP are shown in the upper panel (a) and the lower panel (b) is the same dataset reconstructed with IFFT. The window level for projections 1 to 7 were set to the same value except for projection 8, where the threshold was lowered on both panels revealing the noise floor differences. In (c) the SNR was evaluated for AUTOMAP reconstruction using the RootBox training and compared to IFFT reconstruction. (Left axis) The plotted graph shows the SNR of each of the 8 projections reconstructed with AUTOMAP (▲) and IFFT (△). (Right Axis) The SNR enhancement was evaluated and plotted for each projection in red (*).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [BoostingthesignaltonoiseofLowfieldMRIwithDeepLearningImageReconstructionssuppl.docx](#)