

dStripe: slice artefact correction in diffusion MRI via constrained neural network

Maximilian Pietsch^{1,2,3,✉}, Daan Christiaens^{2,4}, Joseph V Hajnal^{1,2}, and J-Donald Tournier^{1,2}

¹Centre for Medical Engineering, King's College London, London, UK

²Centre for the Developing Brain, King's College London, London, UK

³Department of Forensic & Neurodevelopmental Sciences, King's College London, London, UK

⁴Department of Electrical Engineering, ESAT/PSI, KU Leuven, Leuven, Belgium

MRI scanner and sequence imperfections and advances in reconstruction and imaging techniques to increase motion robustness can lead to inter-slice intensity variations in Echo Planar Imaging. Leveraging deep convolutional neural networks as universal image filters, we present a data-driven method for the correction of acquisition artefacts that manifest as inter-slice inconsistencies, regardless of their origin. This technique can be applied to motion- and dropout-artefacted data by embedding it in a reconstruction pipeline. The network is trained in the absence of ground-truth data on, and finally applied to, the reconstructed multi-shell high angular resolution diffusion imaging signal to produce a corrective slice intensity modulation field. This correction can be performed in either motion-corrected or scattered source-space. We focus on gaining control over the learned filter and the image data consistency via built-in spatial frequency and intensity constraints. The end product is a corrected image reconstructed from the original raw data, modulated by a multiplicative field that can be inspected and verified to match the expected features of the artefact. In-plane, the correction approximately preserves the contrast of the diffusion signal and throughout the image series, it reduces inter-slice inconsistencies within and across subjects without biasing the data. We apply our pipeline to enhance the super-resolution reconstruction of neonatal multi-shell high angular resolution data as acquired in the developing Human Connectome Project.

diffusion MRI | image artefact removal | venetian blind artefact

Correspondence: maximilian.pietsch@kcl.ac.uk, 1st Floor South Wing, St Thomas' Hospital, SE17EH London, UK

1. Introduction

Diffusion MRI (dMRI) provides unique information about the microstructural properties of brain tissue by sensitisation to the motion of water molecules on the order of micrometers via strong gradient amplitudes. However, this poses a major challenge for in-vivo imaging where bulk subject motion or flow can cause severe phase errors.

In single-shot echo planar imaging (EPI) (1, 2), the k-space data of a 2D image can be encoded within typically 100 ms after a single excitation which effectively freezes motion. A 3D image can be formed by acquiring a stack of parallel 2D EPI images at different slice positions. In EPI, interactions with previous pulses (spin-history effects) and interference across slices (stimulated echo artefacts (3, 4)), variations in slice timing, imperfect signal unmixing in simultaneous multi-slice (SMS) imaging, and scanner hardware limitations can all lead to inter-slice inconsistencies.

Sensitivity to subject motion and the push to higher in-plane and slice acceleration exacerbates the potential for inter-slice inconsistencies. This can destabilise super-resolution reconstruction algorithms and affect downstream data analyses. In this work, we address the problem of removing intra-slice inconsistencies in neonatal dMRI data of the brain, a cohort that is particularly prone to motion and, due to relatively long T_1 relaxation times, spin history effects.

1.1. EPI slice intensity inconsistencies. Techniques to reduce acquisition time, such as the use of EPI, often combined with partial Fourier (5), parallel imaging (6–8), and simultaneous multi-slice (SMS) (9–12) are prone to image degradation, particularly in the presence of motion. For example, inter-slice and intra-slice signal leakage in accelerated slice-Grappa EPI acquisitions from neuronal (13, 14) and non-neuronal origin, such as eye blinking (15) and motion (16, 17), cannot be fully suppressed, which can lead to false positive activations in fMRI analysis (14, 15). Variations in slice timings and subject motion orthogonal to the slice plane result in a temporary disruption to the steady-state and yield pose-, motion- and tissue-dependent intensity modulations (spin-history artefacts) (18); this is exacerbated when using SMS techniques, as they can reduce the repetition time (TR) to the order of typical T_1 times for brain tissue. Spin-history artefacts are typically 3 to 7% of the image intensity and difficult to model (17) as they are non-linearly related to motion trajectories (19). They can also be caused by localised motion related to breathing (18) and cardiac pulsation (20).

In this study, we aim to remove inter-slice artefacts in neonatal multi-shell dMRI scans that were acquired as part of the developing Human Connectome Project (www.developingconnectome.org). In this data, we observe inter-slice intensity variations of unknown origin: the shell-average of the raw images in scattered source space shows a clear stripe pattern (see figure 1). On a population level, parts of the artefact pattern seem to be linked with the diffusion gradient strength, acquisition order and multiband boundaries. In the shell-average images of a single subject with the least motion in the cohort as well as in the population average across 700 subjects, we observe that the intensity modulation seems to be relatively smooth in-plane, with sharp transitions in the through-plane direction. The stripe modulation artefact has a similar appearance to motion-related slice-wise intensity artefacts. These are frequently observed in the neonatal cohort and can dominate the slice mod-




Fig. 1. Raw (before motion and distortion correction) signal averaged within shells using data from the AP phase-encode direction. The single-subject data is from the scan with the highest (middle) and with the median motion-correction derived quality assessment score (right) in the cohort. The population-averaged raw signal shows clear slice-wise intensity modulation patterns that exhibit a b-value dependency. The population-average pattern is to some degree discernible in the low-motion subject but, in particular for subjects with more motion, the shell-average raw signal contains stripe patterns from motion-related artefacts.

ulation we aim to correct as demonstrated in the shell-average of the subject with median motion in figure 1. The properties of the artefact are described more comprehensively in the [Results and Discussion](#) section, where we assess it using spatial and angular constraints applied to estimated destripe fields.

1.2. Prior and related work in MRI artefact removal.

In fMRI analysis, motion-induced spin-history artefacts are commonly removed via nuisance regressors (16, 21) but there is currently no standard method for removing spin-history artefacts from diffusion-weighted images (22). Variable slice-timing artefacts are best avoided at the setup of the acquisition sequence but this may not be possible for all applications and scanner software versions (23) and retrospective correction might be required for already acquired data.

A suggested retrospective zeroth-order correction of intensity inconsistencies is to apply a single scale factor to even slices to match average intensities between those and adjacent slices (24). While the computational simplicity and transparency is appealing, this approach does not account for spatially variable or tissue-specific intensity modulation, is potentially biased if the true average image intensities do vary between slices and is not applicable to motion corrupted data. Prospective methods have been proposed to deal with similar artefacts in multi-slab acquisition, but to date these do not eliminate the problem (25–32).

1.3. Problem formulation and contribution. A physical model of the artefact or paired corrupted and uncorrupted data could inform us on how to model the correction; whether the stripe correction is best modelled as a linear or nonlinear,

local or non-local function of the data, what constraints to use and whether the model is tied to acquisition-space or to the subject anatomy or to a mixture of both. In the absence of this prior knowledge, we aim to make informed and data integrity-preserving decisions.

The effect of the artefact correction should be directly inspectable and, if possible, the model should be guaranteed to produce corrections that match the observed patterns of the artefact. To maintain the trustworthiness of the images, any degree of freedom beyond that should ideally be restricted.

We assume the artefact to be a smooth multiplicative voxel-wise correction field per acquired slice. In the absence of constraints, this choice does not limit the space of possible solutions. However, this choice matters when enforcing in-plane smoothness constraints. A multiplicative smooth in-plane field – similar to a B1 inhomogeneity correction field – ensures that the correction approximately preserves local image contrast within each acquired slice.

We build on (33) to create a method that facilitates the data-driven and model-free removal of intra-slice inconsistencies of potentially motion-artefacted diffusion-sensitised images in scattered source space or in motion-corrected space. Our framework:

- can be applied retrospectively
- works in tandem with motion correction techniques to remove stripes in the presence of subject motion
- is not tied to a particular q -space sampling scheme or motion correction technique

- requires no ground truth data for training
- produces directly human-inspectable correction fields, in the same space as the dMRI data
- uses explicit constraints that locally preserve in-plane image contrast

The source code will be made publicly available at <https://github.com/maxpietsch/dStripe>.

2. Methods

2.1. Data. The dMRI data was acquired as part of the developing Human Connectome Project. Each dataset consists of 300 volumes at $1.5 \times 1.5 \times 3 \text{ mm}^3$ resolution, acquired in 64 slices with 1.5 mm slice-overlap; interleave 3, shift 2; multi-band factor 4; TR/TE=3800/90 ms; diffusion weightings $b = 0, 400, 1000$, and 2600 s/mm^2 with 20, 64, 88, and 128 directions, respectively (34–36).

The data was reconstructed to 1.5 mm^3 isotropic voxel size and denoised in the complex domain (37). Field maps and brain masks were estimated with FSL `topup` (38) and `bet` (39).

For visualisation of single-subject data, the data ‘sub-CC00083XX10/ses-30900’ was used as it has clearly visible stripe patterns. This baby’s postmenstrual age (PMA) at scan is 42 weeks and the image lies in the 59th percentile quality score estimated from the motion and outlier weights (40). To create a population-average template from representative data, we randomly selected 32 subjects from 38 to 42 weeks PMA that have quality assessment scores above the 20th percentile (see figure 10). Summary data contains the single subject and population template data and additional 11 subjects ranging from 34.3 to 43.2 weeks PMA and quality scores ranging from 3rd to 100th percentile.

2.2. Slice-to-volume motion correction and reconstruction framework. The dMRI data was processed using a motion correction and reconstruction algorithm (MCR) with integrated slice-to-volume reconstruction (SVR) (40). In brief, the reconstruction is based on an iterative estimation of a data-driven multi-shell low-rank (model-free) data representation (SHARD), slice outlier estimation, and rigid registration algorithm (41). The reconstruction utilises information from overlapping slices and the native slice profiles for super-resolution deconvolution and is formulated as an inverse problem that iteratively estimates the reconstruction coefficients \mathbf{x} , defined in the motion-corrected “anatomical” space (the moving subject-aligned reference frame), and the rigid motion parameters $\boldsymbol{\mu}$ that map between “source” space (the scattered slices in scanner coordinates) and anatomical space by minimising the difference between the acquired signal of a slice in the source space \mathbf{y}_s and its signal prediction

$$\min_{\mathbf{x}, \boldsymbol{\mu}} \sum_s w_s \|\mathbf{y}_s - \mathbf{B}_s \mathbf{M}(\boldsymbol{\mu}_s) \mathbf{Q}_s(\boldsymbol{\mu}_s) \mathbf{x}\|_2^2 + r(\mathbf{x}).$$

The model consists of the q-space (SHARD) basis $\mathbf{Q}_s(\boldsymbol{\mu}_s)$, the linear motion and interpolation operator $\mathbf{M}(\boldsymbol{\mu}_s)$, and the

blurring and slice selection matrix \mathbf{B}_s that also incorporates the slice sensitivity profile. Slice weights w_s are used to reduce the effect of outliers, and a regularisation term $r(\mathbf{x})$ is used to stabilise the inverse problem. EPI distortions are corrected by unwarping the input dMRI data before each reconstruction step using a field map and the subject motion parameters.

To isolate the effects of dStripe from those of the motion correction, motion, slice-weight and SHARD basis parameters were kept constant after their estimation on the original data for subsequent reconstruction when using source-space destriping (see below).

2.3. dStripe approaches: source or anatomical space.

MCR provides a mapping between the “source” space of the scanner and the “anatomical” space of the moving subject and separates dropout and other artefacts from anatomical features. The stripe artefact can be affected by the scanner geometry and the subject tissue properties and motion. This raises a question about the best space to operate in.

2.3.1. Anatomical-space destriping. Destriping can be applied in anatomical space to remove residual stripe patterns in the MCR output. This can be performed with any motion correction and reconstruction framework as it operates solely on its output and produces correction fields directly in the space of interest. However, this approach cannot be used to interactively refine the MCR, and will only work well if the destripe mechanism can cope with stripe patterns that are themselves scattered due to motion.

2.3.2. Source-space destriping. Taking advantage of the mapping MCR provides, we can use information from the motion-corrected anatomical space to correct the raw data in source space. For each excitation (single slice or multiband shot), we generate the corresponding *corrected* signal prediction ($\mathbf{B}_s \mathbf{M}(\boldsymbol{\mu}_s) \mathbf{Q}_s(\boldsymbol{\mu}_s) \mathbf{x}$). This allows modulating the corresponding slice in its “native” orientation, potentially increasing the effectiveness of the dStripe algorithm.

This approach relies on the signal representation in anatomical space capturing the contrast of interest together with the stripe patterns. Similarly to anatomical-space destriping, this approach does not necessarily capture the full source-space stripe pattern as intensity modulations in target space can be attenuated due to interpolation or if they are discarded by the outlier removal or rank-reduced model fit. Moreover, destriping each excitation in “slice-native” space requires destriping a slice-native volume with corresponding motion parameters; for our data, this increases the computational cost by a factor of 16 compared to anatomical-space destriping (64 slices, multiband 4). Finally, to ensure convergence of the output, we need to destripe and perform the subsequent MCR again multiple times (3 times is sufficient in our data), further increasing the computational cost of the source-space destriping approach.

Note that because the dHCP employed overlapping slices with a super resolution reconstruction, stripe patterns are amplified by the slice profile deconvolution in the reconstruction.

tion. When projecting the signal to source-space, we therefore modify the slice selection matrix B_s to preserve the super-resolved contrast by avoiding slice profile blurring. Instead, we subsequently downscale the dStripe-estimated slice modulation fields to account for the difference in slice thickness.

2.4. Metric of stripiness. dMRI images, in particular highly diffusion weighted volumes, have significant anatomical contrast. While stripe patterns can be detected visually with the trained eye, it is not trivial to mathematically separate expected from artefactual through-plane intensity variations.

Inter-slice signal variation can be measured for instance using the standard deviation (SD) of the dMRI signal (or its spherical harmonic representation) over a kernel of neighbouring voxels along the through-plane direction. We expect signal-preserving destriping to reduce the through-plane SD while preserving the equivalent in-plane SD. However, directly optimising a SD-derived measure is not viable because the expected anatomical through-plane SD dominates the artefactual SD by one to two orders of magnitude.

Therefore, we decided to use a neural network trained to remove simulated stripes artificially introduced into the training data. To assess its performance on the original data, we opt to visually rate images using scalar or rotation invariant representations of the diffusion weighted data, as well as derived images, as outlined below.

The angular information of the signal can be expressed in the basis of real, symmetric spherical harmonics (SH) (42), which allows to investigate its angular frequency spectrum (43). We use the l_2 -norm of the coefficients in a particular harmonic band $\ell = [0, 2, 4, \dots]$ defined as $\|S\|_\ell = \sqrt{\sum_{m=-\ell}^{\ell} (c_\ell^m)^2}$, where c_ℓ^m are the SH coefficients of the dMRI signal in the given voxel and shell. These angular measures $\|S\|_\ell$ are proportional to the square root of the power spectral density corresponding to that frequency band.

Finally, we investigate data consistency across shells and in the angular domain using (i) a diffusion tensor representation (44) and (ii) the brain tissue and free water maps obtained using multi-shell multi-tissue constrained spherical deconvolution (MT CSD) fits (45) (using population-averaged basis functions measured in WM and in CSF voxels (46, 47)). Signal changes due to dStripe are assessed quantitatively and visually using the fit residuals and fit-derived measures.

2.5. Destriping method.

2.5.1. The CNN architecture. To allow potentially taking large parts of the spatial context into account, the network operates on the full field of view of individual volumes ($99 \times 99 \times 64$ voxels). To accommodate this data in GPU memory and for performance reasons the network was designed to be relatively small (18,109 trainable parameters) and has intensity and spatial frequency filter constraints directly built into the last layers of the network. An overview of the network architecture, implemented in pytorch (52), is given in table 1.

| # | layer | ch: x,y | param. |
|----|--|--------------------------------|--------|
| | input volume I | 1: X,Y | |
| 1 | <i>SeparableConv</i> (3,3,3) | 16: X,Y | 60 |
| 2 | <i>BatchNorm</i> | | 32 |
| 3 | <i>ConvBlock</i> (3,3,7) | | 2,128 |
| 4 | <i>ConvBlock</i> (3,3,7) & <i>ReLU</i> | | 2,128 |
| 5 | <i>ConcatPool</i> (2,2,1) | 32: $\frac{X}{2}, \frac{Y}{2}$ | |
| 6 | <i>BatchNorm</i> | | 64 |
| 7 | <i>ConvBlock</i> (3,3,7) | | 5,792 |
| 8 | <i>ConvBlock</i> (3,3,7) & <i>ReLU</i> | | 5,792 |
| 9 | <i>ConcatPool</i> (.,.,1) | 64: 16,16 | |
| 10 | <i>SeparableConv</i> | 32: 16,16 | 2,080 |
| 11 | <i>SeparableConv</i> & <i>ReLU</i> | 1: 16,16 | 33 |
| 12 | dynamic range constraint | | |
| 13 | log transform | | |
| 14 | x,y lowpass filter (9,9,1) | | (81) |
| 15 | z highpass filter (1,1,9) | | (9) |
| 16 | exp transform | | |
| 17 | x,y upsample | 1: X,Y | |

Table 1. The dStripe network architecture as used during training. Input to the network is a single 3D volume I of dimensions X, Y, Z (one channel). The output of the network is the multiplicative field of the same dimension as the input image that, when applied to the image, ideally reverses the stripe-producing mechanism. For approximate invariance to the image dimensions, all layers use voxel-wise or convolutional operations. Spatially, we aim to represent slice-specific information at native resolution through-plane while limiting the in-plane resolution to yield a smooth field after upsampling. Hence, the resolution in the stack dimension (Z) is preserved throughout the network but is reduced in-plane via pooling layers (#5,#9) to half its original extent and then adaptively to 16×16 voxels. In all convolution layers, the image extent is preserved via zero padding. Throughout the network, depth-wise separable convolutions (*SeparableConv*) are used to limit the number of parameters while allowing the network to learn global relations between feature maps and their respective spatial extents (48). This module contains 3 stacked dilated convolution layers with dilations $1 \times 1 \times [1, 2, 3]$ (49). *SeparableConv* consist of one $3 \times 3 \times 3$ convolution followed by a pointwise $1 \times 1 \times 1$ convolution (if the number of output channels exceeds 1). *ConvBlock* consist of 3 concatenated spatial convolutions with z dilations $[1, 2, 3]$, followed by a $1 \times 1 \times 1$ convolution and optionally an activation *ReLU* (rectified linear unit, $v \mapsto \max(0, v)$). *ConcatPool* layers concatenate spatial average and maximum pooling layers. The dynamic range constraint layer computes $v \mapsto \frac{2.0}{1+\exp(-v)} + 10^{-4}$. The inplane lowpass filter is implemented as a fixed 2D Gaussian blur filter and the through-plane high-pass filter subtracts the low-frequency filtered image using a fixed 1D Gaussian filter ($v \mapsto v - f(v)$). For faster training, we use batch normalisation (*BatchNorm*) (50, 51). Layers with operations aggregating information across multiple spatial locations are highlighted with their (maximum) spatial extent in brackets, those without are pointwise operations or $1 \times 1 \times 1$ convolutions. For clarity, dimensions are omitted if unchanged. Parameter counts for fixed layers are denoted in brackets. For inference, layer #15 is deactivated and its function replaced by FFT-based high-pass-filtering of the network output (see section 2.5.6).

2.5.2. Modulation field constraints. To restrict the intensity modulation range, the output of the last convolution and *ReLU* layer (#11) is mapped to the range $[1.0001, 2.0001]$ via $v \mapsto \frac{2.0}{1+\exp(-v)} + 10^{-4}$ (layer #12). This scaling also facilitates spatial frequency filtering of the multiplicative destripe field in the log-domain (#13) using numerically stable additive operations (#14,15). The in-plane and subsequent through-plane filters are implemented as immutable convolution filters. To suppress high frequency information in-plane, a 2D Gaussian blur filter ($\sigma = 1.5$ voxels, kernel size 9×9) is used. The through-plane high-pass filter is similarly implemented by subtracting a 1D Gaussian-filtered version of the image ($\sigma = 1$ voxels), which allows factoring out low-frequency background modulations (and global offset) in the

through-plane direction.

Finally, exponential scaling (#16) and in-plane upsampling (#17) yield the multiplicative field at the target resolution while preserving in-plane smoothness and through-plane resolution.

2.5.3. Training and augmentation. Training is performed on 20 dMRI datasets in 500 epochs with a batch size of one, with the Adam optimiser (53) and triangular learning rate scheduler $[1, 5] \times 10^{-4}$ (80 training iterations per half-cycle) (54). From each dMRI dataset, 10 randomly selected volumes are selected for each b -value and all images divided by its 99th percentile intensity.

Data augmentation was performed by random reorientation via dihedral transformations (90 degree rotations and axis-aligned reflections) such that all three axes (AP, LR, IS) are used as the slice direction with equal probability. Rotated versions of the data can be assumed stripe-free, and are included in the augmentation to minimise the influence of any stripes that might have been present prior to image augmentation.

Slice-wise intensity stripe patterns were simulated to allow random multiplicative factors per slice that are correlated in time according to the slice interleave pattern. For all slices i of a temporally continuous excitation sequence block $p \in [0, 1, 2]$ ($i \bmod 3 = p$), slice-wise intensity scalings were drawn from $s_i \sim S = \max(|\mathcal{N}(\mu_p, \sigma_p^2)|, 10^{-5})^2$, with block-specific variance $\sigma_p^2 = 0.05$ and random but fixed centre drawn from a uniform distribution $\mu_p \sim \mathcal{U}(0.9, 1.1)$. This was repeated for all 3 excitation blocks. Finally the scaling vector s was normalised to unit geometric mean to approximately preserve global scaling effects.

2.5.4. Stein's Unbiased Risk Estimator-based image recovery. Consider the image reconstruction problem to infer an unknown image $\mathbf{x} \in \mathbb{R}^n$ from corrupted measurements $\mathbf{y} = \mathbf{H}\mathbf{x} + \delta$, with the linear measurement operator \mathbf{H} and errors δ . This problem is commonly solved by using prior knowledge about the measurement operator and properties of the image. Convolutional neural networks are powerful models to encode image properties but typically require large amounts of paired training data.

In this work, we have only one corrupted measurement and no ground truth data. Therefore, instead of supervised learning with paired data, we use a training technique that is based on Stein's Unbiased Risk Estimator (SURE) (55), that allows reconstructing an unobserved image \mathbf{x} from individual additive Gaussian noise corrupted measurements as for instance demonstrated in (56, 57). SURE is a measure of the expected generalisation loss ("risk") of an estimator of the mean of a data-generating process $\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbb{I}\sigma^2)$ with unknown mean $\boldsymbol{\mu}$. Assuming the measurements are corrupted by standard additive homoscedastic noise, $\mathbf{H} = \mathbb{I}$, and that the estimator f_Θ with parameters Θ is weakly differentiable, we can

express the expected reconstruction error as (55)

$$\mathbb{E} \left[\frac{1}{n} \|\mathbf{x} - f_\Theta(\mathbf{y})\|_2^2 \right] = \mathbb{E} \left[\frac{1}{n} \|\mathbf{y} - f_\Theta(\mathbf{y})\|_2^2 \right] - \sigma^2 + \frac{2\sigma^2}{n} \nabla_{\mathbf{y}} f_\Theta(\mathbf{y}), \quad (1)$$

where the expectation term measures the bias, and the divergence term $\nabla_{\mathbf{y}} f_\Theta(\mathbf{y}) = \sum_{i=1}^n \frac{\partial f_\Theta(\mathbf{y})}{\partial y_i}$ expresses the response of the model to input perturbations (model variance). Remarkably, SURE does not require access to $\boldsymbol{\mu}$ as σ can be estimated from the data (58) and $\|\mathbf{y} - f_\Theta(\mathbf{y})\|_2^2$ is directly accessible for a given model (55). $\nabla_{\mathbf{y}} f_\Theta(\mathbf{y})$ can be hard or impossible to derive analytically for complex f_Θ . However, for bounded functions $f_\Theta(\mathbf{x})$ with intractable derivatives or prohibitively high-dimensional parameter spaces, such as deep convolutional neural networks, Monte-Carlo techniques can be used to estimate the divergence

$$\nabla_{\mathbf{y}} f_\Theta(\mathbf{y}) = \lim_{\epsilon \rightarrow 0} \mathbb{E} \left[\mathbf{p}^T \frac{f_\Theta(\mathbf{y} + \epsilon \mathbf{p}) - f_\Theta(\mathbf{y})}{\epsilon} \right],$$

with the normally distributed noise vector $\mathbf{p} \sim \mathcal{N}(\mathbf{0}, \mathbb{I})$. As shown by (59), due to the high dimensionality of the data, this can be approximated with a single or few samples and a small non-zero ϵ via

$$\nabla_{\mathbf{y}} f_\Theta(\mathbf{y}) \approx \mathbf{p}^T \frac{f_\Theta(\mathbf{y} + \epsilon \mathbf{p}) - f_\Theta(\mathbf{y})}{\epsilon}. \quad (2)$$

2.5.5. Loss functions. The training loss is evaluated inside the brain mask \mathbf{M} and depends on the slice direction after the dihedral transformation d . Along the AP and LR axes, it consists of the sum of \mathcal{L}_{aug} , the mean squared loss between original data \mathbf{S}_o and augmented data $\mathbf{S}_{d,s} = \mathbf{s} \circ (d \circ \mathbf{S}_o)$ corrected by the network f_Θ

$$\mathcal{L}_{\text{aug}} = \|d \circ \mathbf{S}_o - f_\Theta(\mathbf{S}_{d,s})\|_{2,\mathbf{M}}^2$$

and $\mathcal{L}_{\text{const}}$, the mean squared loss that penalises altering the original image (free of stripe artefacts in that direction):

$$\mathcal{L}_{\text{const}} = \|d \circ \mathbf{S}_o - f_\Theta(d(\mathbf{S}_o))\|_{2,\mathbf{M}}^2.$$

For slice directions along IS, we construct the MC SURE-based loss using an augmented image as input. Following equations 1 and 2 and limiting the divergence to the non-negative domain to bound the loss, we use

$$\mathcal{L}_{\text{SURE}} = \|\mathbf{S}_{d,s} - f_\Theta(\mathbf{S}_{d,s})\|_{2,\mathbf{M}}^2 - \sigma_{\text{SURE}}^2 + 2\sigma_{\text{SURE}} |\text{div}_{\text{SURE}}|$$

with $\sigma_{\text{SURE}} = \sigma(\mathbf{S}_{d,s} - d \circ \mathbf{S}_o)_{\mathbf{M}}$ the standard deviation of the signal change due to image augmentation inside the brain mask, and div_{SURE} the model variance penalty to input perturbations which is estimated using

$$\text{div}_{\text{SURE}} = \mathbf{P}^T \frac{f_\Theta(\mathbf{S}_{d,s} + \epsilon \mathbf{P}) - f_\Theta(\mathbf{S}_{d,s})}{\epsilon} \Big|_{\mathbf{M}},$$

with $\epsilon = 10^{-3}$ and the calibrated perturbation image $\mathbf{P} = (\hat{\mathbf{s}} \circ \mathbf{S}_{d,s} - \mathbf{S}_{d,s}) / \sigma(\hat{\mathbf{s}} \circ \mathbf{S}_{d,s} - \mathbf{S}_{d,s})_{\mathbf{M}}$ generated using an additional simulated slice modulation $\hat{\mathbf{s}}$ drawn from S .

2.5.6. Modifications for inference: attention and recursion.

Training is performed as described, however during inference we can improve performance and avoid issues we have observed. First, destripe performance can be improved by recursively applying the neural network for up to three iterations after which the resulting images are approximately converged. However, naïve recursion would increasingly relax the field's frequency constraints. Second, in-plane, the field estimates require additional low-pass filtering due to up-sampling artefacts of the bilinear interpolation layer (#17). Third, the field estimate can be unreliable in areas dominated by noise or ill-defined due to the lack of signal of interest and close to the edge of the field of view. This can, through the field frequency constraints, negatively impact the field in adjacent brain areas.

The approach to address these issues consists of the following modifications to the final stages of the network, applied during inference only:

- 1) We **remove the built-in high-pass filter** (layer #15) and replace its function with a FFT-based high-pass filter (\mathcal{H}_z) that is applied to the (upsampled) field.
- 2) To **prevent frequency spectrum drift** over iterations i , we do not apply the through-plane filter \mathcal{H}_z (described above) to the output of the network, but instead to the composition of the current best estimate of the field \mathbf{F}_i and its update:

$$\mathbf{F}_{i+1} \leftarrow \mathcal{H}_z(\mathbf{F}_i * f_{\Theta}(\mathbf{S}_o * \mathbf{F}_i))$$

For this filter, we use a parametric FFT-based frequency filter which allows high-quality frequency constraints that can be easily adapted retrospectively. We use a 4th order Butterworth filter (B_1 : normalised frequency cutoff: 21/32, padding: 17) to block low-frequency components.

3) To **address edge effects** we implement an attention mechanism based on the brain mask that attenuates high-frequency contributions to the field from outside trusted areas. This modulation has to be applied before high-pass frequency filtering to reduce leakage of the field into areas of interest. Hence, the attention mechanism is incorporated into the through-plane frequency filter \mathcal{H}_z . It uses a through-plane frequency separation into low- and mid- to high-frequency components and a spatial attention map \mathbf{A} to smoothly blend these components via geometric weighted averaging in the image domain.

Specifically, $\mathcal{H}_z(\mathbf{F})$ is defined as

$$\begin{aligned} \text{step 1: } \mathbf{F} &\leftarrow \exp(\log(\mathbf{F}/\mathbf{F}_{lp,3D}) * \mathbf{A}) * \mathbf{F}_{lp,3D} \\ \text{step 2: } \mathbf{F} &\leftarrow \mathbf{F}/\text{FFT}_z^{-1}(\text{FFT}_z(\mathbf{F}) * B_1), \end{aligned}$$

where $\mathbf{F}_{lp,3D}$ is a lowpass-filtered version of the field (3D Gaussian-blur, $\sigma = 5 \times 5 \times 11$ voxels). In step 1, \mathbf{A} is applied to the frequency-separated field to down-scale untrusted areas in the medium- to high-frequency components. In step 2, the recombined field is subsequently high-pass filtered through-plane with the Butterworth filter B_1 . The key is that the Gaussian filter allows estimating the low-frequency background without causing distortion to the high-frequency components of interest and \mathbf{A} is spatially smooth, allowing a

weighted blending of low- with medium- to high-frequency components. In areas with low trust (low \mathbf{A}), the low-frequency field dominates in step 1, hence high-frequencies from these areas contribute little to the filtering in step 2.

We use a brain mask \mathbf{M} to create the attention map \mathbf{A} . In plane, to gradually reduce contributions from outside the brain mask while preserving attention inside the brain, \mathbf{M} is dilated by the full width at one-tenth maximum of a Gaussian filter and subsequently blurred (13 dilations, $\sigma = 9 \times 9$ voxels). Through-plane, this image is smoothed further ($\sigma = 3$ voxels), down-weighting contributions of the most inferior and superior parts of the masked area.

4) As a final step, the field is low-pass filtered in-plane with a cutoff frequency chosen to suppress **in-plane upsampling artefacts** (3rd order Butterworth filter B_2 , normalised frequency cutoff: 2/32, padding: 24).

3. Results and Discussion

3.1. Degrees of freedom of slice modulation field. As outlined in the introduction (1), in this work we assumed that the field needs to be modelled as a function of position and b -value (see figure 1), and direction of the diffusion weighting. Figure 2 shows that this is indeed necessary: restricting the field $f(b, \ell, m, x, y, z)$ that has the same rank as the data, in the angular ($f(b, x, y, z)$) and spatial domain ($f(b, x, y, z)$, $f(b, z)$) leads to increased residual stripes, justifying the design decisions of the method.

In figure 2, the l_2 -norm of the resulting data as well as the difference to the data pre-dStripe are shown for representative b and ℓ value combinations for a representative subject. In the shell-average images, shell-specific but slice-wise constant multiplicative scale factors (figure 2b) reduce the stripe pattern but a smoothly varying field (figure 2c) reduces the stripe pattern further. However, stripe patterns in the angular domain remain nearly unchanged. The full dStripe field attenuates stripe patterns in the angular domain and reduces the power in higher angular frequency components. By design, an inplane reduction in high-frequency SH power can be attributed to smoothly varying intensity modulations of individual dMRI volumes and is therefore likely caused by a reduction of angular variance due to the removal of stripe artefacts.

3.2. Space-dependency and subject motion. If the artefact is tied to scanner space, it is desirable to destripe in source space; similarly if it is fixed to the subject tissue, dStripe should be performed in anatomical space. Given the observation that the field has consistent features in source space of a cohort with relatively prevalent motion, it is reasonable to at least partially model the field in source space. However, it is unclear if it is best to correct stripe patterns in scattered source or in the motion-corrected anatomical space as this depends not just on the nature of the artefact but also on how well it can be corrected given our framework, where we rely on a motion-corrected model of the data.

The properties of the estimated slice modulation artefact depend on the space in which it is estimated. In particular,

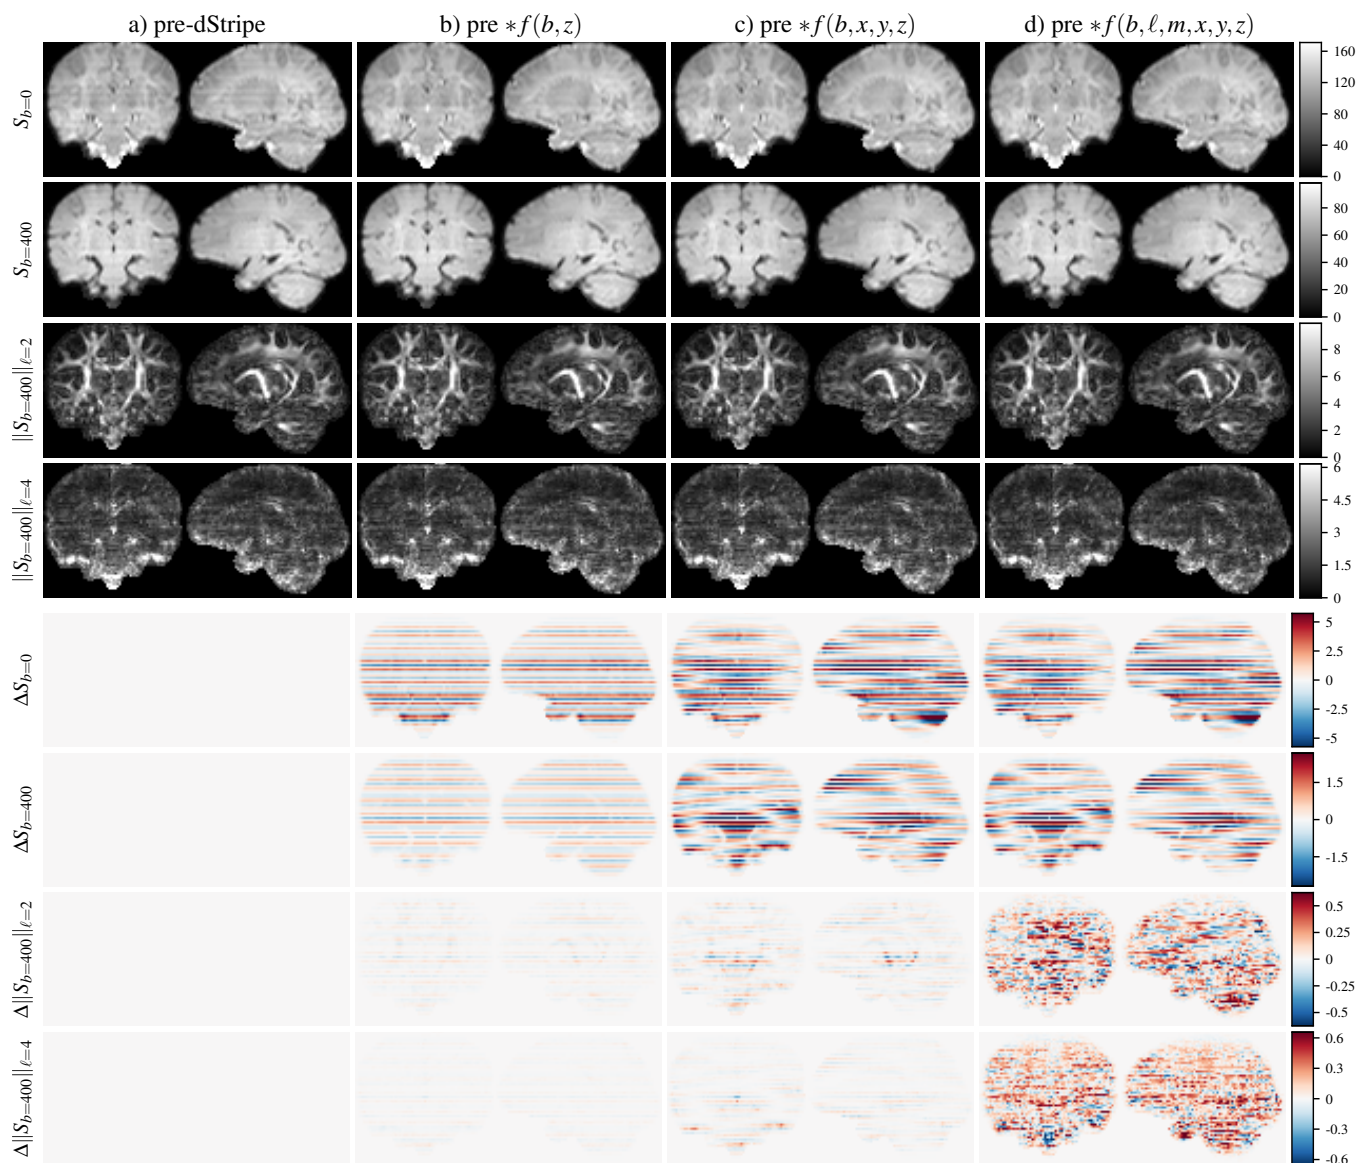


Fig. 2. Exemplary components of the l_2 -norm spectrum with increasing degrees of freedom of the correction field (top) and the difference in the l_2 -norm spectrum due to dStripe (pre-dStripe - post-dStripe) shown below. a): pre-dStripe: original data after motion correction, b) 1D b-value specific correction applied, c) 3D scalar b-value specific field applied, d) full correction without rank constraints. b-values in units of s/mm^2 .

the super-resolution reconstruction increases inter-slice intensity variations compared to the patterns observed in the source data acquired with overlapping slices. The presence of subject motion also affects the spatial frequency content of the artefact: assuming a smooth field in source space, subject rotation can cause higher frequency patterns in motion-corrected space. We rely on a motion-corrected representation of the data, where artefacts originating from source space might be smeared out or might not be fixable given the frequency constraints of the field. Similarly, when projected to source space, stripes observed in the motion corrected signal could be distributed over multiple scattered (oblique) slices.

A possible source of destripe field inconsistencies for source-space methods is a potential field of view or pose-dependency of the method. To estimate the slice-specific field, dStripe requires spatially contiguous adjacent slices, in our case a complete “slice-native” volume for context. For source-space

destriping, intra-volume motion can affect the pose of the slice-native volume for each slice of interest, introducing a possible source of variance into the destripe field estimation, which can potentially be amplified or introduce local artefacts in the subsequently required motion correction, especially if super-resolution algorithms are used. In anatomical-space destriping, the pose is fixed eliminating this source of inter-slice and intra-volume variance.

Finally, the choice of the space in which dStripe is performed has implications for the nature of the change to the motion corrected data. A smooth modulation applied to scattered source space data can improve subsequent (super-resolution) motion correction but also introduces local changes in its output due to the interpolation and aggregation of data from multiple slices; dStripe performed in anatomical space guarantees smoothly varying fields in the space where the data is used for analysis.

In the following sections, we attempt to remove stripes in both spaces independently and compare the data in anatomical space. This does not prove the origin of the artefact but assesses the stripe artefact removal potential for each space, given the framework.

Figure 3 shows the shell-averaged data in motion-corrected space without dStripe, with dStripe in anatomical space and in source space. The shell-average signal changes most in the lower b -values, source-space destriping changes the signal slightly more, but the difference between source- and anatomical-space destriping are much smaller than between data pre- and post-dStripe. The difference between anatomical- and source-space dStripe at the top of the brain is related to the implementation of the attention filter that can be gradually circumvented with further dStripe iterations. In the presence of motion, potential field of view edge effects of the network can appear in more central slices, in particular if motion caused the brain to be cropped during the acquisition. Figure 4 shows raw $b = 0$ data and corresponding fields (estimated in source space) for a single volume and averaged across the shell. By design, the fields are smooth in-plane and relatively sharp through-plane. In source space, using source-space destriping, the average field shares similarity with that in individual volumes but is attenuated overall. This attenuation can originate from a non-stationary of the modulation artefact, for instance caused by blurring due to subject motion if the field is tissue dependent, or assuming a fixed field in source-space, it could indicate variance in the estimated dStripe field due to varying subject pose.

3.3. Performance. To evaluate performance we use the following objectives: post-dStripe, dMRI data has to appear visually less stripy and have reduced variance in the through-plane direction. Furthermore, we use DTI and MT CSD fits to evaluate the consistency of the data across shells and in the angular domain.

3.3.1. spatial dMRI signal variance. The local standard deviation of the (motion corrected) dMRI signal evaluated in 1D patches of 7 voxels inside the brain mask aligned in-plane (x, y) and through-plane (z) of each spherical harmonic volume are shown in figure 5. The analysis is split by b -value and harmonic band. Note that in-plane and through-plane values of (anatomical) variance are not equal (see left column) and anatomical and inter-subject variance dominate changes due to dStripe. Performing a subject-specific comparison, dStripe reduces through-plane variance while approximately preserving in-plane signal variance. Anatomical-space destriping tends to reduce in-plane variance more, source-space destriping reduces through-plane variance in the outer shell more which is partially driven by the difference in attention filters but also due to the amplitude of changes (see figure 3).

3.3.2. Diffusion signal representation fits. Shell-average fit residuals and fit-derived maps are shown for a subject in figure 6. Stripy residual maps for both data representations indicates that the fit residuals are sensitive to stripe patterns;

MD and tissue and fluid volume fraction maps particularly show clear patterns of stripe artefacts prior to dStripe. After (anatomical-space) dStripe, both residual maps and MD and volume fraction maps show clearly reduced stripe patterns while preserving anatomical contrast.

Using diffusion tensor and MT CSD root mean squared fit residuals as a proxy for data consistency across shells and directions, figure 7 shows that dStripe increases data consistency across all shells for both signal representations. Fit residuals are reduced mostly in the $b = 400 \text{ s/mm}^2$ and $b = 1000 \text{ s/mm}^2$ shells. Source-space dStripe compared to anatomical-space dStripe tends to yield slightly lower residuals with the exception of the $b = 0 \text{ s/mm}^2$ and $b = 2600 \text{ s/mm}^2$ shell for the DT fit. For both fits and spaces a larger relative reduction of residuals can be observed for data with higher residuals in the $b = 400 \text{ s/mm}^2$ and $b = 1000 \text{ s/mm}^2$ shells; dStripe improves residuals the most for high-residual data and source-space dStripe slightly outperforms anatomical-space dStripe in this regime.

3.4. Population-level effects and anatomical bias. As demonstrated, dStripe reduces stripe patterns in the dMRI data and overall increases its consistency within subjects. Here, we use data from 32 age-matched subjects to assess the effect of dStripe on data consistency across subjects and whether the dStripe method introduces systematic changes (bias) across subjects related to specific anatomical locations. Specifically, to assess the spatial distribution of the signal changes within each subject in a common reference frame and to measure bias possibly tied to subject anatomy, we jointly coregister each subjects' pre- and post-dStripe data to create a common population-average template for pre- and post-dStripe data. Below we analyse changes to signal properties between groups, first across the whole brain then spatially-resolved.

3.4.1. Whole-brain analysis. Figure 8 shows the relative change of CSD-derived quantities inside the brain that can be attributed to applying dStripe. We display the bias using histograms of the voxel-wise relative change in the population-average fluid and tissue component volume fractions and the tissue angular power (A). After dStripe, on average, the volume fractions in the template components are close to constant (fluid: +0.12%, tissue: -0.07%), the tissue angular energy is slightly reduced in the $\ell = 2$ (-0.30%) and $\ell = 4$ (-0.44%) band and slightly increased in the $\ell = 6$ band (+1.00%). This is in line with an observed overall reduction in angular power in single-subject dMRI signal (see figure 2d).

In the coregistered data, we use the intra-template cross-subject standard deviation as a measure of data consistency. Figure 8 (B) shows the histograms of the relative reduction in voxel-wise cross-subject standard deviation due to dStripe. The relative inter-subject variation is most decreased by dStripe in the fluid volume fraction and in the $\ell = 4$ band of the tissue signal. On average, dStripe decreases inter-subject standard deviation nearly across the entire power spectrum of the fluid and tissue components (figure 8f: -

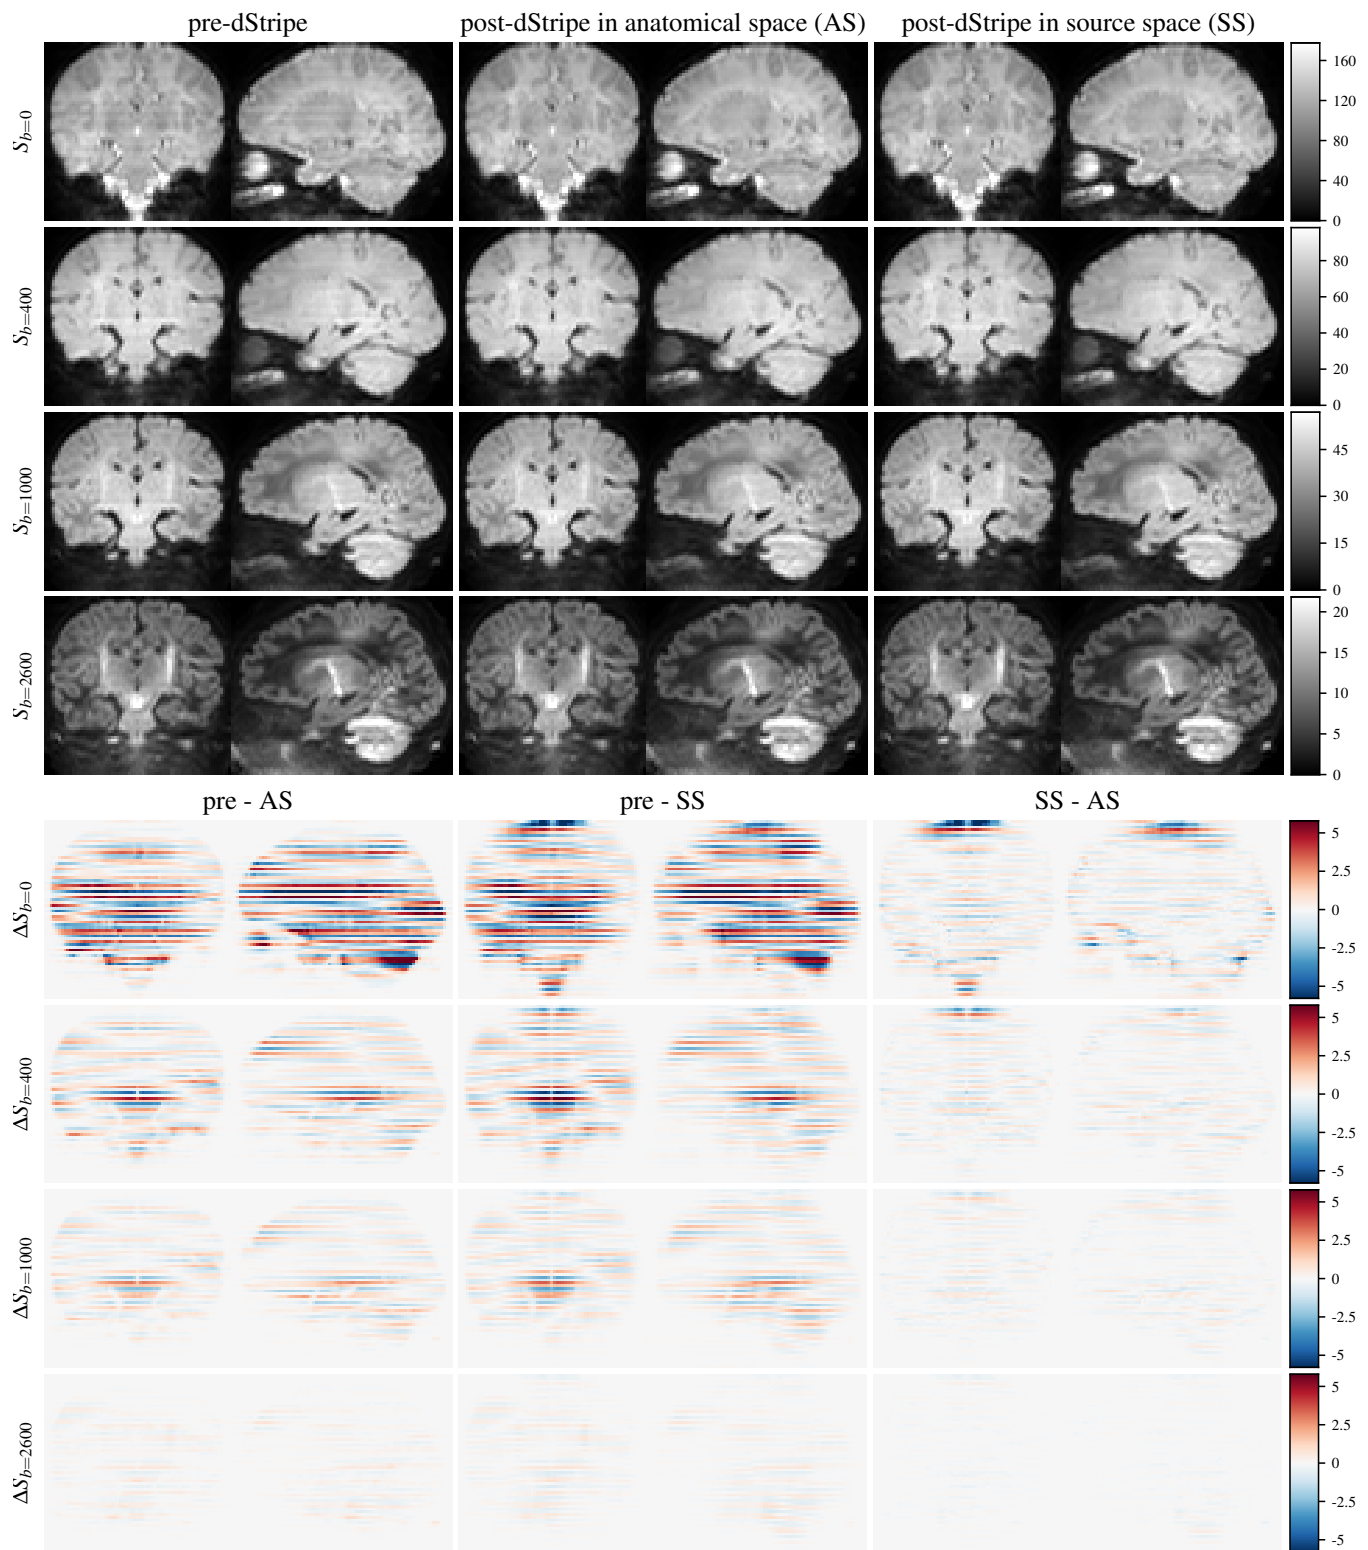


Fig. 3. Exemplary shell-average signal pre- and post-dStripe and shell-average signal difference images displayed in motion-corrected anatomical-space.

1.61%, g: -0.64%, h: -0.97%, i: -2.1%, j: -0.24%) indicating higher data consistency across subjects after dStripe.

Performing dStripe in anatomical space or in source space yields close to identical results but source-space dStripe exhibits a higher dispersion of the voxel-wise population-average tissue volume fractions and exhibits slightly wider

inter-subject standard deviation (compare lines in figure 8).

3.4.2. Spatially-resolved analysis. Here we use data from the relatively homogeneous group of 32 subjects selected for template creation to investigate whether the application of dStripe causes systematic changes (bias) or increased variance of the dMRI signal in certain anatomical locations. Af-

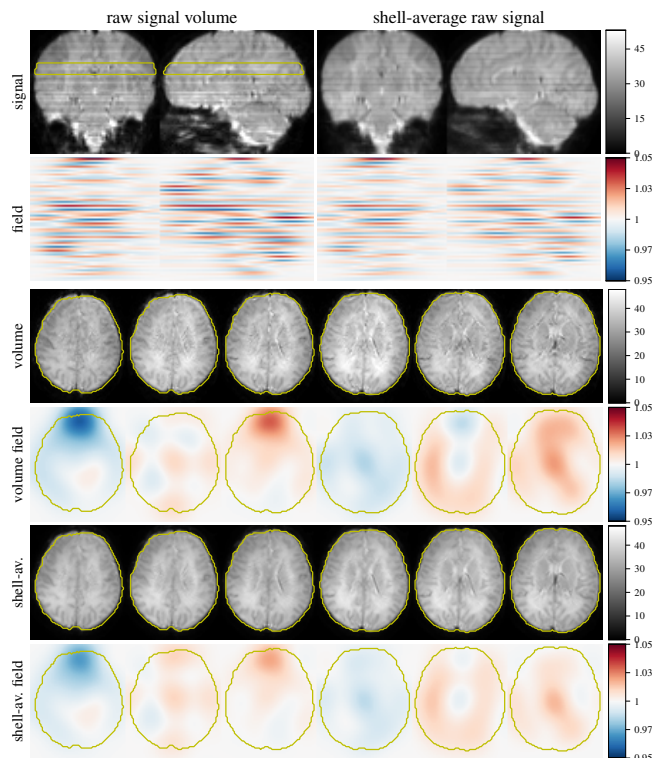


Fig. 4. Signal and slide modulation field in scattered source space. Shown is an individual $b=0$ volume with relatively little motion and the $b=0$ source-space average signal as well as dStripe fields estimated in source space through-plane and inplane. Note that the averaged signal contains signal dropout and other artefacts that are suppressed by the motion correction algorithm and are therefore not accounted for in the dStripe field. Through-plane frequency constraints apply in "slice-native" space, high-frequency variations are not enforced in source space shown here.

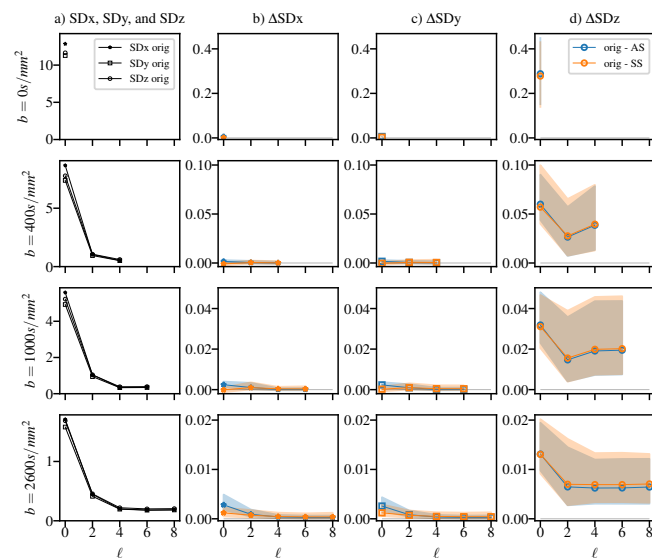


Fig. 5. The local standard deviation of the DWI signal in the SH-basis in plane (SDx, SDy) and through-plane (SDz) by harmonic order of the original data after reconstruction (a) and the subject-wise change after destriping (pre-post) in anatomical space (AS) and source space (SS). The through-plane SD is reduced throughout the spectrum, the in-plane SD is approximately preserved. Both dStripe methods perform similarly, SS dStripe is slightly superior for data with larger residuals. $n=44$.

ter alignment of the data to the population-average template space, the heterogeneity of subject position in source-space should dampen the influence of any subject-specific stripe




Fig. 6. Shell-specific root mean squared DT fit residuals R (A) and DT and MT CSD fit-derived quantities (B). For both fits, dStripe reduces the inter-slice variation within the residuals and in the DT-derived mean diffusivity (MD), as well as CSD-derived fluid and tissue fraction maps. Fractional anisotropy (FA) and tissue anisotropy appear visually unchanged.

patterns. Any systematic difference between coregistered and averaged pre-dStripe data and post-dStripe data beyond random variations due to the finite number of images and subject poses in source-space indicates either systematic anatomical bias of the method or a dependence of the stripe artefact on subject anatomy.

Figure 9 shows the corresponding images of the population-average (a), its absolute change (b) and the change in within-template standard deviation (d) due to dStripe as well as a map of the within-subject standard deviation across subjects

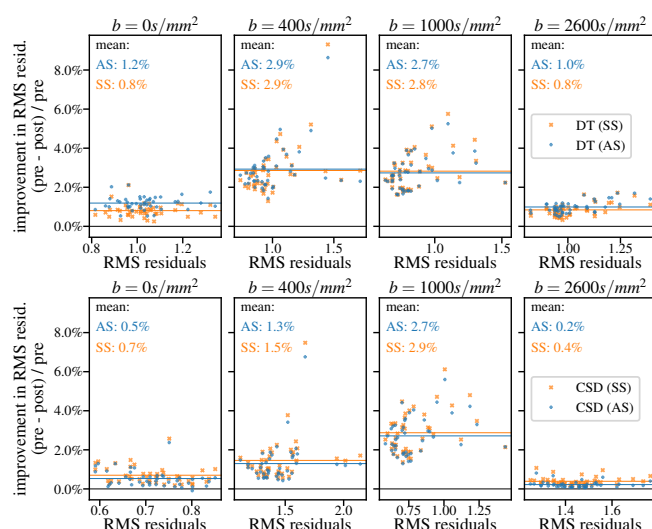


Fig. 7. Relative improvement of the shell-average root mean square residuals for diffusion tensor (top row) and MT CSD (bottom row) fits after destriping compared to fitting non-destriped data for $n=44$ subjects.

in template space (c). We use these to spatially assess variance and bias associated with dStripe.

In the majority of locations, the inter-subject variance is reduced by dStripe, in particular in the tissue $\ell = 4$ band (figure 9d), indicating that dStripe increases consistency throughout the brain and, despite the spatial inhomogeneity of the across-subject standard deviation of signal changes (figure 9c), does not introduce inconsistencies at specific locations. Within individual subjects, the fluid component is more affected by dStripe than the tissue component (figure 9c) but no such systematic effect is observable in the population average, indicating that the variance due to the stripe modulation in the fluid component is larger than in the tissue component.

In the template, systematic effects across subjects due to dStripe (bias) are small (figure 9b) compared to the inter-subject standard deviation of within-subject signal changes (figure 9c). This shows that across subjects, the changes due to dStripe tend to be independent of anatomical location. The most prominent local effect on the average compartment volume fractions occur superior to the body of the corpus callosum and in the vicinity of the pons, medulla and cerebellum and the superior frontal lobe, areas prone to pulsation and distortion artefacts. A consistent local reduction of angular power can be observed superior to the body of the corpus callosum (cingulum and parts of the fornix), most prominent in the template in the tissue $\ell = 2$ and $\ell = 4$ and FA images (b). This systematic reduction in angular power in the population-average template resembles a feature in appearance and amplitude that can be observed in the angular spectrum of the unregistered cohort-average ($N=700$) raw data signal. This creates some ambiguity about whether and how much of this change should be attributed to anatomical bias of the network rather than a genuine need for consistent change as our reference point for zero expected change in the population-average space might not be neutral in this area.

3.5. Limitations. The focus of this work is to develop and evaluate a framework that allows removing stripe artefacts from motion corrupted data without access to ground truth data. Our algorithm choices were guided by observations within our data and are potentially tied to the choice of SVR algorithm. The dependency on the motion corrected prediction can limit the destriping of the source data when subject motion causes stripe patterns to be absent in the anatomical-space – they can be either spatially smeared out, averaged with other data, or removed by the outlier detection mechanism.

This study is a first step at modelling inter-slice intensity modulations. We do not model interactions across slices, volumes (directions) and we do not explicitly use the motion trajectory in the dStripe model. Inter-slice interactions, for instance caused by inhomogeneous g-factor maps biasing the multiband reconstruction, could require unmixing of the signal, which would break the modulation field assumptions made in this work. This and longer-range bias across multiband pack boundaries can be explored in future work.

We do not take tissue-dependencies of the stripe artefact into account. However, if needed, the signal could be decomposed into distinct compartments that separate long T1 from short T1 species, each destriped independently using the same approach as described here, and subsequently recombined. This approach requires the ability to robustly decompose the signal into components relevant for the artefact – in the presence of stripe artefacts.


While our technique allows training on the data of interest, as with other deep learning techniques, the transferability to other and abnormal data remain to be investigated (60).

4. Conclusions and outlook

We presented a data-driven method for the removal of stripe artefacts from dMRI data. dStripe reduces stripe artefacts from the shell-average and the angular signal components, and thus decreasing DTI and MT CSD fit residuals across shells. Single-subject component and DTI-derived images appear visually less stripy and inter-subject variation is reduced indicating improved data consistency across subjects. Applying the dStripe approach in source space (which is likely the space of origin of the slice modulation artefact) slightly outperforms anatomical-space destriping in terms of the reduction of inter-slice signal variation and of signal representation fit residuals. However, it is limited in its applicability as it requires control over the SVR framework's forward-projection and comes with a high computational cost. We showed that anatomical-space dStripe is a suitable substitute as it produces similar results and performs equally in terms of data consistency. For cohorts with less motion we expect the results of anatomical-space dStripe to be even more favourable, in particular for data acquired using lower multiband factors.

While our pipeline is optimised for diffusion MRI, it could be adapted for other modalities and applications such as the detection and removal of stripe artefacts in transcranial ultrasound imaging (61), line and area levelling in atomic force

A) relative template average change



B) relative inter-subject standard deviation change




Fig. 8. 2D histograms of the fluid and tissue components densities and angular energy spectra (band-specific l_2 norm of the spherical harmonic coefficients) using non-destriped data (pre-) and data dstriped in anatomical-space (post-). A shows the population-average change due to dStripe (comparing the template images), B compares the within-template voxel-wise standard deviation across subjects between original and destriped data. Boxplots summarise the 5, 25, 50, 75, and 95 percentile and mean (+) in each band. White lines demark the (conditional) 5%, 50%, 95% probability density of the 2D histograms, yellow dotted lines show the results for source-space destriping.

microscopy (62, 63), or venetian blind artefact removal in multiple overlapping thin slab acquisitions (64).

5. Acknowledgements

This work received funding from the European Research Council under the European Union's Seventh Framework Programme ([FP7/2007-2013/ERC] grant agreement no. [319456] dHCP project), and was supported by the Wellcome/EPSRC Centre for Medical Engineering at King's College London [WT 203148/Z/16/Z]; the Medical Research Council [MR/K006355/1] and by the National Institute for Health Research (NIHR) Biomedical Research Centre based at Guy's and St Thomas' NHS Foundation Trust and King's College London. The views expressed are those of the author(s) and not necessarily those of the NHS, the NIHR or the Department of Health.

- Peter Mansfield. Multi-planar image formation using nmr spin echoes. *Journal of Physics C: Solid State Physics*, 10(3):L55, 1977. doi: 10.1088/0022-3719/10/3/004.
- Wenchuan Wu and Karla L Miller. Image formation in diffusion mri: a review of recent technical developments. *Journal of Magnetic Resonance Imaging*, 46(3):646–662, 2017. doi: 10.1002/jmri.25664.
- Deborah Burstein. Stimulated echoes: Description, applications, practical hints. *Concepts in Magnetic Resonance*, 8(4):269–278, 1996. doi: 10.1002/(sici)1099-0534(1996)8:4<269::aid-cmr3>3.0.co;2-x.
- AP Crawley and RM Henkelman. A stimulated echo artifact from slice interference in magnetic resonance imaging. *Medical physics*, 14(5):842–848, 1987. doi: 10.1118/1.596011.
- G McGibney, MR Smith, ST Nichols, and A Crawley. Quantitative evaluation of several partial fourier reconstruction algorithms used in mri. *Magnetic resonance in medicine*, 30(1):51–59, 1993. doi: 10.1002/mrm.1910300109.
- David J Larkman and Rita G Nunes. Parallel magnetic resonance imaging. *Physics in Medicine & Biology*, 52(7):R15, 2007. doi: 10.1088/0031-9155/52/7/r01.
- Klaas P Pruessmann, Markus Weiger, Markus B Scheidegger, and Peter Boesiger. Sense: sensitivity encoding for fast mri. *Magnetic resonance in medicine*, 42(5):952–962, 1999. doi: 10.1002/(sici)1522-2594(199911)42:5<952::aid-mrm16>3.0.co;2-s.
- Mark A Griswold, Peter M Jakob, Robin M Heidemann, Mathias Nitka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized autocalibrating partially parallel acquisitions (grappa). *Magnetic Resonance in Medicine: An Official Journal of the*

International Society for Magnetic Resonance in Medicine, 47(6):1202–1210, 2002. doi: 10.1002/mrm.10171.

- David J Larkman, Joseph V Hajnal, Amy H Herlihy, Glyn A Coutts, Ian R Young, and Gösta Ehnholm. Use of multicoil arrays for separation of signal from multiple slices simultaneously excited. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 13(2):313–317, 2001. doi: 10.1002/1522-2586(200102)13:2<313::aid-jmri1045>3.0.co;2-w.
- Markus Barth, Felix Breuer, Peter J Koopmans, David G Norris, and Benedikt A Poser. Simultaneous multislice (sms) imaging techniques. *Magnetic resonance in medicine*, 75(1):63–81, 2016. doi: 10.1002/mrm.25897.
- Kawin Setsompop, Borjan A Gagoski, Jonathan R Polimeni, Thomas Witzel, Van J Wedeen, and Lawrence L Wald. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magnetic resonance in medicine*, 67(5):1210–1224, 2012. doi: 10.1002/mrm.23097.
- Martin Blaimer, Morwan Choli, Peter M Jakob, Mark A Griswold, and Felix A Breuer. Multi-band phase-constrained parallel mri. *Magnetic resonance in medicine*, 69(4):974–980, 2013. doi: 10.1002/mrm.24685.
- Stephen F Cauley, Jonathan R Polimeni, Himanshu Bhat, Lawrence L Wald, and Kawin Setsompop. Interslice leakage artifact reduction technique for simultaneous multislice acquisitions. *Magnetic resonance in medicine*, 72(1):93–102, 2014. doi: 10.1002/mrm.24898.
- Nick Todd, Steen Moeller, Edward J Auerbach, Essa Yacoub, Guillaume Flandin, and Nikolaus Weiskopf. Evaluation of 2d multiband epi imaging for high-resolution, whole-brain, task-based fmri studies at 3t: sensitivity and slice leakage artifacts. *Neuroimage*, 124:32–42, 2016. doi: 10.1016/j.neuroimage.2015.08.056.
- Carolyn Beth McNabb, Michael Lindner, Shan Shen, Kou Murayama, Laura G Burgess, and Tom Johnstone. Inter-slice leakage and intra-slice aliasing in simultaneous multi-slice echo-planar images. *bioRxiv*, page 458042, 2018. doi: 10.1101/458042.
- Michael E Kelly, Eugene P Duff, Janine D Bijsterbosch, Natalie L Voets, Nicola Filippini, Steen Moeller, Junqian Xu, Essa S Yacoub, Edward J Auerbach, Kamil Ugurbil, et al. An assessment of motion artefacts in multi band epi for high spatial and temporal resolution resting state fmri. In *Proceedings of the ISMRM, Abstract*, volume 3275, 2013.
- Lucian Muresan, Remco Renken, Jos BTM Roerdink, and Hendrikus Duifhuis. Automated correction of spin-history related motion artefacts in fmri: simulated and phantom data. *IEEE transactions on biomedical engineering*, 52(8):1450–1460, 2005. doi: 10.1109/tbme.2005.851484.
- Karl J Friston, Steven Williams, Robert Howard, Richard SJ Frackowiak, and Robert Turner. Movement-related effects in fmri time-series. *Magnetic resonance in medicine*, 35(3):346–355, 1996. doi: 10.1002/mrm.1910350312.
- Sadie E Yancey, David J Rotenberg, Fred Tam, Mark Chiew, Shawn Ranieri, L Biswas, Kevan JT Anderson, S Nicole Baker, Graham A Wright, and Simon J Graham. Spin-history artifact during functional mri: potential for adaptive correction. *Medical physics*, 38(4):4634–4646, 2011. doi: 10.1118/1.3583814.
- Mandeep S Dagli, John E Ingeholm, and James V Haxby. Localization of cardiac-induced signal change in fmri. *Neuroimage*, 9(4):407–415, 1999. doi: 10.1006/nimg.1998.0424.
- Kevin Murphy, Rasmus M Birn, and Peter A Bandettini. Resting-state fmri confounds and

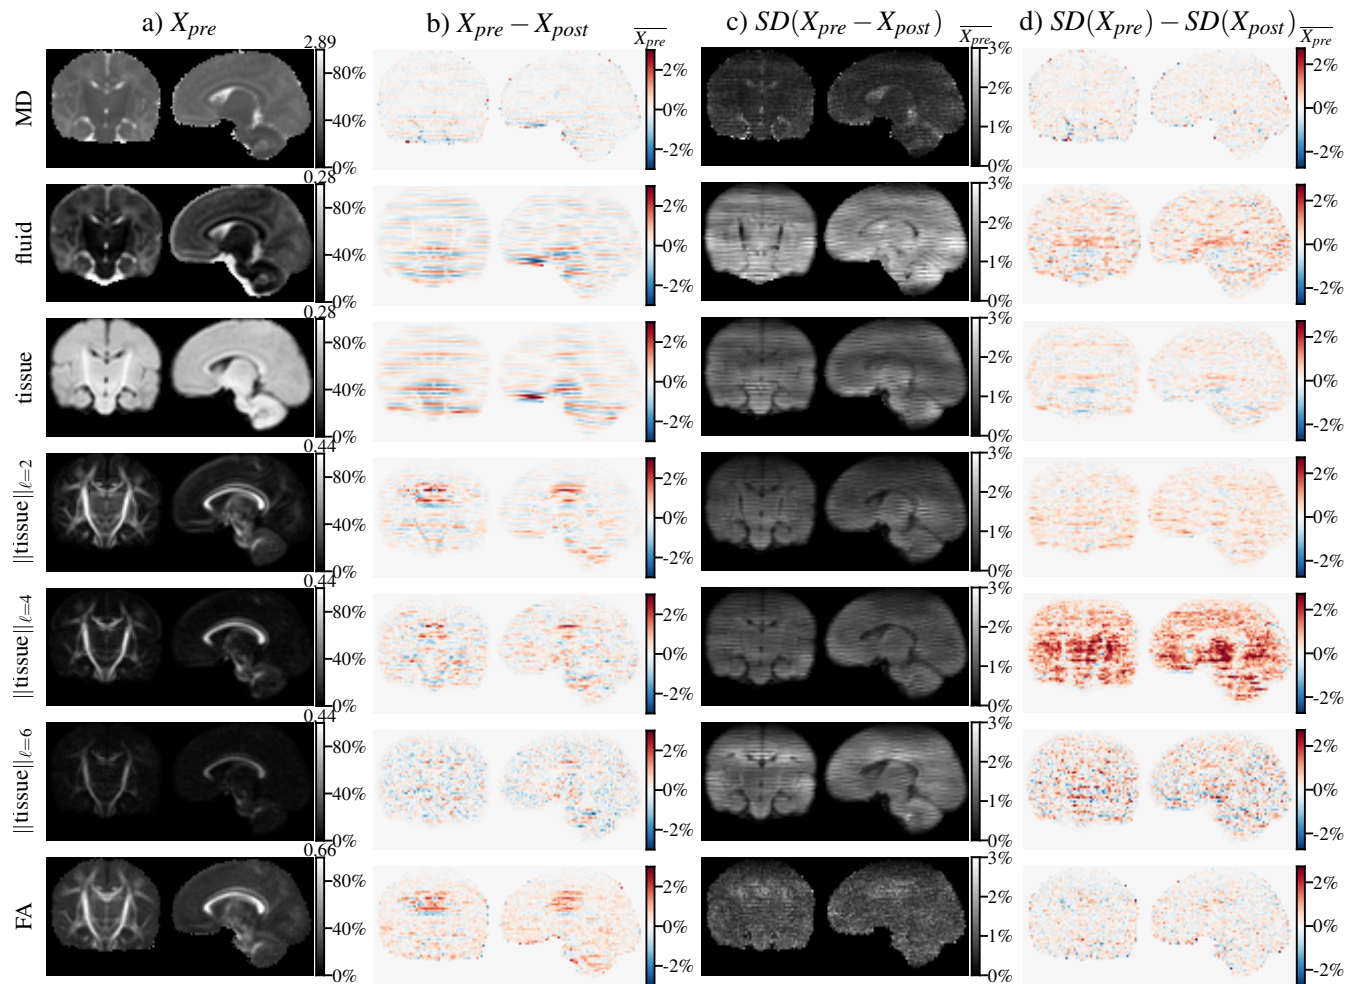


Fig. 9. Effect of dStripe on diffusion tensor and MT CSD-derived quantities aligned to the population-average template ($n=32$). The template (population average \bar{X}) pre-dStripe is shown in column a) and intensity values in all other columns are expressed relative to the maximum intensity in a). The effect of dStripe on the population-average is shown in b). Column c) displays the standard deviation (SD) of the within-subject changes across subjects due to dStripe, indicating the expected effect on a subject-level. The change to the within-group heterogeneity measured as the difference between standard deviation across aligned subjects pre- and post-dStripe data is shown in column d).

- cleanup. *NeuroImage*, 80:349–359, 2013. doi: 10.1016/j.neuroimage.2013.04.001.
22. Jesper LR Andersson and Stamatios N Sotiropoulos. An integrated approach to correction for off-resonance effects and subject movement in diffusion mr imaging. *NeuroImage*, 125: 1063–1078, 2016. doi: 10.1016/j.neuroimage.2015.10.019.
23. M Meyer, A Biber, and MA Koch. Chasing the zebra: the quest for the origin of a stripe artifact in diffusion-weighted mri. In *Student Conference Medical Engineering Science 2014: Proceedings*, volume 3, page 199. GRIN Verlag, 2014.
24. S David, A Heemskerk, M Viergever, and A Leemans. Closing the venetian blinds: A processing strategy for correcting stripe artifacts in diffusion mri data. In *ESMRMB Annual Scientific Meeting*, page 215, 2017.
25. Mathias Engström and Stefan Skare. Diffusion-weighted 3d multislab echo planar imaging for high signal-to-noise ratio efficiency and isotropic image resolution. *Magnetic resonance in medicine*, 70(6):1507–1514, 2013. doi: 10.1002/mrm.24594.
26. Robert Frost, Karla L Miller, Rob HN Tijssen, David A Porter, and Peter Jezzard. 3d multislab diffusion-weighted readout-segmented epi with real-time cardiac-reordered k-space acquisition. *Magnetic resonance in medicine*, 72(6):1565–1579, 2014. doi: 10.1002/mrm.25062.
27. Anh T Van, Diego Hernando, and Bradley P Sutton. Motion-induced phase error estimation and correction in 3d diffusion tensor imaging. *IEEE transactions on medical imaging*, 30(11):1933–1940, 2011. doi: 10.1109/tmi.2011.2158654.
28. Dennis L Parker, Chun Yuan, and Duane D Blatter. Mr angiography by multiple thin slab 3d acquisition. *Magnetic resonance in medicine*, 17(2):434–451, 1991. doi: 10.1002/mrm.1910170215.
29. Anh T Van, Murat Aksoy, Samantha J Holdsworth, Daniel Kopeinigg, Sjoerd B Vos, and Roland Bammer. Slab profile encoding (pen) for minimizing slab boundary artifact in three-dimensional diffusion-weighted multislab acquisition. *Magnetic resonance in medicine*, 73(2):605–613, 2015. doi: 10.1002/mrm.25169.
30. Wenchuan Wu, Peter J Koopmans, Robert Frost, and Karla L Miller. Reducing slab boundary artifacts in three-dimensional multislab diffusion mri using nonlinear inversion for slab profile encoding (npen). *Magnetic resonance in medicine*, 76(4):1183–1195, 2016. doi: 10.1002/mrm.26027.
31. Kavin Setsompop, Qiuyun Fan, Jason Stockmann, Berkin Bilgic, Susie Huang, Stephen F Cauley, Aapo Nummenmaa, Fuyixue Wang, Yogesh Rath, Thomas Witzel, et al. High-resolution in vivo diffusion imaging of the human brain with generalized slice dithered enhanced resolution: Simultaneous multislice (g slider-sms). *Magnetic resonance in medicine*, 79(1):141–151, 2018. doi: 10.1002/mrm.26653.
32. Congyu Liao, Jason Stockmann, Qiuyun Fan, Berkin Bilgic, Mary Kate Manhard, Lawrence L Wald, and Kavin Setsompop. High-fidelity, high-isotropic resolution diffusion imaging through gslider acquisition with b1+ and t1 corrections and multi-coil b0 shim array. *arXiv preprint arXiv:1811.05473*, 2018.
33. Maximilian Pietsch, Daan Christiaens, J-Donald Tournier, and Joseph V Hajnal. D-stripe: correction for stripe artefacts in diffusion mri using a combined deep neural network and svr approach. In *Proceedings of the ISMRM, Abstract*, volume 3411, 2019.
34. Jana Hutter, J Donald Tournier, Anthony N Price, Lucilio Cordero-Grande, Emer J Hughes, Shaihan Malik, Johannes Steinweg, Matteo Bastiani, Stamatios N Sotiropoulos, Saad Jbabdi, et al. Time-efficient and flexible design of optimized multishell hardi diffusion. *Magnetic resonance in medicine*, 79(3):1276–1292, 2018. doi: 10.1002/mrm.26765.
35. Emer J Hughes, Tobias Winchman, Francesco Padormo, Rui Teixeira, Julia Wurie, Maryanne Sharma, Matthew Fox, Jana Hutter, Lucilio Cordero-Grande, Anthony N Price, Joanna Allsop, Jose Bueno-Conde, Nora Tusor, Tomoki Arichi, A D Edwards, Mary A Rutherford, Serena J Counsell, and Joseph V Hajnal. A dedicated neonatal brain imaging system. *Magnetic resonance in medicine: official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 78(2):794–804, August 2017. ISSN 0740-3194, 1522-2594. doi: 10.1002/mrm.26462.
36. J-Donald Tournier, Daan Christiaens, Jana Hutter, Anthony N Price, Lucilio Cordero-Grande, Emer Hughes, Matteo Bastiani, Stamatios N Sotiropoulos, Stephen M Smith, Daniel Rueckert, et al. A data-driven approach to optimising the encoding for multi-shell diffusion mri with application to neonatal imaging. *NMR in Biomedicine*, 33(9):e4348, 2020. doi: 10.1002/nbm.4348.
37. Jelle Veraart, Dmitry S Novikov, Daan Christiaens, Benjamin Ades-Aron, Jan Sibjers, and Els Fieremans. Denoising of diffusion mri using random matrix theory. *NeuroImage*, 142: 394–406, 2016. doi: 10.1016/j.neuroimage.2016.08.016.
38. Jesper LR Andersson, Stefan Skare, and John Ashburner. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *Neu-*

- roimage*, 20(2):870–888, 2003. doi: 10.1016/s1053-8119(03)00336-7.
39. Stephen M Smith. Fast robust automated brain extraction. *Human brain mapping*, 17(3): 143–155, 2002. doi: 10.1002/hbm.10062.
40. Daan Christiaens, Lucilio Cordero-Grande, Maximilian Pietsch, Jana Hutter, Anthony N. Price, Emer J. Hughes, Katy Vecchiato, Maria Deprez, A. David Edwards, Joseph V. Hajnal, and J-Donald Tournier. Scattered slice SHARD reconstruction for motion correction in multi-shell diffusion MRI. *NeuroImage*, page 117437, oct 2020. doi: 10.1016/j.neuroimage.2020.117437.
41. Daan Christiaens, Lucilio Cordero-Grande, Jana Hutter, Anthony N Price, Maria Deprez, Joseph V Hajnal, and J-Donald Tournier. Learning compact q-space representations for multi-shell diffusion-weighted mri. *IEEE transactions on medical imaging*, 38(3):834–843, 2019. doi: 10.1109/TMI.2018.2873736.
42. mrtrix.readthedocs.io. *Spherical Harmonics documentation*, 2020.
43. Michael Kazhdan, Thomas Funkhouser, and Szymon Rusinkiewicz. Rotation invariant spherical harmonic representation of 3 d shape descriptors. In *Symposium on geometry processing*, volume 6, pages 156–164, 2003.
44. Peter J Basser, James Mattiello, and Denis LeBihan. Mr diffusion tensor spectroscopy and imaging. *Biophysical journal*, 66(1):259–267, 1994. doi: 10.1016/s0006-3495(94)80775-1.
45. Ben Jeurissen, Jacques-Donald Tournier, Thijs Dhollander, Alan Connelly, and Jan Sijbers. Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion mri data. *NeuroImage*, 103:411–426, 2014. doi: 10.1016/j.neuroimage.2014.07.061.
46. T Dhollander, R Mito, D Raffelt, and A Connelly. Improved white matter response function estimation for 3-tissue constrained spherical deconvolution. In *Proc. Intl. Soc. Mag. Reson. Med*, page 555, 2019.
47. Maximilian Pietsch, Daan Christiaens, Jana Hutter, Lucilio Cordero-Grande, Anthony N Price, Emer Hughes, A David Edwards, Joseph V Hajnal, Serena J Counsell, and J-Donald Tournier. A framework for multi-component analysis of diffusion mri data over the neonatal period. *NeuroImage*, 186:321–337, 2019. doi: 10.1016/j.neuroimage.2018.10.060.
48. Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018. doi: 10.1109/cvpr.2018.00745.
49. Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.
50. Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
51. Shiban Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. How does batch normalization help optimization? In *Advances in Neural Information Processing Systems*, pages 2483–2493, 2018.
52. Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
53. Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
54. Leslie N Smith. Cyclical learning rates for training neural networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 464–472. IEEE, 2017. doi: 10.1109/wacv.2017.58.
55. Charles M Stein. Estimation of the mean of a multivariate normal distribution. *The annals of Statistics*, pages 1135–1151, 1981. doi: 10.1214/aos/1176345632.
56. Shakarim Soltanayev and Se Young Chun. Training deep learning based denoisers without ground truth data. In *Advances in Neural Information Processing Systems*, pages 3257–3267, 2018.
57. Christopher A Metzler, Ali Mousavi, Reinhard Heckel, and Richard G Baraniuk. Unsupervised learning with stein’s unbiased risk estimator. *arXiv preprint arXiv:1805.10531*, 2018.
58. Bradley Efron. The estimation of prediction error: covariance penalties and cross-validation. *Journal of the American Statistical Association*, 99(467):619–632, 2004. doi: 10.1198/016214504000000692.
59. Sathish Ramani, Thierry Blu, and Michael Unser. Monte-carlo sure: A black-box optimization of regularization parameters for general denoising algorithms. *IEEE Transactions on image processing*, 17(9):1540–1554, 2008. doi: 10.1109/tip.2008.2001404.
60. Vegard Antun, Francesco Renna, Clarice Poon, Ben Adcock, and Anders C Hansen. On instabilities of deep learning in image reconstruction and the potential costs of ai. *Proceedings of the National Academy of Sciences*, 2020. doi: 10.1073/pnas.1907377117.
61. Francois Vignon, William T Shi, Xiangtao Yin, Thilo Hoelscher, and Jeffrey E Powers. The stripe artifact in transcranial ultrasound imaging. *Journal of Ultrasound in Medicine*, 29(12): 1779–1786, 2010. doi: 10.7863/jum.2010.29.12.1779.
62. M Raposo, Q Ferreira, and PA Ribeiro. A guide for atomic force microscopy analysis of soft-condensed matter. *Modern research and educational topics in microscopy*, 1:758–769, 2007.
63. Yuliang Wang, Tongda Lu, Xiaolai Li, and Huimin Wang. Automated image segmentation-assisted flattening of atomic force microscopy images. *Beilstein journal of nanotechnology*, 9(1):975–985, 2018. doi: 10.3762/bjnano.9.91.
64. DD Blatter, DL Parker, and RO Robison. Cerebral mr angiography with multiple overlapping thin slab acquisition. part i. quantitative analysis of vessel visibility. *Radiology*, 179(3):805–811, 1991. doi: 10.1148/radiology.179.3.2027996.

A. Data overview




Fig. 10. Top: sagittal cross-sections through b=0 data pre-dStripe (left) and post-dStripe (right) for the 44 diffusion MRIs used for the analysis of the dStripe method. Images are sorted by percentile quality assessment score (higher is better). Bottom: Age at scan versus percentile quality assessment score. Data used for single-subject analysis and in the template is marked as such.