NeuroImage 125 (2016) 1079–1094

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Realistic simulation of artefacts in diffusion MRI for validating post-processing correction techniques

Mark S. Graham *, Ivana Drobnjak, Hui Zhang

Centre for Medical Image Computing & Department of Computer Science, University College London, UK

ARTICLE INFO

Article history:
Received 18 May 2015
Accepted 4 November 2015
Available online 10 November 2015

Keywords:
Diffusion MRI
Simulation
Artefact correction
Eddy currents
Motion
POSSUM

ABSTRACT

In this paper we demonstrate a simulation framework that enables the direct and quantitative comparison of post-processing methods for diffusion weighted magnetic resonance (DW-MR) images. DW-MR datasets are employed in a range of techniques that enable estimates of local microstructure and global connectivity in the brain. These techniques require full alignment of images across the dataset, but this is rarely the case. Artefacts such as eddy-current (EC) distortion and motion lead to misalignment between images, which compromise the quality of the microstructural measures obtained from them. Numerous methods and software packages exist to correct these artefacts, some of which have become de-facto standards, but none have been subject to rigorous validation. In the literature, improved alignment is assessed using either qualitative visual measures or quantitative surrogate metrics. Here we introduce a simulation framework that allows for the direct, quantitative assessment of techniques, enabling objective comparisons of existing and future methods. DW-MR datasets are generated using a process that is based on the physics of MRI acquisition, which allows for the salient features of the images and their artefacts to be reproduced. We apply this framework in three ways. Firstly we assess the most commonly used method for artefact correction, FSL's eddy_correct, and compare it to a recently proposed alternative, eddy. We demonstrate quantitatively that using eddy_correct leads to significant errors in the corrected data, whilst eddy is able to provide much improved correction. Secondly we investigate the datasets required to achieve good correction with eddy, by looking at the minimum number of directions required and comparing the recommended full-sphere acquisitions to equivalent half-sphere protocols. Finally, we investigate the impact of correction quality by examining the fits from microstructure models to real and simulated data.

Introduction

Diffusion-weighted magnetic resonance (DW-MR) imaging is a powerful, non-invasive technique that allows us to probe the microstructure of biological tissue (Assaf and Cohen, 2009). The technique is well suited to the brain, and is used by clinicians and researchers studying its structure in health and disease (Bodini and Ciccarelli, 2009; Scholz et al., 2009).

A typical DW measurement is made by applying a diffusion-sensitising gradient waveform in a particular direction across a sample, before acquiring an MR image. The image contains information on the diffusion of water in this direction. This diffusion is influenced by the underlying microstructure, and by acquiring a range of images with varying gradient strength and direction we can probe this structure. Many techniques, both model-based (Assaf and Basser, 2005; Zhang et al., 2012) and model-free (Tournier et al., 2004; Özarslan et al., 2013) use the rich information provided by such datasets to characterise brain microstructure.

Unfortunately images acquired with DW-MRI are susceptible to a number of artefacts (Le Bihan et al., 2006). For example, susceptibility differences at the air-tissue boundary lead to alterations of the B_0 field that can cause spatial displacements of several pixels. This adversely affects analysis of the DW data itself (Irfanoglu et al., 2012), and prevents comparison between these DW images and others that do not contain these artefacts, such as T1- and T2-weighted images. Some artefacts lead to spatial offsets between the DW-MR images in a dataset, which further undermines the estimates of microstructure obtained from them. For example, motion can lead to rigid offsets between images, and eddy currents (EC) lead to distortions of the image in the phase encoding (PE) direction that vary according to the amount of diffusion sensitisation used (typically summarised by the *b*-value) and the direction it is applied in.

Techniques for dealing with these artefacts can broadly be divided into those implemented at acquisition time (Reese et al., 2003; Jezzard and Balaban, 1995; Andersson et al., 2003), involving either some modification to the acquisition process or the collection of supplementary

* Corresponding author at: Centre for Medical Image Computing, University College London, Gower Street, London, WC1E 7JE, United Kingdom.
E-mail address: mark.graham.13@ucl.ac.uk (M.S. Graham).

http://dx.doi.org/10.1016/j.neuroimage.2015.11.006

data, and post-processing methods implemented after acquisition time (Haselgrove and Moore, 1996; Jenkinson and Smith, 2001). Post-processing techniques are the most widely used, as they have several advantages: they can be applied retrospectively to already acquired data, a user can revert to the original data if the technique does not work as hoped, and they don't require additional scan-time, which is often expensive.

The literature contains a vast body of post-processing techniques and software packages for correcting artefacts in DW-MRI (Oguz et al., 2014; Jenkinson and Smith, 2001; Andersson et al., 2003). Ideally their corrections would be validated by comparison to the ground truth, i.e., a map of the spatial deformations caused by the artefacts, but these cannot be obtained for real data. As a result the literature relies on either qualitative visual assessments of image alignment (Mangin et al., 2002; Horsfield, 1999), or quantitative assessments of surrogate measures of alignment, such as tract length (Zhuang et al., 2013), fractional anisotropy (FA) values (Bai and Alexander, 2008) or reduced residuals from fits to microstructural models (Ben-Amitay et al., 2012; Andersson and Skare, 2002). The lack of an objective ground truth means existing techniques cannot be systematically assessed, preventing end-users from making an informed choice. The development of new methods is also hindered, as any improvements over existing ones are difficult to demonstrate.

Simulation could provide us with a ground truth that would enable us to assess methods objectively, allowing researchers to make informed decisions when selecting post-processing methods. Simulation systems exist for MRI (Kwan et al., 1999; Drobnjak et al., 2006) but there is nothing satisfactory for assessment of artefact correction in DW-MRI. There has been some limited work on the simulation of full-brain DWIs, but they all exhibit at least one of two serious limitations. The first is the failure to model the full process of image acquisition (Bastin, 1998, 2001), i.e., the recording of a signal in frequency space which is Fourier transformed to generate a spatial image, which precludes the inclusion of realistic spatial-encoding artefacts. The second is the use of a heavily simplified model to create the DW contrast, which means the simulations do not capture some of the features of DWIs that makes their processing uniquely challenging, such as the variation of contrast with the direction of diffusion weighting (Nunes et al., 2011; Bastin, 2001). More recently, a number of systems have been developed to simulate DWIs of white matter (WM) bundles (Neher et al., 2013; Close et al., 2009; Caruyer et al., 2014), and are intended to help validate model fitting and tractography techniques. These systems allow for the use of advanced models to predict the diffusion attenuation, and some are able to simulate the full process of image acquisition (Neher et al., 2013). However, these models rely on an underlying description of WM structures to generate the signal, which means they are unable to produce full-brain images that realistically model the signal in the grey matter (GM) and cerebrospinal fluid (CSF). To enable effective validation of artefact correction, there is a need for a simulation system that combines realistic full-brain DW-MR contrast with a modelling of the full MR acquisition process.

In this work we introduce a framework that simulates realistic DW-MR images with artefacts, along with displacement fields that describe the ground truth (GT) correction of each image, enabling the effectiveness of correction techniques to be assessed objectively and directly. DW images are simulated by extending POSSUM, which simulates the physics of MRI acquisition by solving Bloch's and Maxwell's equations (Drobnjak et al., 2006, 2010). This ensures that the images and their artefacts capture the key features of their real-world counterparts. In contrast to existing approaches that use microstructure-based models to estimate the diffusion signal, we use a model-free approach to obtain the signal from real data, allowing the full complexity of the diffusion signal in GM, WM and CSF to be captured. Correction techniques are assessed by comparing the displacement fields they predict to the ground truth fields provided by the framework. We use this framework to compare two EC and motion correction techniques, FSL's eddy_correct and eddy, and demonstrate that eddy_correct provides particularly poor correction. We then turn the framework to a more thorough investigation of the most promising of these techniques, eddy. We evaluate eddy's performance as the number of directions acquired is varied, and also compare the recommended full-shell acquisition to equivalent half-shell protocols, in order to make practical recommendations for its use. Finally we look at the impact of correction quality by examining microstructure fits to real and simulated data. This work has appeared in a preliminary form in Graham et al. (2015). Simulated datasets and code will be made available at http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/POSSUM.

Methods

In this section we describe our simulation framework for producing realistic DW-MR images along with the displacement fields that map them into undistorted space. We first outline the framework in Framework overview. Our implementation of the framework is detailed in Framework implementation — producing DWIs and Framework implementation — ground truth displacement fields.

Framework overview

The framework (Fig. 1) combines a physics-based approach to the MR image acquisition process with a model-free representation of diffusion in order to simulate realistic DW-MR datasets. To provide a meaningful validation a simulation must capture the key characteristics of DW-MR images and their artefacts. Many of the artefacts are introduced during the acquisition of the MR signal in k-space, so a simulation can only be deemed faithful if it can reproduce this process. We emphasise the distinction between this and methods that introduce artefacts by applying geometric transformations in image space. Our method is able to capture non-geometric effects such as blurring due to EC decay.

The framework takes four main inputs. The first is a geometric object that specifies the proton density and location of WM, GM and CSF along with their T1 and T2 values. The second is a representation of diffusion-weighting. The third is a pulsed-gradient spin-echo (PGSE) sequence, detailing RF pulses and gradients (the framework can be straightforwardly extended to take other diffusion sequences, such as twice-refocused spin-echo (Reese et al., 2003)). The first two inputs are combined with diffusion parameters extracted from the third (direction and magnitude of diffusion weighting) to produce a geometric object with its proton density reduced by a diffusion attenuation factor. The PGSE sequence is converted to a standard echo-planar imaging (EPI) sequence for simulation, so the diffusion attenuation is introduced solely through the input object. The fourth input is any details that will lead to the simulation of artefacts, such as motion parameters. The effects of eddy currents are included in the EPI sequence.

The framework creates two outputs. The first is a DWI. The MR simulator takes the attenuated object, pulse sequence and details pertaining to artefacts, and solves Bloch's and Maxwell's equations at each point in the object, summing the resultant signal in order to generate the k-space measurements. This is Fourier transformed to produce the output DWI. The second output is a displacement field that describes the mapping of this DWI from a distorted to undistorted space.

Framework implementation — producing DWIs

We extended the POSSUM (Drobnjak et al., 2006, 2010) MRI simulation in order to produce realistic DW-MR images. POSSUM takes as inputs a pulse sequence and a geometric object. It simulates the full MR acquisition process by solving for the magnetization vectors over time at every voxel in the object, which allows for the faithful reproduction of artefacts that are introduced during acquisition. POSSUM is designed for the simulation of MRI/fMRI data, so a number of extensions were

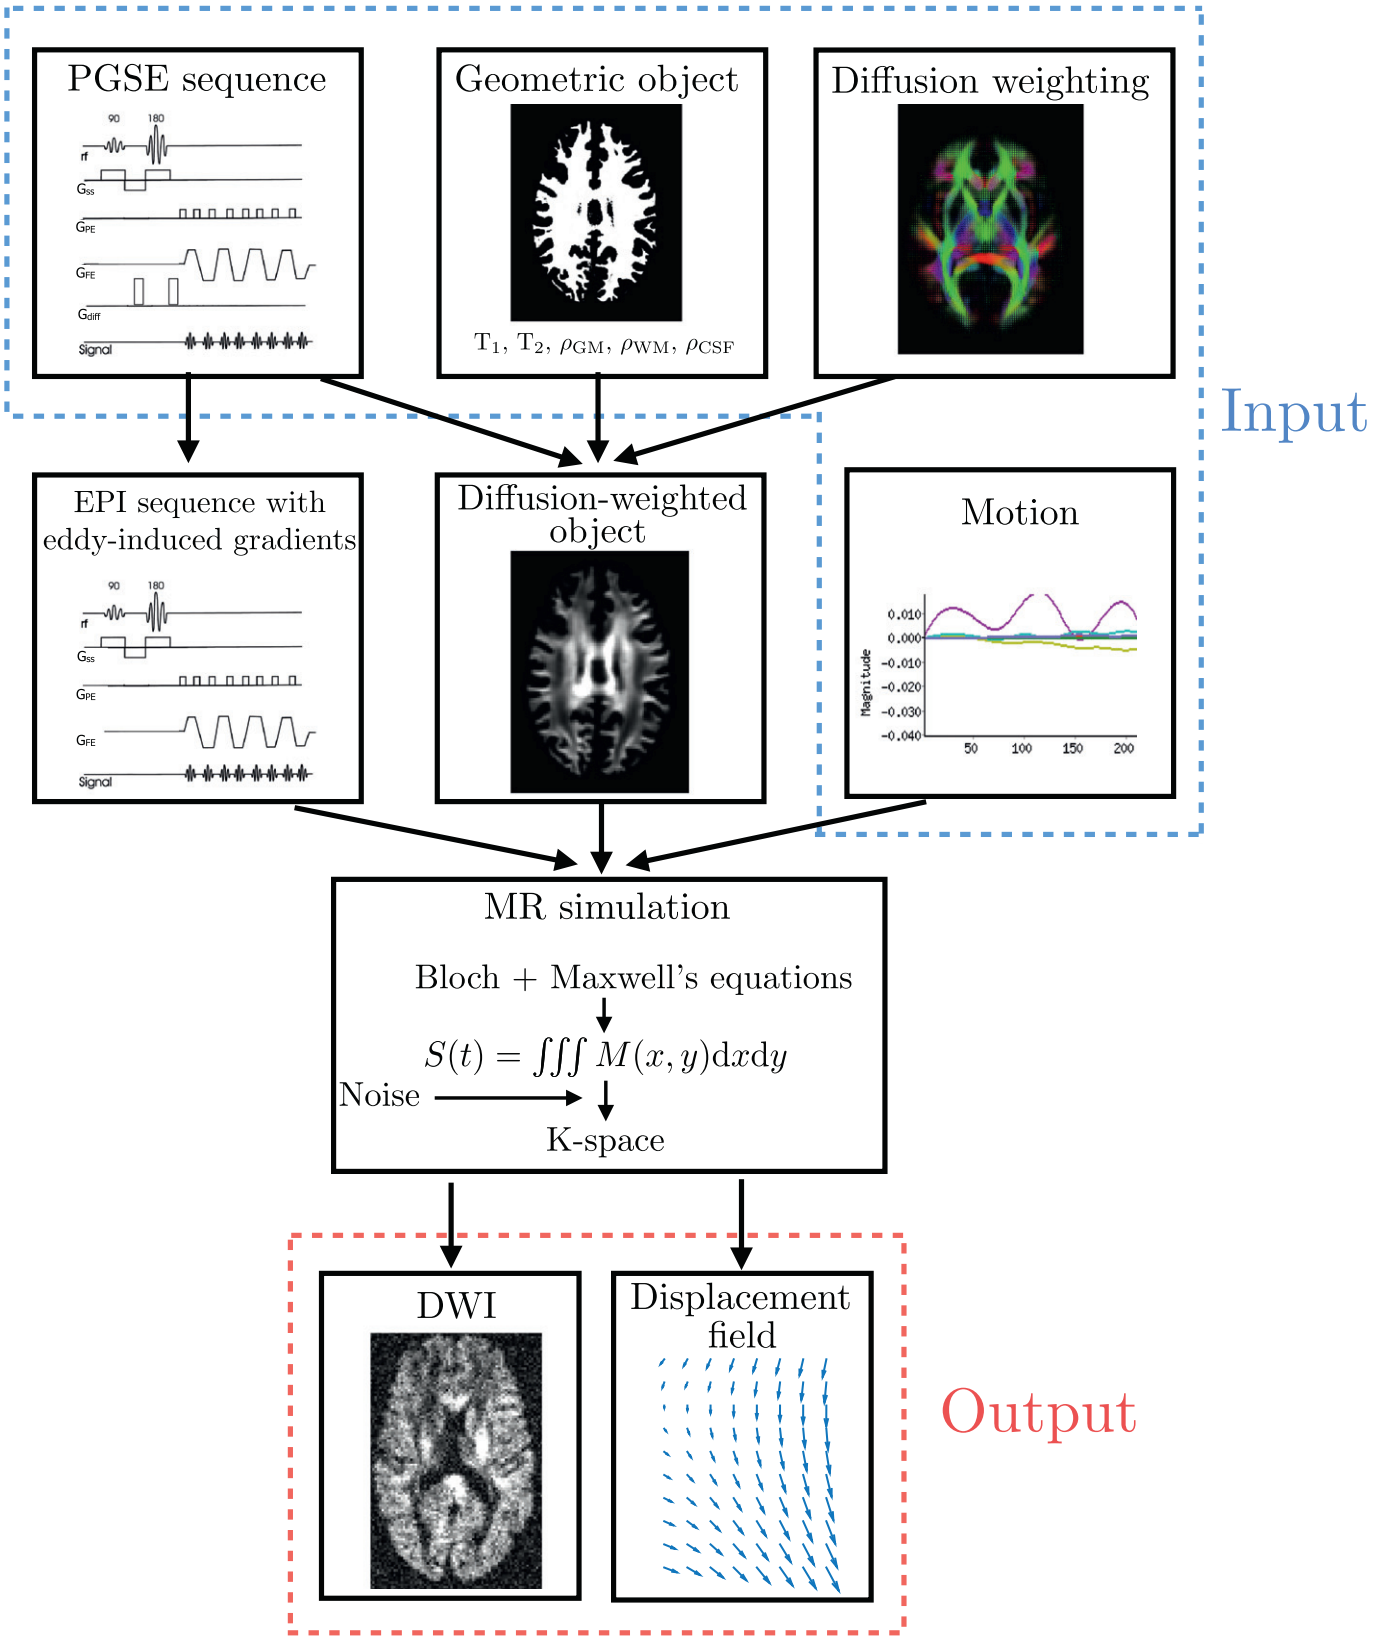

Fig. 1. The pipeline for simulating DWIs. Details of DWI and displacement field generation are in Framework implementation — producing DWIs and Framework implementation — ground truth displacement fields.

made to enable DW-MRI simulations. These include the introduction of diffusion weighting, the addition of EC-induced gradients and the inclusion of spin-echo contrast. They are detailed below.

A full-brain segmentation was used as POSSUM's geometric object input. It was created with T1- and T2-weighted images from a single subject from the WU-Minn HCP dataset (Van Essen et al., 2012), using

FSL's FAST (Zhang et al., 2001). Diffusion weighting was achieved using a voxel-wise spherical harmonic (SH) fit to the subject's diffusion data (Alexander et al., 2002). An order $n = 8$ fit was used to fully capture the angular information available in these datasets (Tournier et al., 2013), with constraints placed on the coefficients to ensure the signal is real and exhibits antipodal symmetry. Separate fits were performed on the $b = 1000$ s/mm^2 and $b = 2000$ s/mm^2 shells. These SH fits enable us to predict the diffusion-weighting at each voxel along any gradient direction, $\hat{b}$. The fit to the $b = 1000$ s/mm^2 shell was used to predict signal attenuation for simulated DWIs with $b = 1000$ s/mm^2, and the $b = 2000$ s/mm^2 fit to predict attenuation for simulations at $b = 2000$ s/mm^2. For some experiments DWIs with $b = 700$ s/mm^2 were simulated. A mono-exponential assumption was used to predict the signal attenuation at $b = 700$ s/mm^2 in direction $\hat{b}$, $A_{700}(\hat{b})$ from the predicted attenuation at $b = 1000$ s/mm^2, $A_{1000}(\hat{b})$:

$$A_{700}\left(\hat{b}\right) = \exp\left(\frac{1000}{700}\ln\left(A_{1000}\left(\hat{b}\right)\right)\right). \tag{1}$$

The prediction of diffusion-weighting at $b = 700$ s/mm^2 is no longer strictly model-free, but relies on minimal assumptions about the form of the diffusion signal that are reasonable in practice (Niendorf et al., 1996).

By default, POSSUM uses gradient-echo EPI sequences. To simulate the spin-echo EPI sequences typically used for DWI acquisition, we replace the default tissue-specific T_2^* values with their corresponding T_2 values: 75 ms and 70 ms in the GM and WM, respectively. These values were obtained by adjusting the average estimates from the literature (Wansapura et al., 1999) until our simulations best matched the contrast seen in a real dataset. All simulations were performed at 3 T. K-space was acquired with a readout bandwidth of 100 kHz, using a linear-ordered, cartesian sampling trajectory with full coverage. We used a matrix size of 72 × 86, which was chosen along with the image voxel size (2.5 isotropic) to strike a balance between minimising computation time and ensuring full-brain coverage. K-space was apodized using a Hamming window, and no zero-filling was performed.

In this work we focus on the addition of EC and instantaneous motion artefacts. Eddy-induced gradients were added to the EPI pulse sequence using the spatially linear model in Nunes et al. (2011), by superposing a sum of decaying exponentials on each gradient field:

$$G^E_{x,y,z} = \sum_i \pm\varepsilon G^{\text{diff}}_{x,y,z}\exp[-(t-t_i)/\tau] \tag{2}$$

where t_i corresponds to the time each diffusion gradient is turned on or off (determined by the pulse width δ and diffusion time Δ, obtained from the input PGSE sequence), τ is the decay time, ε is a constant determining the relationship between the strength of eddy and diffusion gradients and a + or − is selected depending on whether the gradient is being turned on or off. We performed simulations with a maximum value of $G^{\text{diff}} = 40$ mT m^{-1}, and selected $\varepsilon = 0.009$ and $\tau = 100$ ms to represent typical values found in a clinical scanner (Jezzard et al., 1998). Motion was added by applying rigid rotations to the geometric object, before the MR simulation began, in order to recreate instantaneous motion between volumes. Motion was accounted for when applying the diffusion-weighting, such that rotations of the head affected the diffusion contrast. Physically this is equivalent to simulating motion as occurring before the diffusion sensitisation gradients are applied.

Whilst we only demonstrate EC + motion artefacts in this work, we note that the simulations are flexible, and able to recreate many more artefacts. Nyquist ghosting, chemical shifting, motion during signal read-out and Gibbs ringing have all been demonstrated in POSSUM (Drobnjak et al., 2006), and can be readily included in our simulations.

Framework implementation — ground truth displacement fields

The framework generates the GT displacement field for each DWI, detailing its mapping from distorted to undistorted space, as comparison of predicted and GT spatial displacement fields is the most direct way to evaluate post-processing methods. Here we describe how we obtain the combined displacement field that corrects for the EC + motion artefacts in our simulations. The EC displacement fields are calculated in a general way that can be straightforwardly extended to other off-resonance artefacts.

The EC displacement field can be calculated from a knowledge of the off-resonance frequency that the EC give rise to. In the presence of any off-resonance conditions, the spin frequency can be written:

$$\omega(\mathbf{r}, t) = \omega_0 + f(\mathbf{r}, t) \tag{3}$$

where ω_0 is the system frequency and $f(\mathbf{r}, t)$ is the off-resonance term, that may vary with spatial location $\mathbf{r}$ and time t. When the off-resonance term is time-invariant, i.e., $f(\mathbf{r}, t) = f(\mathbf{r})$, we can express the displacement field caused by these off-resonance effects, $\boldsymbol{\psi}^O$, as (Zaitsev et al., 2004):

$$\boldsymbol{\psi}^O(\mathbf{r}) = t_s f(\mathbf{r})\mathbf{p} \tag{4}$$

where t_s is the echo spacing and $\mathbf{p}$ is a vector pointing in the phase-encoding direction with a magnitude equal to the inverse of the k-space line spacing.

In the case of linear EC gradients assumed in our simulations, we can express the off-resonance term:

$$f(\mathbf{r}, \mathbf{t}) = \gamma\mathbf{G}^E(t)\cdot\mathbf{r} \tag{5}$$

where $\mathbf{G}^E = (G^E_x, G^E_y, G^E_z)$ and γ is the gyromagnetic ratio. These gradients are time-varying according to Eq. (2). We evaluate this off-resonance term at the centre of k-space, $t = \text{TE}$, as this dominates the gross structure of the image:

$$f(\mathbf{r}) = \gamma\mathbf{G}^E(\text{TE})\cdot\mathbf{r}. \tag{6}$$

We can now combine Eqs. (4) and (6) to find the displacement field:

$$\boldsymbol{\psi}^O(\mathbf{r}) = \gamma t_s\left(\mathbf{G}^E(\text{TE})\cdot\mathbf{r}\right)\mathbf{p}. \tag{7}$$

The motion displacement field, $\boldsymbol{\psi}^M(\mathbf{r})$ is straightforwardly obtained from the rigid transformation $\mathbf{R}$:

$$\boldsymbol{\psi}^M(\mathbf{r}) = \mathbf{R}\mathbf{r} - \mathbf{r}. \tag{8}$$

We can obtain the total displacement field, $\boldsymbol{\psi}^T(\mathbf{r})$, from a composition of the two fields:

$$\boldsymbol{\psi}^T(\mathbf{r}) = \boldsymbol{\psi}^O \circ \boldsymbol{\psi}^M(\mathbf{r}). \tag{9}$$

The ground truth displacement field gives us our mapping from undistorted to distorted space. An artefact correction method predicts a displacement field, $\boldsymbol{\psi}^P(\mathbf{r})$, that attempts to map a volume from distorted to undistorted space. We define the error field, $\boldsymbol{\psi}^E(\mathbf{r})$, as the displacement field that describes the mapping of each voxel in undistorted space into corrected space:

$$\boldsymbol{\psi}^E(\mathbf{r}) = \boldsymbol{\psi}^P \circ \boldsymbol{\psi}^T(\mathbf{r}). \tag{10}$$

A zero error displacement field indicates perfect correction has been achieved. In this work we make use of these error fields to quantitatively assess the effectiveness of artefact correction schemes.

Experiments and results

In this section we explain our experiment design and show results. Firstly we discuss our validation of the simulations. We then undertake a comparison of two artefact correction methods, eddy_correct and eddy, before evaluating the performance of eddy as a function of the quality of the dataset. Finally we assess the impact of correction by fitting models of microstructure to real and simulated data.

Validation of the simulations

We first assess how well the simulated images capture the most important characteristics of real images. POSSUM has been shown to provide realistic MR simulation without diffusion weighting (Drobnjak et al., 2006, 2010), so here we focus on assessing the simulation of diffusion weighting. In the case of DW-MR the key characteristic is the variation in contrast as the strength and direction of diffusion weighting changes. To test this we compared a real and simulated dataset with identical parameters: a 3 T scanner with three shells, $b = 300/700/2000$ s/mm^2, 8/32/64 directions with 1/4/8 $b = 0$ images, TR/TE = 7500/109 ms. Fig. 2a compares the changes in contrast with varying b-value. Fig. 2b compares changes in contrast with varying direction of b-vector. A comparison with Fiberfox, the current state-of-the-art in model-based approaches, is included in Supplementary material. Finally we demonstrate that our EC model produces realistic artefacts. We simulated EC distorted images with diffusion weighting along the x-, y- and z-axes. Fig. 3 shows the results.

The results demonstrate that our simulation framework is able to capture both the increased attenuation with b-value and the variation in contrast with varying b-vector that are present in real data. Note that the real data and the data used to generate the input object for simulation were obtained from different subjects, so the slices are not perfectly matched, which may account for some of the differences in the appearance of WM tracts. Finally we show that we are able to reproduce EC distortions seen in real data. For linear ECs, gradients along the x-, y- and z-axes should lead to shears, scaling and translations in the images respectively (assuming here that y is the PE direction and z is slice-select). In our model, an applied diffusion gradient in a given direction gives rise to EC gradients in the same direction, so we expect to see pure shears, scaling and translation in the three examples in Fig. 3, which we do.

Comparison of correction methods

In this section we compare the quality of correction obtained from two post-processing techniques. The first is FSL's eddy_correct, which registers each volume in a dataset to a $b=0$ image in order to simultaneously correct for motion and EC distortions. We also test a more sophisticated method, FSL's eddy, which registers each volume to a model-free prediction of how it should look in undistorted space (Andersson and Sotiropoulos, 2015). Firstly we compare them quantitatively using our simulation framework, and then we compare them qualitatively on a real dataset to demonstrate that our findings in simulation are present in real data.

The simulated dataset consists of two shells, $b=700/2000$ s / mm^2, 32/64 directions with 4/8 $b=0$ images, TR/TE = 7500/109 ms, $72 \times 86 \times 55$ with isotropic voxel size of 2.5 mm. Diffusion directions were distributed isotropically on the sphere. EC gradients were added to the pulse sequence according to the model in the Framework implementation — producing DWIs section. One dataset was simulated with just EC artefacts, and one was created with both EC and motion artefacts. In this dataset, a translation along each axis was selected for each volume randomly from the range −5 to 5 mm, in addition to a rotation about each axis taken from the range −5 to 5°. Each dataset had Rician noise added at two different levels to produce datasets with SNR = 10 and SNR = 20, measured on the $b = 0$ images in a region of interest (ROI) in the WM (the centrum semiovale).

The simulated datasets were corrected using eddy_correct and eddy. Default settings were used for eddy_correct: correlation ratio as the similarity measure and trilinear interpolation. Default settings were mostly used for eddy: 1000 voxels for estimating the Gaussian Process hyperparameter, spline interpolation, quadradic first-level modelling of the EC, and no second-level modelling. However we used 10 iterations rather than the default 5, because we sometimes found 5 was not sufficient to ensure convergence. We also performed correction using eddy_correct with normalised mutual information (NMI) as a cost function, to test the claim that it is more robust than other cost-functions (Rohde et al., 2004). Results for correction of the datasets with EC artefacts are shown in Fig. 4a, and that dataset with both EC and motion artefacts are in Fig. 4b. Fig. 5 shows how these displacement field errors are spatially distributed across the brain.

The real dataset was acquired on a Siemens PET-MR 3 T with similar parameters to the simulated dataset: two shells, $b = 700/2000$ s/mm^2, 32/64 directions with 4/8 $b = 0$ images, TR/TE = 7500/103 ms, isotropic voxels of size 2.5 mm. The only differences were the dimensions, here $96 \times 96 \times 55$, and the TE (103 ms here vs 109 ms in the simulations). The SNR of the data was 25, measured on the $b = 0$ images in an ROI in the centrum semiovale. This dataset was also corrected using eddy_correct and eddy with the default settings. Fig. 6 shows the results.

Fig. 4 demonstrates that eddy_correct is unable to correct the data well, even for DWIs acquired with $b = 700$ s/mm^2, whilst eddy is able to provide good correction across the dataset. Volumes corrected with eddy_correct have average errors of one voxel at $b = 700$ s/mm^2, rising to 1.5 voxels at $b = 1000$ s/mm^2. This is consistent with expectations; the increasing contrast differences between DWI and $b = 0$ as b-value is increased makes direct registration progressively worse. These findings are in agreement with previous work (Bastin, 1998) which found that DWIs can only be successfully corrected by registration to $b = 0$ for $b \leq 300$ s/mm^2. By contrast, eddy is able to correct with errors of less than 0.2 voxels across the dataset at SNR = 20, and 0.5 voxels at SNR = 10. Fig. 5 shows the errors are spatially distributed as we would expect. The data corrected with eddy_correct shows a consistent over-scaling of the data. This is likely caused by the attenuation of the CSF rim around the brain in DWIs. This makes the DWIs look smaller then the $b = 0$ images, which causes registration to enlarge them. Interestingly, Fig. 4 shows that at SNR = 10 eddy provides better correction for the dataset with EC and motion than the dataset with EC and no motion. For the dataset with no motion, eddy was not able to detect the EC artefacts, and essentially left the data uncorrected, whilst for the dataset with motion eddy was able to accurately estimate both the motion and EC artefacts. At SNR = 20 eddy was able to estimate the correct parameters for the dataset with EC and no motion. Thus it seems that the combination of low SNR and no motion artefacts caused eddy's optimisation process to fail. We found that both methods showed little sensitivity to noise for SNR> 20: results on a dataset with SNR = 30 (not shown) were very similar to those found on the SNR = 20 data.

The results for real data corroborate with our findings for simulation. Fig. 6 demonstrates an over-scaling of data corrected by eddy_correct, noticeable at $b = 700$ s/mm^2 and extremely clear at $b = 2000$ s/mm^2. The datasets corrected by eddy are much better aligned with the $b=0$ outlines. This figure also serves to highlight the difficulty in the application of qualitative methods to the assessment of artefact correction: the results are sensitive to exactly how the outline is drawn on the $b = 0$, which involves a subjective judgement of how much of the CSF to exclude as the b-value varies.

Evaluation of eddy

We investigated the dependence of eddy's performance on the quality of the dataset being corrected, both by varying the number of DW directions and by comparing the recommended full-shell acquisition

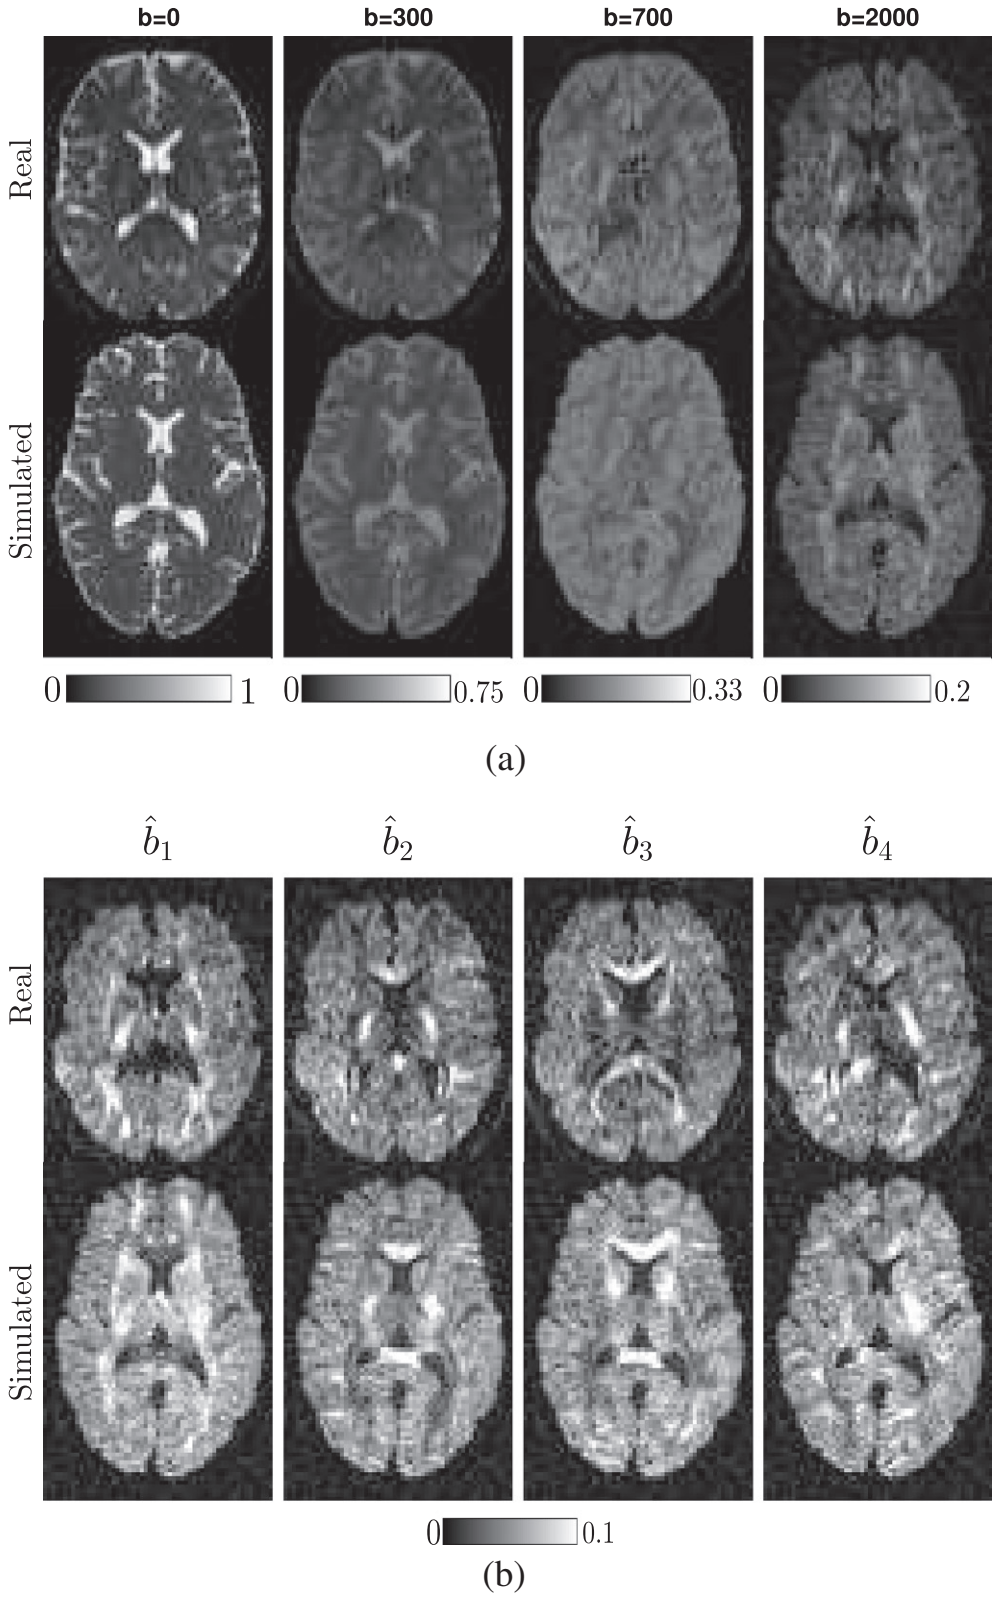

Fig. 2. Comparison of real and simulated data. a. Variation in contrast with respect to *b*-value. Both real and simulated datasets normalised against their respective $b = 0$ images. The direction of diffusion weighing is the same for both datasets. b. Variation in contrast with respect to direction of diffusion weighting. Each column represents an image acquired at $b = 2000$ s/mm^2 with a different b-vector, $\hat{b}_1 - \hat{b}_4$. Real and simulated datasets are normalised against their respective $b = 0$ images and shown on one intensity scale.

scheme to a half-shell scheme. Eddy makes use of information from 'similar' volumes in order to create a registration target for each volume in the dataset, and so, unlike eddy_correct, its ability to correct each volume is dependent on the full dataset. Thus it is recommended that datasets have a sufficient number of diffusion directions and are sampled on the full-sphere (or alternatively, with a blip-up blip-down

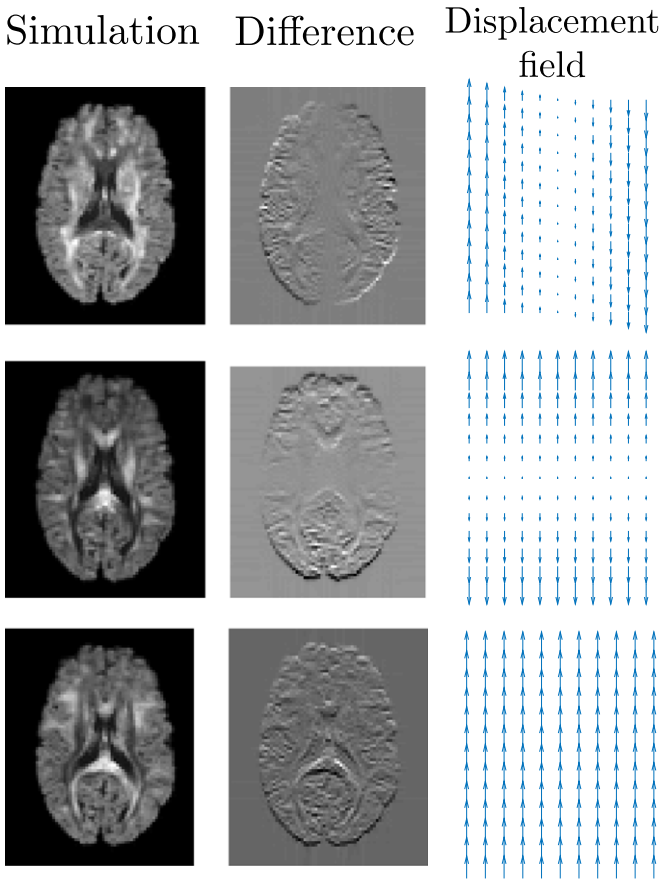

Fig. 3. Demonstration of DWIs simulated with EC artefacts. Left column are DWIs simulated at $b = 2000$ s/mm^2, with gradients pointing along the *x*-, *y*- and *z*-axes. Middle column shows difference images between these DWIs and their counterparts simulated without EC distortion. Right column shows displacement fields, representing the transformations from undistorted to distorted space. Displacement fields downsampled for clarity.

acquisition). However there is no information available on the minimum number of diffusion directions required for effective correction, or on the performance penalty incurred when correcting data acquired on the half-sphere.

To test eddy's dependence on the number of DW directions, a number of datasets were simulated to represent common diffusion tensor imaging (DTI) and high angular resolution (HARDI) acquisition protocols. Single-shell datasets with 16, 32, 48 and 64 diffusion directions were generated at both $b = 1000$ s/mm^2 and $b = 2000$ s/mm^2, with one $b = 0$ image for every 8 DWIs. Multi-shell datasets were made by combining the single-shell acquisitions, to create sets with 16/16, 16/32, 32/32, 32/64 and 64/64 directions in the $b = 1000$ s/mm^2 and $b = 2000$ s/mm^2 shells respectively. Diffusion directions were obtained from a minimisation of electrostatic energy as implemented in Camino (Cook et al., 2006; Jones et al., 1999), and optimised on the full-sphere. Each dataset had TR/TE = 7500/109 ms, $72 \times 86 \times 55$ with isotropic voxel size 2.5 mm, and Rician noise was added to produce datasets with both SNR = 10 and 20. EC and motion artefacts were added in the manner described in the Comparison of correction methods section. Each dataset was corrected using eddy with the same settings described in Comparison of correction methods. The results are reported in Fig. 7.

To test eddy's ability to cope with half-shell datasets, half-sphere sampling schemes were obtained from the full-sphere schemes mentioned previously, by negating each *b*-vector with a *z*-component less than 0. Datasets were generated from these schemes with the same acquisition parameters as their full shell counterparts. Rician noise was added to each dataset to create an SNR of 20. Each dataset was corrected three times using eddy: the first using the default settings as described in the Comparison of correction methods section, and the second and third times using linear and quadratic models that relate the parameters that define the EC distortion field to the applied *b*-vector (the default places no constraints on the relationship between these parameters and the *b*-vector). The results are shown in Fig. 8.

The results indicate that eddy is able to provide good correction down to 16 diffusion directions at SNR = 20, at both $b = 1000$ s/mm^2 and $b = 2000$ s/mm^2. At SNR = 10 good correction is achieved for the $b = 1000$ s/mm^2 datasets, but eddy struggles to correct the $b = 2000$ s/mm^2 single-shell datasets. The addition of more information might improve correction at low SNR-results for the combined 64/64 dataset are better than the 64 direction dataset at $b = 2000$ s/mm^2.

Results also indicate that, whilst full-shell sampling is optimal, it is still possible to obtain good correction on datasets acquired on the half-sphere. For most datasets the correction could be marginally improved by enforcing linear second-level modelling of the EC artefacts, which is to be expected as our simulations use a linear EC model.

Impact of correction on microstructure estimation

In this section we investigate the impact of artefact correction on the estimation of microstructural features from diffusion data. Firstly we fit two models, the diffusion tensor (DT) and NODDI, to simulated datasets before and after correction. The advantage of simulations is that we can compare the results to a 'ground truth' obtained by fitting these models to a dataset simulated without artefacts. Secondly we fit these models to a real dataset to demonstrate consistency with our findings on simulation. Finally, we examine the use of fitting residuals as a surrogate marker of improved correction.

The simulated dataset was the same as the one used in the Comparison of correction methods section: two shells with $b = 700$/

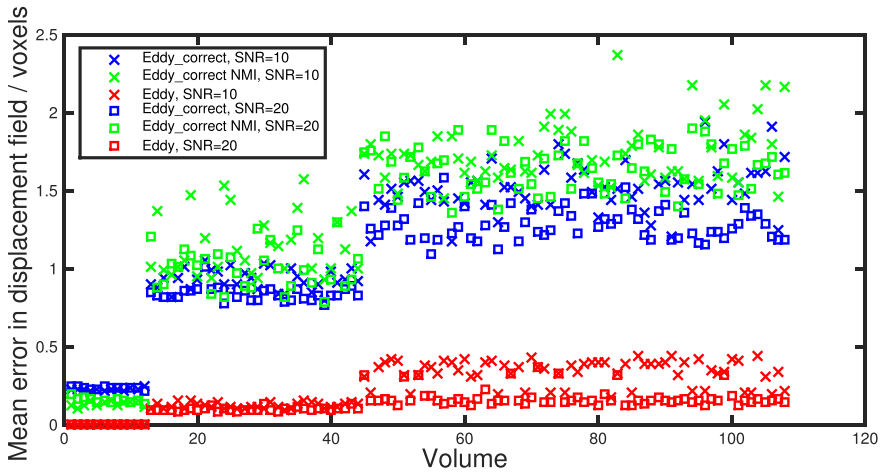

(a) Eddy currents

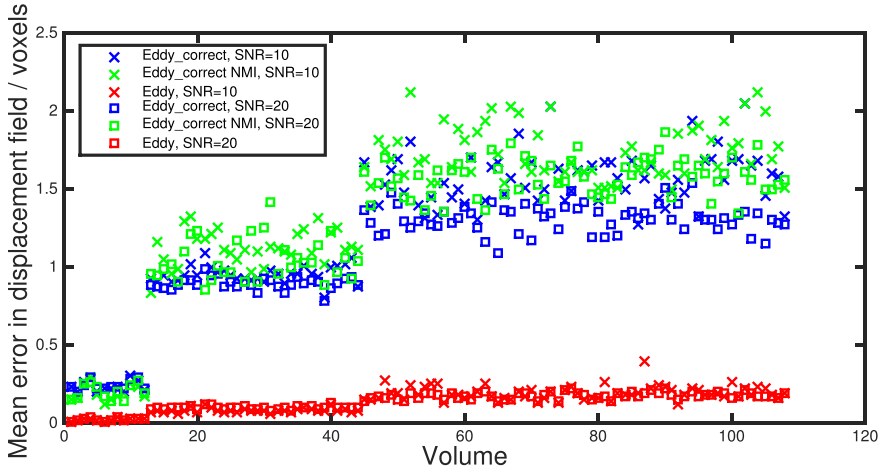

(b) Eddy currents and motion

Fig. 4. Mean error in displacement field across the brain. The first 8 volumes are $b = 0$, the next 32 are $b = 700$ s/mm^2 and the remaining 64 are $b = 2000$ s/mm^2.

2000 s/mm^2, 32/64 directions and isotropic voxel size 2.5 mm. We fit the DT to the $b = 700$ s/mm^2 shell of five datasets: the ground truth, both with and without noise, one distorted with motion and EC artefacts, and this distorted dataset corrected by eddy and eddy_correct. All datasets had SNR = 20. We changed the interpolation used by eddy_correct to spline, to match that used by eddy. FSL's DTIFIT was used to fit the tensor. We fit NODDI to both shells of each of these datasets, using the NODDI Matlab Toolbox. Fig. 9 shows the resulting FA and sum-squared error (SSE) residual maps, and Fig. 10 shows the NODDI parameter maps.

The real dataset was the same as used in the Comparison of correction methods section: two shells, $b = 700/2000$ s/mm^2, 32/64 directions with isotropic size 2.5 mm. We fit the DT to the $b = 700$ s/mm^2 shell of the original data, and the data after correction by eddy and eddy_correct. We fit NODDI to the multi-shell dataset. Fig. 11 shows the resulting parameter and residual maps.

The results from simulation in Figs. 9 and 10 demonstrate the impact of these post-processing techniques on estimating microstructure. Whilst FA maps in Fig. 9 are hard to distinguish, the difference maps in Fig. 9b are more informative. They show anatomical structure in the data corrected by eddy_correct which is not apparent in the data corrected by eddy. The NODDI maps in Fig. 10a show a smoothing of the parameters when data is corrected with eddy_correct, particularly noticeable in the orientation dispersion index (ODI). There is overestimation of the intra-cellular volume fraction (V_{IC}) in the dataset corrected by eddy_correct, particularly in GM regions, and an underestimation of the isotropic volume fraction (V_{ISO}) around the edges of the brain, likely caused by the overscaling of the DWIs. The difference maps in Fig. 10b make these problems more clear. There is noticeable structure in the difference maps for eddy_correct. By contrast, the parameter maps estimated from data corrected by eddy are much more similar to the ground truth maps.

Whilst there is no GT available for the fits to real data, the parameter maps show features consistent with those found on simulation. Compared to the data corrected by eddy, the dataset corrected by eddy_correct shows a loss of sharpness in some of the FA structures in Fig. 11a. There is also an overestimation of the V_{IC} and underestimation of V_{ISO} relative to the results from eddy, indicated on the figure. The general smoothing of the ODI map is also noticeable.

The results allow us to investigate the effectiveness of using fitting residuals as a surrogate marker of image alignment. Both corrected datasets show lower residuals than the original data. Residuals from the two correction methods appear similar, and are best compared using difference images, shown in Fig. 12. Despite results from the Comparison of correction methods section showing that eddy provides better correction, eddy_correct gives lower residuals from DT fits in a rim around the brain, in both real and simulated data. There are also some regions in the middle of the brain where eddy_correct shows lower residuals from DT fits. We speculate this is caused by the smoothing that results from the over-scaling of the DWIs, which makes the signal easier to fit to. It seems that the smoothing from interpolation also has an effect: in Fig. 10a the

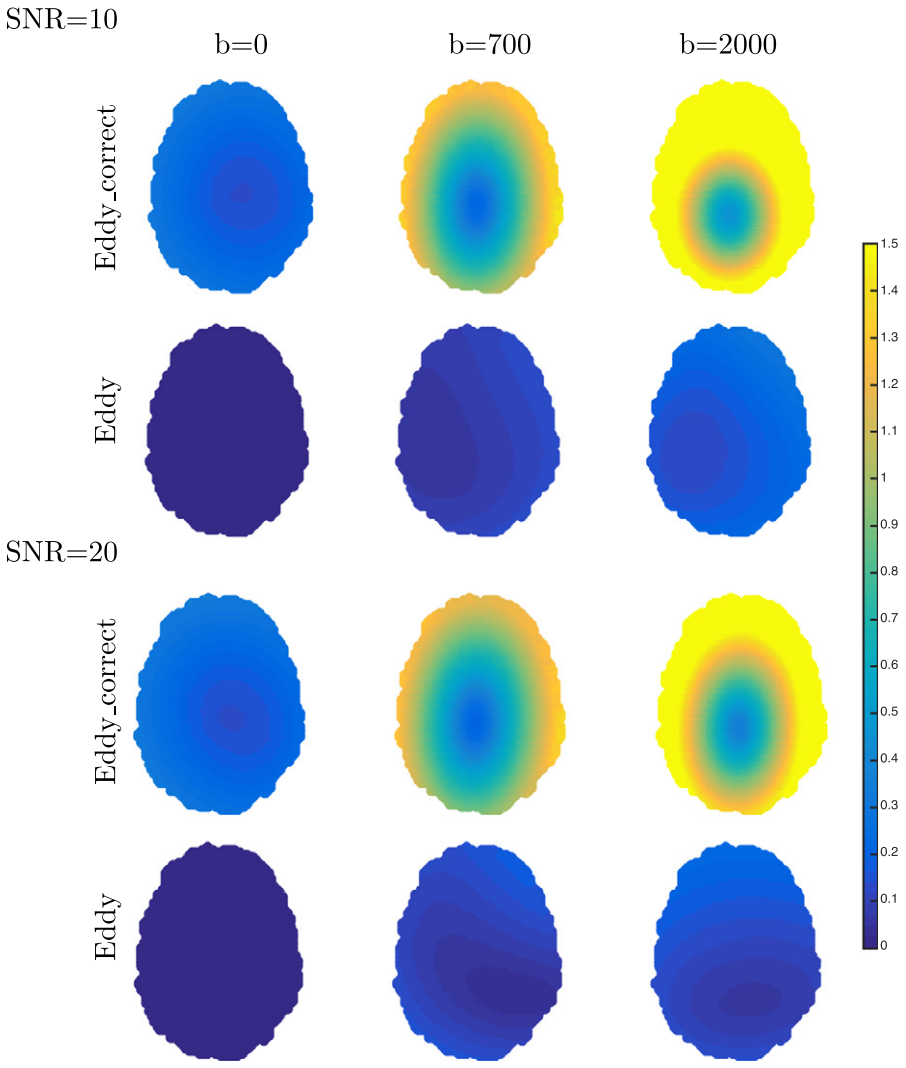

Fig. 5. Error in displacement field, in voxels (each 2.5 mm isotropic), shown across one slice of the brain. The errors are a mean across all of the volumes with the same *b*-value. Results shown for eddy_correct were obtained using the default cost function.

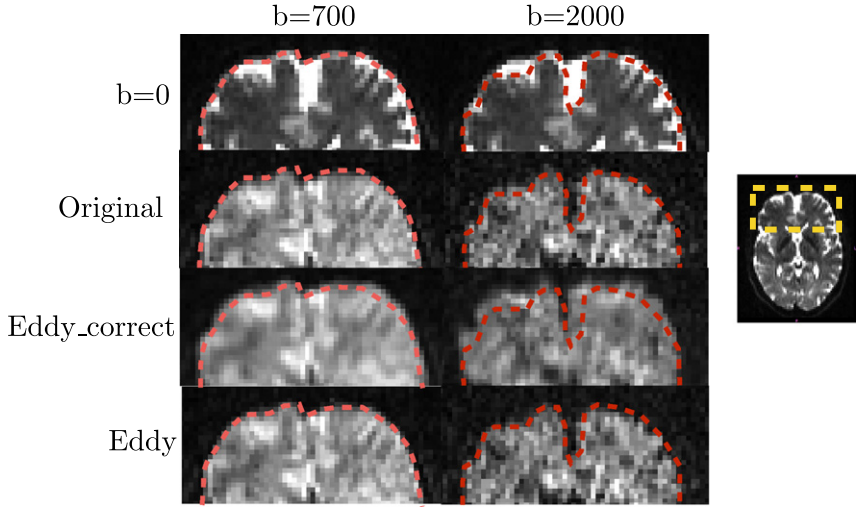

Fig. 6. Correction errors on real data. Anterior portion of the brain in an axial slice is shown, corresponding to the yellow region on the inset image. An outline was drawn around the undistorted $b = 0$ image, then superposed on a DWI. The boundaries of an undistorted DWI should align with this outline. Different outlines were drawn on the $b = 0$ for use on the $b = 700$ and $b = 2000$ volumes, to account for the different amounts of CSF attenuation present.

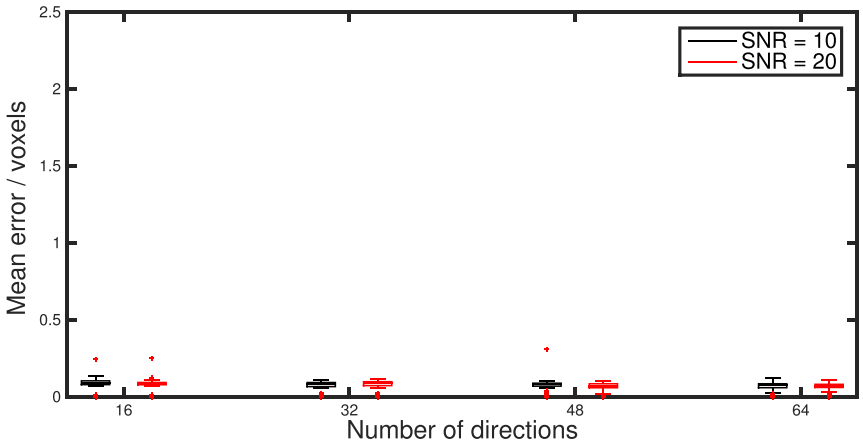

(a) Single shell, b=1000 s/mm^2

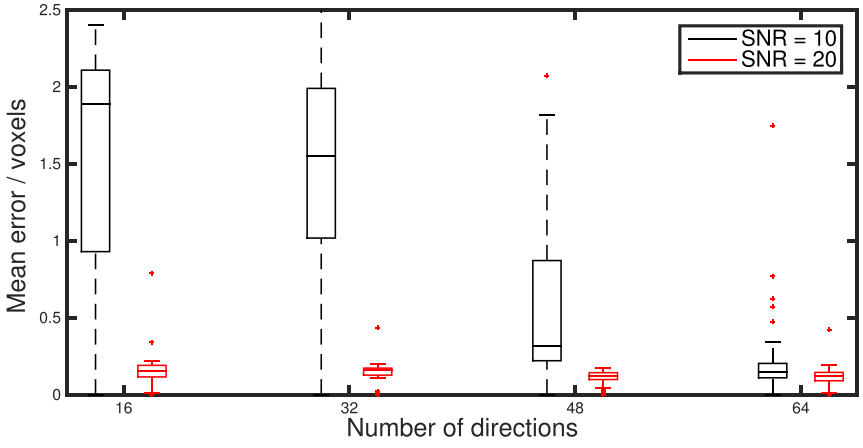

(b) Single shell, b=2000 s/mm^2

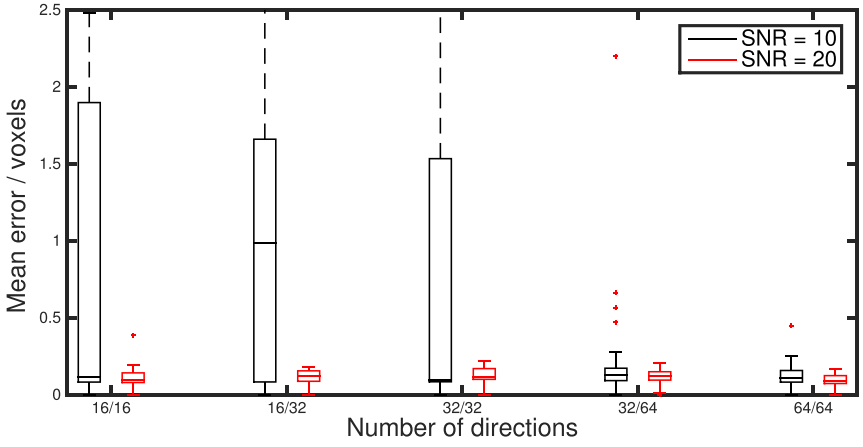

(c) Multi-shell, b=1000 s/mm^2 and 2000 s/mm^2

Fig. 7. Errors in the displacement fields in datasets corrected with eddy, as the number of directions in the acquisition is varied. Each data point is a mean over the voxels in the brain for a volume, boxplots show the distribution of these means across the dataset.

residuals from the corrected datasets are lower than those in the GT + noise data.

Fig. 12 shows that residuals from NODDI fits are lower for data corrected with eddy than data corrected with eddy_correct. The increased residuals for data corrected by eddy_correct seems to be due to the multi-shell nature of the data. Results in the Comparison of correction methods section show that eddy_correct overscales DWIs by increasing amounts with increasing b-values, so this internal misalignment in the dataset is likely the reason that NODDI fits data corrected using eddy_correct badly. These results indicate that reduced residuals from model fits can be indicative of increased image alignment but can also be confounded by other factors, such as image smoothing and expansion into background regions, that mean these measures need to be interpreted carefully.

Discussion

We have presented a framework that allows for the simulation of full-brain DW-MR datasets with artefacts, along with their corresponding ground truth displacement fields, providing an objective and

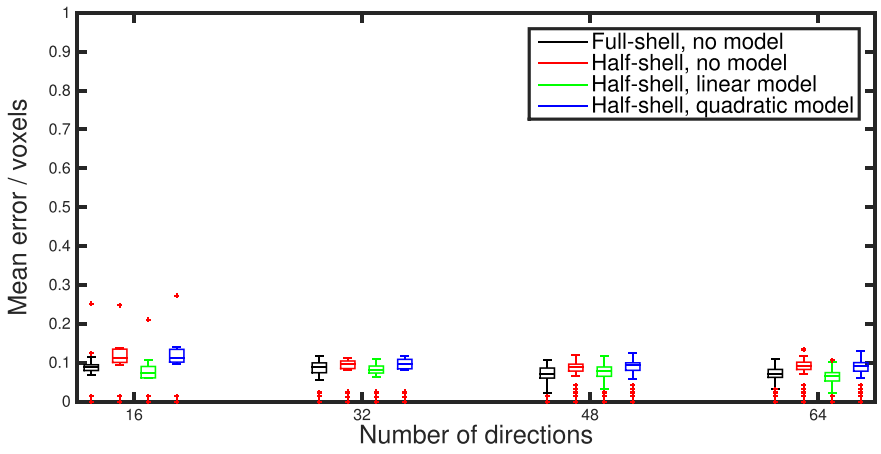

(a) Single shell, b=1000 s/mm^2

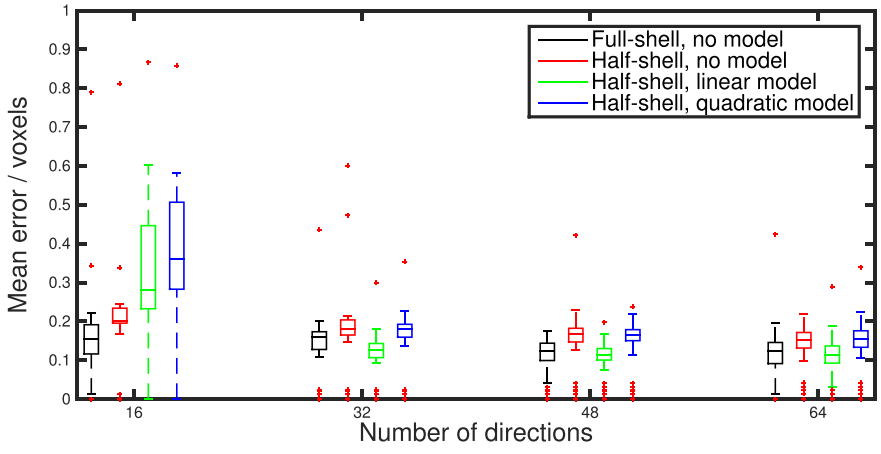

(b) Single shell, b=2000 s/mm^2

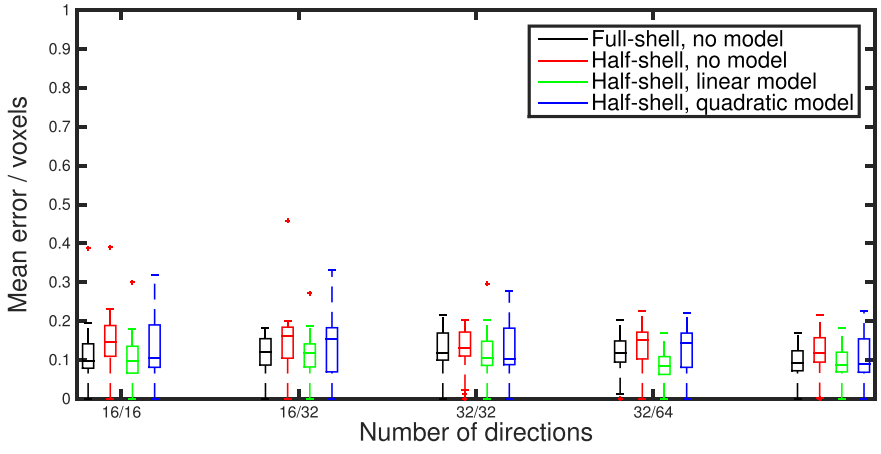

(c) Multi-shell, b=1000 s/mm^2 and 2000 s/mm^2

Fig. 8. Errors in the displacement fields in datasets corrected with eddy for full-shell and half-shell acquisitions. Each data point is a mean over the voxels in the brain for a volume, boxplots show the distribution of these means across the dataset. The model refers to the relationship between the applied b −vector and the parameters that determine the EC field. All datasets have SNR = 20.

quantitative means of assessing post-processing techniques. Our framework combines two elements necessary for the simulation of realistic full-brain DWIs. Firstly, our simulated images are able to provide a realistic representation of the contrast differences found across DW-MR datasets. We achieve this using a model-free approach that obtains the signal from real data. Previous simulations have used a single representative mean diffusivity or diffusion tensor for each tissue type to provide diffusion-weighting (Bastin, 1998; Nunes et al., 2011), which leads to vastly oversimplified contrast, or have used underlying WM structures to generate the signal (Neher et al., 2013; Close et al., 2009; Caruyer et al., 2014), preventing realistic contrast from being achieved in non-WM regions of the brain. The second element is the modelling of the MR acquisition process. Without simulating the full image generation process certain artefacts cannot be introduced, such as blurring due to the decay of EC, certain motion artefacts (blurring, signal drop-out and distortion), and the effect of the Fourier transform on a complex signal.

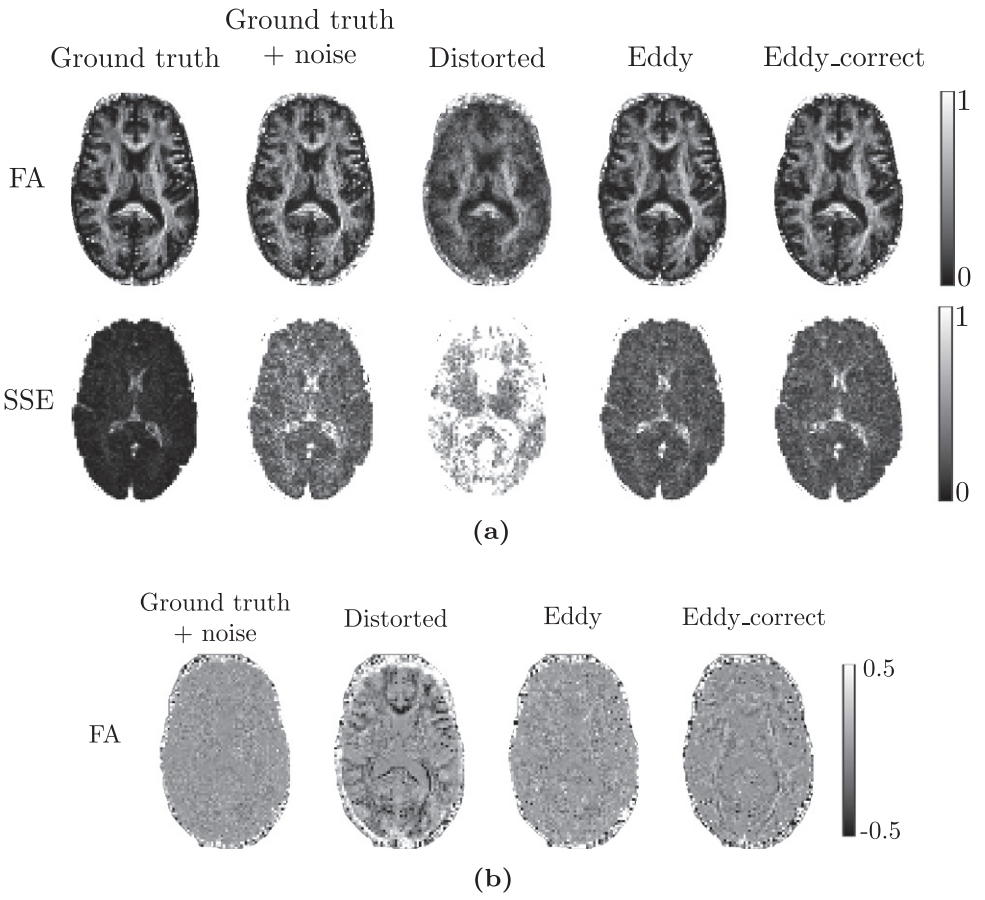

Fig. 9. FA maps resulting from fits to ground truth data (i.e., data simulated with no distortions), ground truth data with added noise, data distorted with both EC + motion, and distorted data corrected by eddy and eddy_correct. a. FA and signal residual maps. Residuals are the sum-squared difference between actual and predicted signal. b. Difference in FA compared to ground truth estimates. SNR = 20 for all noisy data.

The simulation framework allowed us to quantitatively compare a widely used technique for correcting EC and motion artefacts in DWI, eddy_correct, and its recently proposed alternative, eddy. We were able to provide quantitative evidence that eddy_correct systematically overscales DWIs, which corroborates with findings in the literature (Maniega et al., 2007). Other similar techniques may also exhibit this tendency to overscale, and the framework could be readily applied to test this. We also demonstrated that eddy is able to provide significantly better correction. These findings corroborate with a previous study (Yamada et al., 2014), which applied the two methods to real datasets and assessed them using a combination of visual inspection and comparison of FA values.

We were also able to investigate the quality of dataset needed to ensure good correction with eddy. For data with SNR = 20, good correction can be obtained on as few as 16 directions, though denser sampling is needed for data with lower SNR, such as $b = 2000$ s/mm^2 data at SNR = 10. We also demonstrated that, whilst a full-sphere acquisition scheme is ideal, it is still possible to achieve good correction on half-sphere datasets. These findings are useful for retrospective studies where full-shell acquisition has not been considered. The datasets simulated contained particularly severe artefacts, including large, random movements between each volume and large EC distortions, so these findings could be considered to be an upper-bound on the error that can be expected when correcting DW datasets. However, our findings indicate that the performance of eddy is robust to the severity of artefact, as supported by the similar corrections achieved for a dataset with just EC artefacts and a dataset with both EC and severe motion in Fig. 4. It may be that the severity of these artefacts explains why we found more than the recommended 5 iterations were necessary to achieve good correction with eddy.

These results are important in the context of techniques that use DW-MR data. Data is most commonly acquired at $b = 1000$ s/mm^2, and our results indicate we can expect errors of more than 1 voxel in such images if they are corrected using registration to $b = 0$. These are enough to cause anatomical misalignment in regions of partial volume, such as the boundaries between GM and CSF which will compromise any information on microstructure obtained from such data, as demonstrated by the FA maps in Fig. 9. Fig. 10 shows this effect is even more severe for data acquired at $b = 2000$ s/mm^2, which is becoming more common with the increasing popularity of HARDI techniques. In this work we concentrated on the assessment of EC and instantaneous motion artefacts, but we are able to reproduce a number of artefacts with our framework, including Nyquist ghosting, chemical shifts, motion during signal read-out, Gibbs ringing and signal dropout. The presence of such artefacts will confound attempts to correct for EC + motion, and further work could investigate the extent to which they do.

We use comparison of displacement fields in order to assess the effectiveness of correction techniques. This directly measures the desired outcome of such techniques, i.e., the mapping of all images into an undistorted space. We contrast this with some of the quantitative surrogate metrics that have been used to assess techniques, such as FA values in WM tracts (Yamada et al., 2014), the length of tracts obtained from tractography (Zhuang et al., 2013) or the size of the residuals from a model-based fit to the data (Ben-Amitay et al., 2012; Andersson and Skare, 2002). These methods are often the best available when testing correction on real datasets where a GT is not available, but they are indirect metrics: increased image alignment is only one of many factors that can affect them. For example, the results indicate that residuals from a DT fit can be lowered by smoothing of the data due to overscaling of the DWIs. Current methods only attempt to correct for

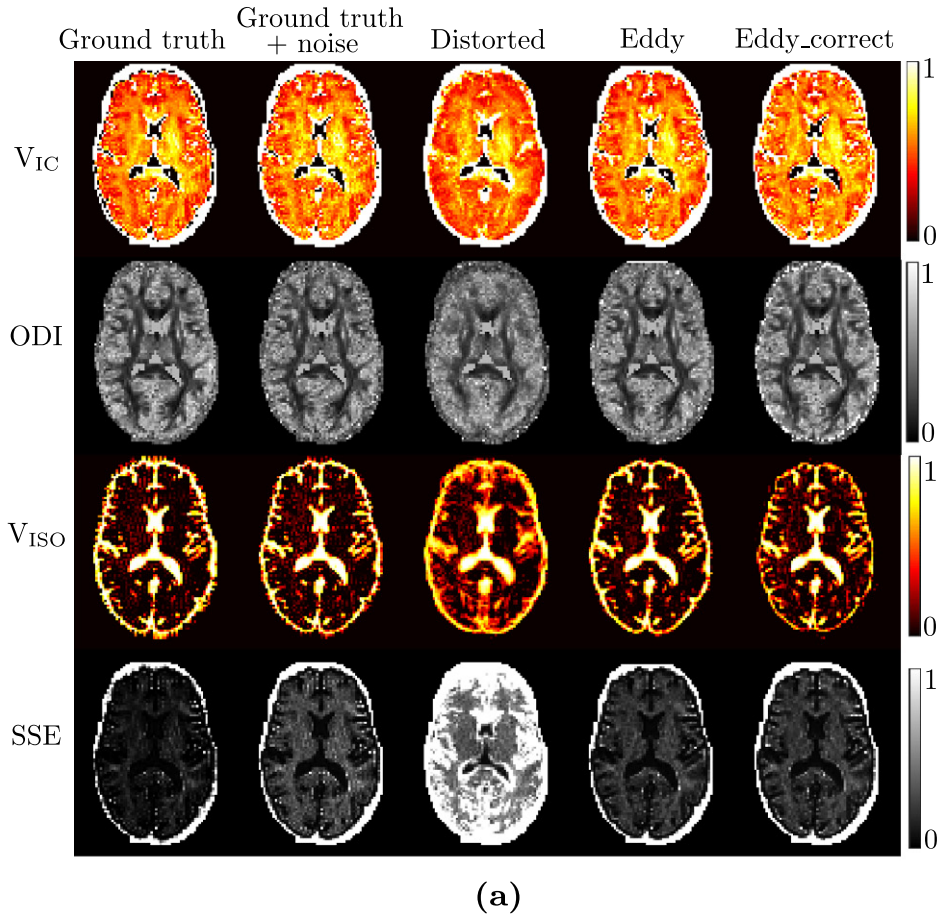

Fig. 10. NODDI parameters resulting from fits to ground truth data (i.e., data simulated with no distortions), ground truth data with added noise, data distorted with both EC + motion, and distorted data corrected by eddy and eddy_correct. a. NODDI parameter and signal residual maps. Residuals are the sum-squared difference between actual and predicted signal. b. Difference in NODDI parameters compared to ground truth estimates. SNR = 20 for all noisy data. Parameters are: V_{IC} — intracellular volume fraction, ODI — orientation dispersion index, V_{ISO} — isotropic volume fraction.

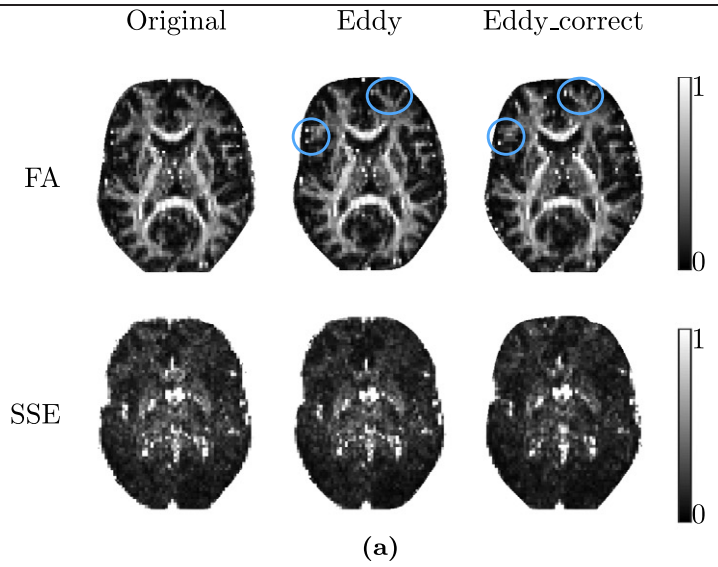

(a)

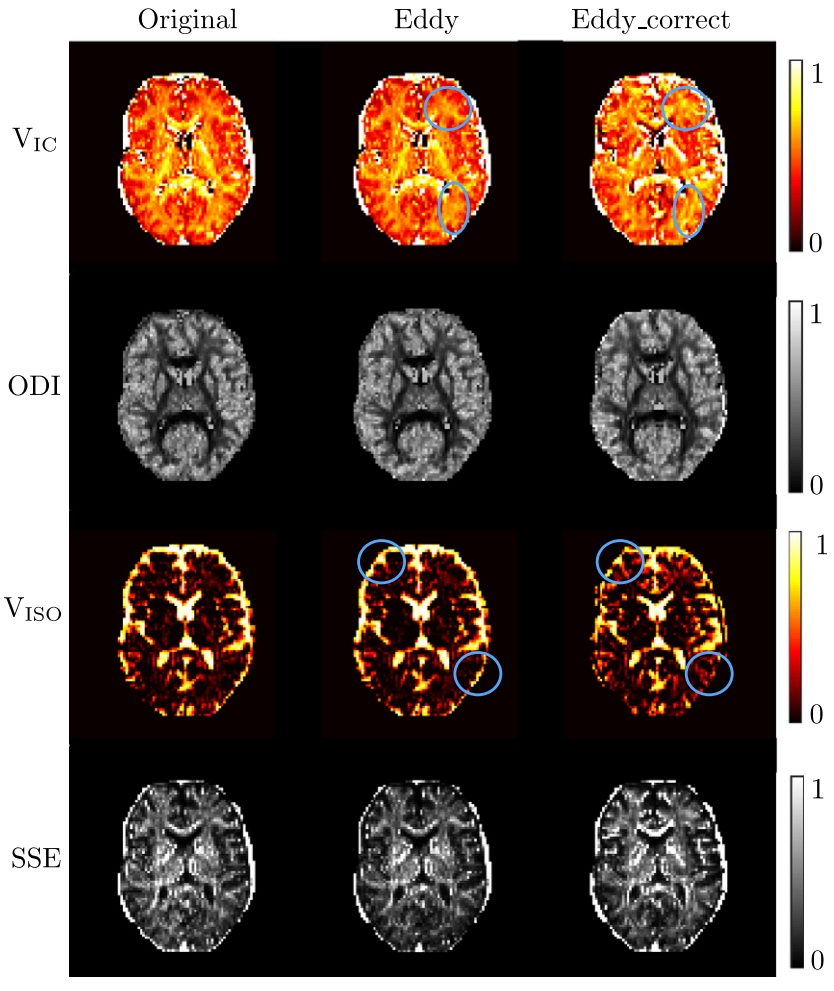

(b)

Fig. 11. Parameter maps resulting from fits to real data. a. FA and SSE residual maps. Blue regions highlight blurring of FA structures by eddy_correct compared to eddy. b. NODDI parameter and SSE residual maps. Blue regions highlight increase in V_{IC} and decrease in V_{ISO} in data corrected by eddy_correct when compared to eddy. Also note eddy_correct causes smoothing of the ODI map.

geometric distortions, so we predicted geometric displacement fields for assessment. Future methods may also try to correct for non-geometric effects such as the blurring due to decay of ECs, and our framework can be extended to provide GT estimates of how these may be corrected, which could assist both the development and testing of such methods.

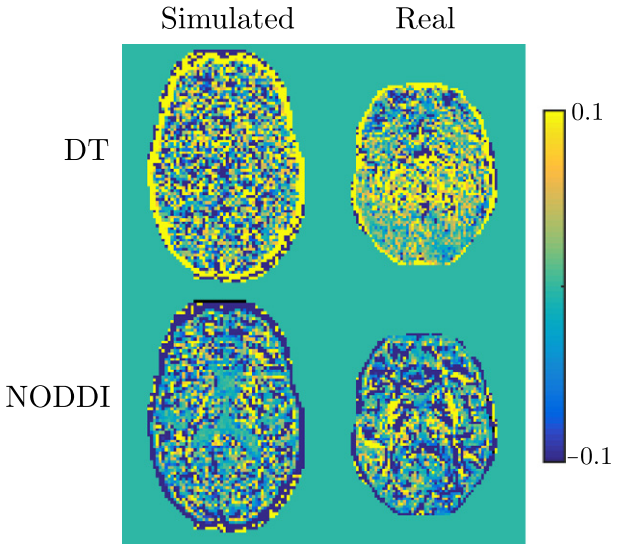

Fig. 12. Difference in residuals from eddy and eddy_correct: SSE_{eddy}–$SSE_{eddy\_correct}$, meaning positive values indicate voxels where data corrected using eddy_correct had lower residuals.

There are some limitations to the framework presented. Typical processing pipelines include correction of susceptibility artefacts in addition to EC and motion, so B_0 inhomogeneities need to be included to enable testing of more complete pipelines. Currently these cannot be accurately simulated, due to our use of a GE-EPI sequence with T_2^* values replaced by T_2 values to simulate SE-EPI. The lack of a refocusing pulse means that any B_0 inhomogeneities will lead to overly severe distortions and signal dropout. Implementation of true SE-EPI will be the subject of future work. Additionally, the linear EC model used, which causes simple scalings, shearings and translations of each DWI, has been found to be inadequate for describing the EC fields in some scanners. Rather they are better described by a second or third order polynomial (Rohde et al., 2004). Furthermore we used a single value of T1 and T2 for each tissue type, whilst they have been shown to vary spatially across the brain (Wansapura et al., 1999). Failure to model this will account for the lack of subtlety in the contrast compared to real data, noticeable in e.g., the comparison at $b = 700$ s/mm^2 in Fig. 2a. Further work could incorporate spatial maps that will allow for variations within tissue types. Finally we are currently unable to simulate parallel imaging, and any of the artefacts it can give rise to, such as spatially varying SNR across the brain. Future work implementing parallel imaging will enable us to assess how much these additional considerations impact the results found in this work.

Using a real dataset to determine the signal attenuation allows us to achieve more realistic contrast throughout the brain than is currently possible with model-based techniques, but this approach has two limitations. The spherical harmonic approach means we can only predict attenuation for b-values we already have data for — in our HCP dataset this is $b = 1000$, 2000 and 3000 s/mm^2. Whilst we were able to predict attenuation at $b = 700$ s/mm^2 using data from the $b = 1000$ s/mm^2 shell and a mono-exponential assumption, this is only appropriate at low b-values and could not be used to predict attenuation at higher b-values. Employing an approach such as MAP-MRI (Özarslan et al., 2013) could allow us to simulate datasets over a wider range of b-values. Furthermore, any artefacts such as ghosting, ringing and noise that are present in the data will be projected into our simulations. We tried to minimise this effect by using data from the HCP, which is acquired using bespoke scanners and sequences, then carefully processed to produce high quality data. We also visually inspected the dataset to check for the presence of artefacts. We believe that artefacts that were not detected by these checks will have minimal impact on the experiments performed. An interesting avenue for future work might be to build upon the model-based techniques that generate their signal from a description of underlying microstructure, such as Fiberfox (Neher et al., 2013), in order to generate realistic full-brain contrast. Whilst these techniques do not currently provide realistic contrast in non-WM regions (see Supplementary Fig. 1), this is not an inherent limitation of such methods: better models and descriptions of underlying microstructure in non-WM regions could offer improvements. This would have the advantage of providing a GT for model fitting in addition to the current deformation field GT that our method provides.

Using the simulation framework, we can quantitatively assess the effectiveness of artefact correction schemes. In demonstrating the framework's application we have shown that one of the most commonly used correction techniques, eddy_correct, introduces a systematic error that is significant enough to undermine any analysis performed on data corrected using this scheme. We were further able to demonstrate that eddy provides a much better alternative, and give an idea of the level of correction that it can be expected to achieve for commonly acquired datasets. Finally, we were able to demonstrate that poor correction can adversely impact model fitting to the data. We hope that this framework will become a key aspect of the validation of any post-processing schemes, which will allow users to make decisions on their choice of processing techniques that are informed by objective, quantitative evidence.

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.neuroimage.2015.11.006.

Acknowledgments

We would like to thank Jesper Andersson for providing us with a version of eddy that outputs displacement fields, and for helpful discussions. MG is supported by the EPSRC (EP/L504889/1) and the EPSRC Centre for Doctoral Training (EP/L016478/1). HZ is supported by the EPSRC (EP/L022680/1), the MRC (MR/L011530/1) and the Royal Academy of Engineering Research Exchanges with China and India. ID is supported by the Leverhulme Trust. MG and HZ are additionally supported by the Royal Society International Exchange Scheme with China.

Data were provided [in part] by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

References

Alexander, D.C., Barker, G.J., Arridge, S.R., 2002. Detection and modeling of non-Gaussian apparent diffusion coefficient profiles in human brain data. Magn. Reson. Med. 48, 331–340. http://dx.doi.org/10.1002/mrm.10209.

Andersson, J.L.R., Skare, S., 2002. A model-based method for retrospective correction of geometric distortions in diffusion-weighted EPI. NeuroImage 16, 177–199. http://dx.doi.org/10.1006/nimg.2001.1039.

Andersson, J.L., Sotiropoulos, S.N., 2015. An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging. NeuroImage http://dx.doi.org/10.1016/j.neuroimage.2015.10.019.

Andersson, J.L.R., Skare, S., Ashburner, J., 2003. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. NeuroImage 20, 870–888. http://dx.doi.org/10.1016/S1053-8119(03)00336-7.

Assaf, Y., Basser, P.J., 2005. Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. NeuroImage 27, 48–58. http://dx.doi.org/10.1016/j.neuroimage.2005.03.042.

Assaf, Y., Cohen, Y., 2009. Inferring microstructural information of white matter from diffusion MRI. Diffusion MRI: From Quantitative Measurement to In-vivo Neuroanatomy. 127.

Bai, Y., Alexander, D., 2008. Model-based registration to correct for motion between acquisitions in diffusion MR imaging. 5th IEEE Int Symp Biom, Imaging, pp. 947–950.

Bastin, M.E., 1998. Correction of eddy current induced artefacts in MR diffusion iterative cross-correlation. Magn. Reson. Imaging 17, 1998-1998.

Bastin, M.E., 2001. On the use of the FLAIR technique to improve the correction of eddy current induced artefacts in MR diffusion tensor imaging. Magn. Reson. Imaging 19, 937–950. http://dx.doi.org/10.1016/S0730-725X(01)00427-1.

Ben-Amitay, S., Jones, D.K., Assaf, Y., 2012. Motion correction and registration of high b-value diffusion weighted images. Magn. Reson. Med. 67, 1694–1702. http://dx.doi.org/10.1002/mrm.23186.

Bodini, B., Ciccarelli, O., 2009. Diffusion MRI in neurological disorders. In: Johansen-Berg, H., Johansen-Berg, H., Behrens, T.E.J. (Eds.), Diffusion MRI: From Quantitative Measurement to In-vivo, Neuroanatomy, pp. 175–204.

Caruyer, E., Daducci, A., Descoteaux, M., Houde, J.c., Thiran, J.p., Verma, R., 2014. Phantomas: A Flexible Software Library to Simulate Diffusion MR Phantoms. 17 p. 20013.

Close, T.G., Tournier, J.D., Calamante, F., Johnston, L.a., Mareels, I., Connelly, A., 2009. A software tool to generate simulated white matter structures for the assessment of fibre-tracking algorithms. NeuroImage 47, 1288–1300. http://dx.doi.org/10.1016/j.neuroimage.2009.03.077.

Cook, P.a., Bai, Y., Seunarine, K.K., Hall, M.G., Parker, G.J., Alexander, D.C., 2006. Camino: open-source diffusion-MRI reconstruction and processing. 14th Scientific Meeting of the International Society for Magnetic Resonance in Medicine 14, p. 2759.

Drobnjak, I., Gavaghan, D., Süli, E., Pitt-Francis, J., Jenkinson, M., 2006. Development of a functional magnetic resonance imaging simulator for modeling realistic rigid-body motion artifacts. Magn. Reson. Med. 56, 364–380. http://dx.doi.org/10.1002/mrm.20939.

Drobnjak, I., Pell, G.S., Jenkinson, M., 2010. Simulating the effects of time-varying magnetic fields with a realistic simulated scanner. Magn. Reson. Imaging 28, 1014–1021. http://dx.doi.org/10.1016/j.mri.2010.03.029.

Graham, M.S., Drobnjak, I., Zhang, H., 2015. A simulation framework for quantitative validation of artefact correction in diffusion MRI. Information Processing in Medical, Imaging, pp. 638–649.

Haselgrove, J.C., Moore, J.R., 1996. Correction for distortion of echo-planar images used to calculate the apparent diffusion coefficient. Magn. Reson. Med. 36, 960–964. http://dx.doi.org/10.1002/mrm.1910360620.

Horsfield, M.a., 1999. Mapping eddy current induced fields for the correction of diffusion-weighted echo planar images. Magn. Reson. Imaging 17, 1335–1345. http://dx.doi.org/10.1016/S0730-725X(99)00077-6.

Irfanoglu, M.O., Walker, L., Sarlls, J., Marenco, S., Pierpaoli, C., 2012. Effects of image distortions originating from susceptibility variations and concomitant fields on diffusion MRI tractography results. NeuroImage 61, 275–288. http://dx.doi.org/10.1016/j.neuroimage.2012.02.054.

Jenkinson, M., Smith, S., 2001. A global optimisation method for robust affine registration of brain images. Med. Image Anal. 5, 143–156.

Jezzard, P., Balaban, R.S., 1995. Correction for geometric distortion in echo planar images from B0 field variations. Magn. Reson. Med. 34, 65–73. http://dx.doi.org/10.1002/mrm.1910340111.

Jezzard, P., Barnett, A.S., Pierpaoli, C., 1998. Characterization of and correction for eddy current artifacts in echo planar diffusion imaging. Magn. Reson. Med. 39, 801–812. http://dx.doi.org/10.1002/mrm.1910390518.

Jones, D.K., Horsfield, M.A., Simmons, A., 1999. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. Magn. Reson. Med. 42, 515–525. http://dx.doi.org/10.1002/(SICI)1522-2594(199909)42:3<515::AID-MRM14>3.0.CO;2-Q.

Kwan, R.S., Evans, A.C., Pike, G.B., 1999. MRI simulation-based evaluation of image-processing and classification methods. IEEE Trans. Med. Imaging 18, 1085–1097.

Le Bihan, D., Poupon, C., Amadon, A., Lethimonnier, F., 2006. Artifacts and pitfalls in diffusion MRI. J. Magn. Reson. Imaging 24, 478–488.

Mangin, J.F., Poupon, C., Clark, C., Le Bihan, D., Bloch, I., 2002. Distortion correction and robust tensor estimation for MR diffusion imaging. Med. Image Anal. 6, 191–198.

Maniega, S.M., Bastin, M.E., Armitage, P.A., 2007. A quantitative comparison of two methods to correct eddy current-induced distortions in DT-MRI. Magn. Reson. Imaging 25, 341–349. http://dx.doi.org/10.1016/j.mri.2006.09.009.

Neher, P.F., Laun, F.B., Stieltjes, B., Maier-Hein, K.H., 2013. Fiberfox: facilitating the creation of realistic white matter software phantoms. Magn. Reson. Med. 1470, 1460–1470. http://dx.doi.org/10.1002/mrm.25045.

Niendorf, T., Dijkhuizen, R.M., Norris, D.G., van Lookeren Campagne, M., Nicolay, K., 1996. Biexponential diffusion attenuation in various states of brain tissue: implications for diffusion-weighted imaging. Magn. Reson. Med. 36, 847–857. http://dx.doi.org/10.1002/mrm.1910360607.

Nunes, R.G., Drobnjak, I., Clare, S., Jezzard, P., Jenkinson, M., 2011. Performance of single spin-echo and doubly refocused diffusion-weighted sequences in the presence of eddy current fields with multiple components. Magn. Reson. Imaging 29, 659–667. http://dx.doi.org/10.1016/j.mri.2011.02.015.

Oguz, I., Farzinfar, M., Matsui, J., Budin, F., Liu, Z., Gerig, G., Johnson, H.J., Styner, M., 2014. DTIPrep: quality control of diffusion-weighted images. Front. neuroinformatics 8.

Özarslan, E., Koay, C.G., Shepherd, T.M., Komlosh, M.E., Irfanoglu, M.O., Pierpaoli, C., Basser, P.J., 2013. Mean apparent propagator (MAP) MRI: a novel diffusion imaging method for mapping tissue microstructure. NeuroImage 78, 16–32. http://dx.doi.org/10.1016/j.neuroimage.2013.04.016.

Reese, T.G., Heid, O., Weisskoff, R.M., Wedeen, V.J., 2003. Reduction of eddy-current-induced distortion in diffusion MRI using a twice-refocused spin echo. Magn. Reson. Med. 49, 177–182. http://dx.doi.org/10.1002/mrm.10308.

Rohde, G.K., Barnett, a.S., Basser, P.J., Marenco, S., Pierpaoli, C., 2004. Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI. Magn. Reson. Med. 51, 103–114. http://dx.doi.org/10.1002/mrm.10677.

Scholz, J., Tomassini, V., Johansen-Berg, H., 2009. Individual differences in white matter microstructure in the healthy brain. Diffusion MRI. Academic Press, London, pp. 237–250.

Tournier, J., Calamante, F., Gadian, D.G., Connelly, A., 2004. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. NeuroImage 23, 1176–1185.

Tournier, J.D., Calamante, F., Connelly, A., 2013. Determination of the appropriate b value and number of gradient directions for high-angular-resolution diffusion-weighted imaging. NMR Biomed. 26, 1775–1786. http://dx.doi.org/10.1002/nbm.3017.

Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S.W., et al., 2012. The human connectome project: a data acquisition perspective. NeuroImage 62, 2222–2231.

Wansapura, J.P., Holland, S.K., Dunn, R.S., Ball, W.S., 1999. NMR relaxation times in the human brain at 3.0 Tesla. J. Magn. Reson. Imaging 9, 531–538. http://dx.doi.org/10.1002/(SICI)1522-2586(199904)9:4<531::AID-JMRI4>3.0.CO;2-L.

Yamada, H., Abe, O., Shizukuishi, T., Kikuta, J., Shinozaki, T., Dezawa, K., Nagano, A., Matsuda, M., Haradome, H., Imamura, Y., 2014. Efficacy of distortion correction on diffusion imaging: comparison of FSL eddy and eddy_correct using 30 and 60 directions diffusion encoding. PLoS One 9, e112411. http://dx.doi.org/10.1371/journal.pone.0112411.

Zaitsev, M., Hennig, J., Speck, O., 2004. Point spread function mapping with parallel imaging techniques and high acceleration factors: fast, robust, and flexible method for echo-planar imaging distortion correction. Magn. Reson. Med. 52, 1156–1166. http://dx.doi.org/10.1002/mrm.20261.

Zhang, Y., Brady, M., Smith, S., 2001. Segmentation of brain MR images through a hidden markov random field model and the expectation–maximization algorithm. IEEE Trans. Med. Imaging 20, 45–57.

Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C., 2012. NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. NeuroImage 61, 1000–1016. http://dx.doi.org/10.1016/j.neuroimage.2012.03.072.

Zhuang, J., Lu, Z.L., Vidal, C.B., Damasio, H., 2013. Correction of eddy current distortions in high angular resolution diffusion imaging. J. Magn. Reson. Imaging 37, 1460–1467. http://dx.doi.org/10.1002/jmri.23929.