Published in final edited form as:

IEEE Trans Med Imaging. 2014 February ; 33(2): 272–289. doi:10.1109/TMI.2013.2284014.

A Unified Approach to Diffusion Direction Sensitive Slice Registration and 3-D DTI Reconstruction From Moving Fetal Brain Anatomy

Mads Fogtmann,

Biomedical Image Computing Group, Departments of Pediatrics, Bioengineering and Radiology, University of Washington, Seattle, WA 98195 USA, and also with the DTU Informatics, Technical University of Denmark, 2800 Kgs-Lyngby, Denmark

Sharmishta Seshamani,

Biomedical Image Computing Group, Departments of Pediatrics, Bioengineering and Radiology, University of Washington, Seattle, WA 98195 USA

Christopher Kroenke,

Oregon Health and Science University, Advanced Imaging Research Center and Department of Behavioral Neuroscience, University of Washington, Portland, OR 97239 USA

Xi Cheng,

Biomedical Image Computing Group, Departments of Pediatrics, Bioengineering and Radiology, University of Washington, Seattle, WA 98195 USA

Teresa Chapman,

Department of Radiology, Seattle Children's Hospital, Seattle, WA 98105 USA

Jakob Wilm,

DTU Informatics, Technical University of Denmark, 2800 Kgs-Lyngby, Denmark

François Rousseau, and

ICube, Universit de Strasbourg, CNRS, 67412 Illkirch, France

Colin Studholme [Senior Member IEEE]

Biomedical Image Computing Group, Departments of Pediatrics, Bioengineering and Radiology, University of Washington, Seattle, WA 98195 USA

Colin Studholme: colin.studholme@ieee.org

Abstract

This paper presents an approach to 3-D diffusion tensor image (DTI) reconstruction from multi-slice diffusion weighted (DW) magnetic resonance imaging acquisitions of the moving fetal brain. Motion scatters the slice measurements in the spatial and spherical diffusion domain with respect to the underlying anatomy. Previous image registration techniques have been described to estimate

the between slice fetal head motion, allowing the reconstruction of 3-D a diffusion estimate on a regular grid using interpolation. We propose Approach to Unified Diffusion Sensitive Slice Alignment and Reconstruction (AUDiSSAR) that explicitly formulates a process for diffusion direction sensitive DW-slice-to-DTI-volume alignment. This also incorporates image resolution modeling to iteratively deconvolve the effects of the imaging point spread function using the multiple views provided by thick slices acquired in different anatomical planes. The algorithm is implemented using a multi-resolution iterative scheme and multiple real and synthetic data are used to evaluate the performance of the technique. An accuracy experiment using synthetically created motion data of an adult head and a experiment using synthetic motion added to sedated fetal monkey dataset show a significant improvement in motion-trajectory estimation compared to a state-of-the-art approaches. The performance of the method is then evaluated on challenging but clinically typical *in utero* fetal scans of four different human cases, showing improved rendition of cortical anatomy and extraction of white matter tracts. While the experimental work focuses on DTI reconstruction (second-order tensor model), the proposed reconstruction framework can employ any 5-D diffusion volume model that can be represented by the spatial parameterizations of an orientation distribution function.

Index Terms

Diffusion tensor image (DTI); fetal imaging; motion-estimation; multi slice MR; reconstruction

I. Introduction

Fetal brain diffusion imaging is emerging as a new and valuable tool to assist in the clinical evaluation of pregnancies. The use of slice based measurements of scalar diffusion properties of the fetal brain has allowed the mapping of the developmental trajectory of white matter tissue properties *in utero* [1]–[4]. In challenging clinical conditions, *in utero* diffusion weighted (DW) magnetic resonance imaging (MRI) has been shown to provide a valuable marker for acute hypoxic ischemic fetal brain lesions [5], [6]. More recently it has been used in the mapping of abnormalities of the laminar structure of the fetal brain in Cobbelstone complex [7]. Berman [8] also reported higher diffusivity in parietal white matter and the thalamus in fetuses with congenital heart defects when compared to controls. In addition to scalar microstructural properties from postmortem studies of the fetal brain, full 3-D diffusion direction and tractography measurements have illustrated the possibility of mapping the development of white matter connectivity using MRI [9], [10]. This research has motivated the first attempts at *in utero* diffusion tensor based tractography studies [11], [12] that have shown the possibility of mapping the emergence of white matter connections *in utero* in cases of limited fetal head motion. However for practical clinical applications, more robust approaches that can deal with fetal head motions are required to allow reliable estimation of tissue microstructure. In this work, we address the problem using postprocessing to estimate slice to slice motions of fetal head anatomy occurring during multi-slice DW acquisitions. Here we implicitly assume that the use of a fast echo planar imaging (EPI) acquisition freezes within-slice motion for the majority of cases. Signal corruption arising from within-slice motion such as spin history effects are then addressed by a robust model fitting. The key challenge is then to estimate the changes in position and

orientation of the DW slices with respect to the underlying anatomy, and to then form a regularly sampled estimate of the diffusion profile across the fetal brain volume for the motion scattered slice data.

Fetal motion, limited spatial resolution and low signal-to-noise (due to the distance to the coils and the use of 1.5T imaging in clinical fetal studies) are the biggest challenges for the successful application of diffusion MRI to study the fetal brain *in utero*. Typical resolutions of clinical multi-slice diffusion weighted imaging (DWI) in adult studies are $2 \times 2 \times 2$ mm or $1 \times 1 \times 4$ mm which make low resolution a major limiting factor given the small scale of the anatomy studied.

A. Between Slice Motion Estimation and Reconstruction From Multi-Slice Diffusion MRI

Unlike adult MRI, where techniques such as optical tracking [13] or navigator echoes can be used to update the scanner measurement coordinates during imaging to compensate for motion, accurate fetal head motion estimation is more challenging because of the scale of the motion and the problem of localizing motion signal to the fetal head. As a result, methods for fetal head motion estimation have focussed on postprocessing of multi-slice data. For conventional structural studies, clinical imaging has employed fast multi-slice snapshot acquisitions. Registration-based structural volume reconstruction of fetal brain anatomy from these rapid 2-D MR sequences was first proposed by Rousseau *et al.* [14] and later by Jiang *et al.* [15]. Rousseau *et al.* suggested to alternate between volume reconstruction through Gaussian weighted averaging (GWA) with an anisotropic kernel and rigid slice-to-volume registration by maximizing the normalized mutual information (NMI) [16] between slices and the reconstructed volume, while Jiang *et al.* used B-spline regression for reconstruction and cross correlation to drive the registration. A limitation with these methods is that it is difficult to gain any knowledge on the convergence of the combined system as it alternates between the two independently formulated problems of volume reconstruction and slice alignment. Kim *et al.* [17] circumvented this problem by estimating relative motion between slices using intersection-driven-registration to align 2-D slices, which was followed by a single image reconstruction step.

Unlike the simple case of multi-slice MRI, a DW MRI sequence collects a set of $N + 1$ slice stacks $\mathcal{S} = \{S_0, S_1, \dots, S_N\}$ that, under ideal conditions, are related through the Stejskal-Tanner (ST) equation

$$S_i = S_0 e^{-bD_i} = \widetilde{S}_0 e^{-TE/T_2} e^{-bD_i} \quad (1)$$

to the underlying diffusion of water at a given location in the brain, where b is a sequence dependent constant (specifying the diffusion sensitivity), D_i is the apparent diffusion coefficient (magnitude of the diffusion in the $\mathbf{g}_i$ direction), S_0 is the nondiffusion weighting (nDW) volume, and $\widetilde{S}_0$ is nonattenuated signal, T_2 is an intrinsic tissue property, and TE is the echo time. Thus, given a diffusion model, $D(\cdot)$ and a sufficient number of DW volumes, $\{S_i\}$, it is possible to reconstruct a diffusion volume with an orientation distribution function (ODF) describing the diffusion in each voxel for a given b-value. Motion between

acquisition of the individual slices corrupts the spatial and spherical correspondence in the image coordinate system and thus the relationship in (1).

Jiang *et al.* [18] were the first to extend the structural registration-based reconstruction to DWI acquisitions by 1) reconstructing a nDW volume [15], 2) registering of the DW slices to the reconstructed nDW volume using NMI, and finally 3) fitting a second-order tensor diffusion model to the scattered data. However, they reported that the fractional anisotropy estimation for the *in utero* fetal brain was not robust due to low signal-to-noise ratio (SNR). Oubel *et al.* [19] suggested a spatial and spherical approximation framework to form a regularly sampled 5-D image. The nDW and the DW slices were registered separately using NMI to their own common reference, and the two sets of registrations were brought into the same coordinate system by global linear registration between the resulting nDW and DW reconstructions. A diffusion volume was then reconstructed by fitting spatio-angular radial basis functions (RBFs), which were defined as the product of a spatial RBF and an angular RBF, to the scattered diffusion data. A limitation of the techniques of [18] and [19] is that they do not take the resolution of the acquired slices into account when reconstructing a diffusion volume, which implicitly limits the resolution of the reconstructed diffusion volume.

B. Resolution Enhancement From Multiple Diffusion MRI Acquisitions

In conventional MRI, super resolution reconstruction (SR), that aims to solve the inverse problem of optimal image formation from multiple acquisitions, was introduced by both Gholipour *et al.* [20] who used the mean-squared-difference (MSD) to drive the slice-to-volume registration, and by Rousseau *et al.* [21] who used NMI. Recently, Fogtmann *et al.* [22] proposed an unified approach for between-slice-motion estimation and SR reconstruction of conventional structural images, where both problems are solved by minimizing the same robust fit functional.

SR approaches to DWI rely on multiple (likely anisotropic) DW acquisitions that are spatially located in a way that allows for the recovery of a high-resolution representation of the underlying anatomy. Scherrer *et al.* [23] proposed SR reconstruction from orthogonal anisotropic DW acquisitions by 1) volume-to-volume registration between the nDW volume and each of the DW volumes, 2) resampling of the DW volumes using Kriging to correct for the rotation in the rigid slice transformation, and 3) applying the SR technique in [20] on each sensitivity direction separately. The application of this technique is likely to fail in fetal diffusion MRI as it only accounts for between-volume-motion and neglects between-slice-motion. The acquisition time of each DW volume is typically around 10 s on a 1.5T clinical scanner, and in such a timeframe, the fetal head can exhibit appreciable motion. The SR approach in [23] accounts for the slice resolution but employs a sequence of independent 3-D reconstructions (one for each diffusion sensitivity direction), rather than a full 5-D reconstruction of the scattered diffusion data and thereby does not exploit correlation between diffusion sensitivity directions.

C. Unifying 3-D Image Formation and Slice Alignment

To address these limitations we propose a unified formulation which employs the Skejtal–Tanner equation and an image acquisition model to build a *fit* function between the estimated motion of the subject, the estimated underlying anatomy (high-resolution representation) and the acquired DWI data. By maximizing the fit with respect to subject motion and the reconstruction, we obtain a 5-D SR reconstruction framework from scattered DW slice data. This explicitly accounts for the change in diffusion measurement direction of a slice relative to the other slices, as its orientation with respect to the reconstruction frame is adapted. In this work, we model the 5-D diffusion volume by spatially parametrizing the rank two tensor model with a linear spline so that we focus on diffusion tensor imaging (DTI) with the data.

The proposed approach has similarities to the approaches in [18], [19], and [23] but differs significantly in a number of aspects and in its more general formulation. To obtain an initial motion estimate we apply a registration concept similar to that of [19] where the motion estimates for the nDW and DW slices are obtained independently using a structural registration-based approach separately on the set of nDW and a set of DW slices. The applied 5-D diffusion model also spatially parametrizes the second-order tensor model as in [18], but with the difference that we represent the tensors in Log-Euclidean space which avoids negative estimates of diffusion. In contrast to [18] and [19], the proposed unified approach models the spatial resolution of the DW slices which enables the estimation of a higher resolution diffusion representation of the underlying anatomy. While Scherrer *et al.* [23] performs SR reconstruction on DWI acquisitions, our method distinguishes itself from that of Scherrer *et al.* as it reconstructs a high-resolution 5-D diffusion volume instead of a sequence of independent 3-D reconstructions of each diffusion sensitivity direction.

To summarize, the novelties of the paper are the unification of the diffusion sensitive registration and the 5-D diffusion volume estimation as well as the SR estimation of the 5-D diffusion volume. Although this work is not the first to perform registration of diffusion volumes (examples include [24] and [25]), but to the best of our knowledge this is the first method described to align DW slices to a diffusion volume using a unified model-based reconstruction formulation. Although the application focus of this work is fetal imaging, the suggested approach is not limited to this application.

II. Slice Positioning, Diffusion Sensitivity, and Motion

A. Spatial Coordinate Systems

Under optimal conditions MR scanners collect multi-slice DWI data uniformly in both the spherical and the spatial domains, allowing spatially local and independent diffusion estimation through the ST equation. In practice, each scan consists of a number of multi-frame stacks of 2-D DW and nDW images (slices), where each frame has its own diffusion sensitivity direction or no diffusion sensitivity. Each slice has its own 3-D mm coordinate-system *slice mm* (Fig. 2), where the center of the slice is placed in the origin of the coordinate-system in the *xy*-plane. This *slice mm* coordinate-system has a corresponding 2-D voxel coordinate-system *slice vox* which accounts for the in-plane sampling rate (voxel

size). The intended anatomically consistent 3-D coordinate frame of the underlying anatomy is then denoted as *ref mm*. In the simplest case, it is simply the scanner coordinate system, but it may then be adjusted to a global reorientation to standardized anatomical coordinates (AC-PC) collective corrections of the slice stacks away from their intended location. The *init mm* coordinate-system is the initial reference coordinate system.

In scanner coordinates, the acquired data are regularly positioned, as can be observed from the evenly spaced and parallel slices in Fig. 1(a) and the regularly spiked spherical histogram of the diffusion sensitivity directions in Fig. 1(c). For a sufficiently fast acquisitions such as single shot EPI, within-slice motion can be mostly neglected (or slice measurements excluded during robust model fitting), and the spatial and the underlying problem limited to estimating the true slice positioning and orientation in the anatomy. Due to subject motion, the slices do not correspond to the underlying anatomy [as illustrated by the mismatch between the head surface and head outlines in Fig. 1(a)]. To allow for a meaningful and geometrically correct reconstruction, the slices must be mapped to a consistent anatomical coordinate frame *ref mm* accounting for slice location [Fig. 1(b)] and diffusion measurement orientation [Fig. 1(d)]. The slices, when correctly mapped to the underlying anatomy, are no longer regularly spaced or parallel to each other causing the DWI data to become scattered in both the spherical (orientation) and the spatial domains in the reference anatomical coordinate system, as illustrated in Fig. 1. An interesting observation from Fig. 1(d) is that in this reference anatomical coordinate system, the DWI sensitivity can be significantly scattered in the spherical domain forming a dense sampling of the diffusion profile from repeated collection of only a small number of diffusion directions. This scattering means that a 5-D diffusion volume model, which spans both the spherical and spatial domain, would be more appropriate to approximate the acquired data.

B. Spatial Transformations

A set of DW slices is typically acquired with different spatial locations and different gradient-sensitization directions fixed in some global coordinate system. Thus, if the patient moves during acquisition, the slices are perturbed relative to the patient in both the spherical and spatial domain. In the case of rigid motion, we can correct the spatial location, $\mathbf{x}$, with the rigid transformation

$$\phi(\mathbf{x}; \mathbf{p}) = \mathbf{y} = \mathbf{R}(\boldsymbol{\alpha})\mathbf{x} + \mathbf{t} \quad (2)$$

where $\mathbf{p} = [\boldsymbol{\alpha}^T \quad \mathbf{t}^T]$ are the transformation parameters and $\mathbf{y}$ are the corrected spatial locations. Assuming that the diffusion measurement direction of the slice, $\mathbf{g}$, has been mapped to the slice coordinate system, we can compute the corrected direction with

$$\tilde{\mathbf{g}} = \mathbf{R}(\boldsymbol{\alpha})\mathbf{g}. \quad (3)$$

DW image geometry within each slice may also be corrupted by eddy-current effects that induce in-plane shearing, magnification and/or translation. In cases of severe eddy-current distortion, the simple rigid-motion may be inadequate to map the slices accurately to a

common anatomical coordinate system. Within the slice alignment process, it is possible to account for eddy-current effects by replacing the rigid transformation of a slice with a more general linear transformation

$$\phi(\mathbf{x}, \mathbf{p}) = \mathbf{A}(\mathbf{p})\mathbf{x} + \mathbf{t}, \quad \mathbf{A}(\mathbf{p}) = \mathbf{R}(\boldsymbol{\alpha})\mathbf{S}(c) \quad (4)$$

where $\mathbf{p} = [\mathbf{c} \ \boldsymbol{\alpha}^T \ \mathbf{t}^T]$ are the transformation parameters of the slice, and $\mathbf{S}$ is a scale-skewness matrix. As we make use of an eddy current nulling acquisition for most of the diffusion data acquired, this effect is negligible and we will not consider it in the following experiments. For the rest of the paper, we will use the short notation $\phi^k(\mathbf{x}) = \phi(\mathbf{x}; \mathbf{p}_k)$ for the transformation of the k th slice.

We let ϕ_i^k define the transformation of the k th slice to the *init mm* coordinate-system and let the estimated slice motion of the k th slice be given by ϕ_m^k , which maps the slice from *init mm* to *ref mm*. The combined slice transformation from the *slice mm* coordinate-system to the *ref mm* coordinate-system is denoted ϕ^k (Fig. 2). The set of slice motion transformations ordered over time can be termed a motion-trajectory. As all the transformations are linear, we can introduce a homogenous mapping matrix $\mathbf{M}^k$ which maps between *slice vox* of the k th slice and *ref vox*. The mapping matrix $\mathbf{M}^k$ is an important element to ensure a stable discretization.

III. Modeling Diffusion From Scattered DW Measurements

Jiang *et al.* [18] used a rank 2 tensor image (DTI) model where the tensor elements were parameterized spatially with a cubic B-spline. This is a simple and intuitive model that can be evaluated at a moderate computational cost. Potential problems with this representation are that the Euclidean averaging of tensors tends to introduce swelling effects, and the estimated tensors may become negative-definite (i.e., negative diffusion in certain directions) and therefore not physically meaningful. The 5-D spatio-angular RBF diffusion model in [19] is constructed from the Kronecker product between spatial Gaussian and angular geodesic distance RBFs. This is an elegant way of combining the spatial and spherical domains, and thus avoiding issues with averaging tensors. However, the angular RBF model is not rotation-invariant, which implies that the resulting diffusion volume estimate might be significantly biased by the chosen knots in the spherical domain so that the fewer the knots the larger the bias. In addition, the RBF model still allows physically implausible negative diffusion, and is computationally more expensive than the spatial parameterization of the rank 2 tensor. Here, we therefore prefer a model similar to that of Jiang *et al.*, by generalizing their model. Let $\Gamma(\mathbf{g}, \mathbf{q})$ be an ODF that is defined by the parameters $\mathbf{q}$ and that describes diffusion in the direction $\mathbf{g}$. Given an ODF, we define the diffusion at the spatial location $\mathbf{y}$ by

$$D(\mathbf{y}, \mathbf{g}, \mathbf{w}) = \Gamma(\mathbf{g}, \mathbf{W}\mathbf{b}(\mathbf{y})) \quad (5)$$

where $\mathbf{b}$ is any set of spatial basis functions, $\mathbf{g}$ is the diffusion direction, $\mathbf{w}$ are the weight parameters/coefficients that define the diffusion volume, and $\mathbf{W} = [\mathbf{w}]_{K \times L}$ is the diffusion

image parameter matrix of size $K \times L$ (where K is the number of ODF parameters, and L is the number of basis functions). In the case of the cubic B-spline parameterized rank 2 tensor model, $\mathbf{b}$ contains the evaluation of the cubic B-spline coefficients (mostly zero elements, and a maximum 64 nonzero elements), and is the rank 2 tensor model, i.e.,

$$\Gamma_2(\mathbf{g}, \mathbf{q}) = \mathbf{g}^T \begin{bmatrix} q_1 & q_4 & q_5 \\ q_4 & q_2 & q_6 \\ q_5 & q_6 & q_3 \end{bmatrix} \mathbf{g} = \mathbf{g}^T \mathbf{Q} \mathbf{g} \quad (6)$$

which is simplest and most widely used ODF, but is considered insufficient to model diffusion in adult studies [26], where imaging times can allow high angular resolution acquisitions to separate complex crossing fiber structures.

In fetal imaging, the underlying anatomy is much simpler and the number of specific directions that can be acquired is limited by fetal motion, imaging field strength, and acquisition time.

For this reason this work focuses on the rank 2 tensor model. However, we would like to emphasize that the approach presented in this paper is not limited to a specific ODF model. In contrast to Jiang *et al.*, we represent the tensor in Log-Euclidean space which guarantees only positive diffusion, i.e.,

$$\Gamma_{\log}(\mathbf{g}, \mathbf{q}) = \mathbf{g}^T \exp(\mathbf{Q}) \mathbf{g} \quad (7)$$

where $\exp$ is the matrix exponential.

To ensure a fast evaluation we apply a linear spline model with uniformly spaced knots. A 3-D point $\mathbf{y} = [y_0, y_1, y_2]^T$ can be normalized by eight knots surrounding it to give $\hat{\mathbf{y}} = [\hat{y}_0, \hat{y}_1, \hat{y}_2]$. With $\mathcal{N}_{\mathbf{y}}$ defining the set of knot points surrounding $\mathbf{y}$ and $\mathcal{N}'_{\mathbf{y}}$ defining those that do not, we have

$$\mathbf{b}_{\mathcal{N}_{\mathbf{y}}}(\mathbf{y}) = \begin{bmatrix} 1 - \hat{y}_0 \\ \hat{y}_0 \end{bmatrix} \otimes \begin{bmatrix} 1 - \hat{y}_1 \\ \hat{y}_1 \end{bmatrix} \otimes \begin{bmatrix} 1 - \hat{y}_2 \\ \hat{y}_2 \end{bmatrix} \quad (8)$$

and

$$\mathbf{b}_{\mathcal{N}'_{\mathbf{y}}}(\mathbf{y}) = 0. \quad (9)$$

IV. Unified Motion-Estimation and Reconstruction in MRI

In this section, we present a general approach (not limited to diffusion MRI) which unifies motion estimation and reconstruction into a single formulation. This is essentially a generalization of the structural reconstruction approach we presented in [22]. Let I denote the underlying image intensity map that we are trying to reconstruct. In the case of diffusion

MRI, we have $I = D$. An approximation to the underlying image and an estimate of the motion-trajectory of the subject can be computed through the maximum a posteriori (MAP) estimate

$$\arg \max_{I, \mathcal{Q}} p(I, \mathcal{Q} | \mathcal{S}) = \arg \max_{I, \mathcal{Q}} p(\mathcal{S} | I, \mathcal{Q}) p(I, \mathcal{Q}) \quad (10)$$

where $\mathcal{S} = \{S_i\}_{i=1}^N$ is the set of N slices and $\mathcal{Q} = \{\phi_m^k\}_{k=1}^N$ is the motion-trajectory represented by the set of individual slice motion transformations. Assuming statistical independence between the noise of the acquired slices and independent image and motion priors, the problem simplifies to

$$\arg \max_{\mathcal{Q}, I} p(\mathcal{Q}, I | \mathcal{S}) = p(\mathcal{Q}) p(I) \prod_k^N p(S_k | \phi_m^k, I) \quad (11)$$

where $p(S_k | \phi_m^k, I)$ is the likelihood function, $p(I)$ is the image prior and $p(\mathcal{Q})$ is the motion prior.

A. Likelihood Function

As the noise in MR modulus imaging data transferred from the scanner is Rician distributed [27] the *optimal* likelihood function is

$$p(S_k = s_k | I) = \prod_{\mathbf{x} \in s_k} \frac{s_k(\mathbf{x})}{\sigma^2(\mathbf{x})} \exp \left(-\frac{s_k(\mathbf{x})^2 + E[S_k | I, \phi_m^k](\mathbf{x})^2}{2\sigma^2(\mathbf{x})} \right) \times I_0 \left(\frac{s_k(\mathbf{x}) E[S_k | I, \phi_m^k](\mathbf{x})}{\sigma^2(\mathbf{x})} \right)$$

where I_0 is the zeroth-order Bessel function, σ is the noise level, and $E[S_k | I, \phi_m^k]$ is the expected value of the i th acquired slice, S_k , given the underlying image, I , and the slice motion transformation, ϕ_m^k . Note that the type of the reconstruction is governed by the expected slice value, $E[S_k | I, \phi_m^k]$.

The noise level is unfortunately unknown and, unlike adult head studies, can be significantly spatially varying in abdominal imaging where the distance between the anatomy and the coils can vary significantly. The noise therefore needs to be estimated simultaneously with the image and transformation parameters. Additional motion artifacts such as spin history, within slice motion artifacts on the diffusion and residual slice registration errors might also corrupt the Rician noise model. We therefore use the robust Huber likelihood function in the model fitting procedure to reduce the number of parameters to be estimated and to increase the robustness to within-slice motion artifacts. The Huber likelihood function is given by

$$p(S_k = s_k | I) = \frac{1}{\sigma} \prod_{\mathbf{x}} \exp \left(-\gamma(\mathbf{x}) \mathcal{L}_\delta(s_k(\mathbf{x}) - E[S_k | I, \phi_m^k](\mathbf{x})) \right)$$

where $\gamma(\mathbf{x})$ is a masking image function approximately separating the fetal brain from the surrounding maternal tissue, and $\mathcal{L}_\delta$ is the Huber loss function, i.e., for otherwise.

$$\mathcal{L}_\delta(a) = \begin{cases} \frac{1}{2}a^2, & \text{for } |a| \leq \delta \\ \delta \left(|a| - \frac{\delta}{2}\right), & \text{otherwise.} \end{cases}$$

B. Image Prior

The prior term $p(I)$ which asserts knowledge on the image volume in the form of an uncertainty distribution. Most biological image reconstruction applications apply some form of spatial homogeneity prior such as smoothness [23] or total variation [28]. In this work, the prior term we use in (11) is

$$p(I) = \exp \left(-\alpha L_\delta \left(\int_{\Omega_y} \nabla_y I(\mathbf{y}) d\mathbf{y} \right) \right) \quad (12)$$

for nDW reconstruction and

$$p(I) = \exp \left(-\alpha \mathcal{L}_\delta \left(\int_{\Omega_g} \int_{\Omega_y} \nabla_y I(\mathbf{y}, \mathbf{g}) d\mathbf{y} d\mathbf{g} \right) \right) \quad (13)$$

for DW reconstruction, where α determines the strength of the prior and ∇_y is the spatial gradient operator. The use of the Huber norm ensures a piecewise smooth reconstructed image that is allowed to have local sharp transitions that can represent tissue boundaries, despite the penalty of the gradient operator. The prior is a mixture of a total variation and a quadratic gradient prior, and it can be converted into the form

$$p(I) = \exp \left(-\alpha \|L\mathbf{w}\|_{h,\delta} \right) \quad (14)$$

where $\mathbf{w}$ are the parameters of I (or D), L is the Tikhonov matrix and $\|\mathbf{a}\|_h, \delta = \sum_i \mathcal{L}_\delta(a_i)$ is the Huber norm.

C. Motion Prior

A simple model of fetal motion is the random walk model [29] where the motion is modeled by a succession of random rigid transformations. Assuming that the slices are temporally ordered from the timing of the DICOM data files that account for slice interleaves, the motion transformation of the k th slice can be computed by $\phi_m^k = \Delta\phi_m^k \circ \phi_m^{k-1}$, where $\Delta\phi_m^k$ is drawn from some random distribution. Given such a model, a reasonable prior would be

$$p(\mathcal{Q}) = \prod_k^{N-1} \exp \left(-\beta d \left(\phi_m^k, \phi_m^{k+1} \right)^2 \right) \quad (15)$$

where β determines the strength of the prior and d measures the distance between the two transformations. Assuming that the parameters, $\mathbf{p}_k$, of the ϕ_k are drawn from an independent normal distribution, we can define the distance between the transformations with

$$d(\phi_m^k, \phi_m^{k+1}) = \|\mathbf{p}_k - \mathbf{p}_{k+1}\|_2. \quad (16)$$

As fetal motion is in general smooth with occasionally large abrupt motions, the Euclidean distance measure will not always produce a meaningful evaluation of the disparity. Instead, we apply the Huber norm

$$p(\mathcal{Q}) = \prod_k^{N-1} \exp\left(-\beta \|\mathbf{p}_k - \mathbf{p}_{k+1}\|_{h,\delta}\right). \quad (17)$$

This motion prior was introduced for fetal imaging by Fogtmann *et al.* [22] as an alternative to motion parameterization in registration-based reconstruction in structural MRI.

V. Modeling the Image Point-Spread Function and Its Deconvolution

A challenge for the successful application of diffusion MRI on fetuses is the limited spatial resolution of the acquired DWI slices, especially given the significantly smaller fetal brain and the relatively low signal level in abdominal imaging which leads to the acquisition of multiple thick slices. To make use of the repeated acquisition of data with different slice orientations it is essential to incorporate a deconvolution component to the process of image formation. Simply assuming that slice voxels are simple points in space with no spatial distribution of sensitivity (as in [18], [19]), slices can be simulated from a diffusion volume and a motion-trajectory directly from the ST equation, i.e.,

$$\mathbb{E}[S_k | D, \phi_m^k](\mathbf{x}) = S_0(\phi^k(\mathbf{x})) e^{-b_k D(\phi^k(\mathbf{x}), \tilde{\mathbf{g}}_k \mathbf{w})} \quad (18)$$

where S_0 is the nDW volume and $\phi^k(\mathbf{x}) = \phi_m^k \circ \phi_i^k(\mathbf{x})$ (Fig. 2). This model however limits the resolution of the reconstructed volume to the resolution of the acquired data by dismissing the point spread function. To improve the resolution, we model the slice profile and the in-plane point spread function with a 3-D Gaussian kernel (as in Jiang *et al.* [15]), i.e.,

$$\mathbb{E}[S_k | D, \phi_m^k](\mathbf{x}) = \int G(\mathbf{z}) S_0(\phi^k(\mathbf{x} + \mathbf{z})) e^{-b_k D(\phi^k(\mathbf{x} + \mathbf{z}), \mathbf{w}, \tilde{\mathbf{g}}_k)} d\mathbf{z}. \quad (19)$$

If S_k is a nDW slice it can then be simulated directly from the DW volume using

$$\mathbb{E}[S_k | S_0, \phi_m^k](\mathbf{x}) = \int G(\mathbf{z}) S_0(\phi^k(\mathbf{x} + \mathbf{z})) d\mathbf{z}. \quad (20)$$

A stable discretization of the convolution in (19) and (20) is vital to ensure a stable optimization of (11). In this work, we discretize the convolution with the Gaussian kernel using the midpoint quadrature rule in *slice vox*. This discretization demands that the sampling is sufficiently dense 1) to represent the multi-dimensional Gaussian and 2) not to *jump over or miss too many voxel* (knots) S_0 in and D . We empirically determined that a sufficient sample spacing can be computed by

$$\Delta_i^k = \min \left(2 \sum_{i,i}, \frac{1}{2 \|\mathbf{m}_i^k\|_2} \right) \quad (21)$$

where $\mathbf{m}_i^k$ is the i th column of $\mathbf{M}^k$ (see Fig. 2) and Σ is the slice resolution in *slice vox*. The first argument to the min function ensures that the discretized Gaussian kernel sufficiently represents the Gaussian shape, and the last argument ensures that the sampling spacing in *slice vox* correspond to no more than half the voxel size in *Ref vox*, which sufficiently limits the number of missed voxels.

VI. Implementation

The implementation computes the MAP estimate in (11) using a gradient-based numerical optimization strategy that alternates between optimizing the motion-trajectory parameters while freezing the diffusion image parameters, and vice versa. We apply the L-BFGS[30] method with analytic first-order derivatives to solve (11) in a LOG-MAP configuration, i.e.,

$$\arg \min_{\mathcal{Q}, D} F = \arg \min_{\mathcal{Q}, D} -\log(p((\mathcal{Q}, D | \mathcal{S}))). \quad (22)$$

In practice, this means that alternating between the problems

$$\hat{\mathbf{w}} = \arg \min_{\mathbf{w}} \sum_{k=1}^N \sum_{\mathbf{x} \in \Omega_k} \gamma(\mathbf{x}) \mathcal{L}_{\delta_1} \left(s_k(\mathbf{x}) - \mathbb{E}[S_k | D, \phi_m^k](\mathbf{x}) \right) + \alpha \|\mathbf{L}\mathbf{w}\|_{h, \delta_2} \quad (23)$$

and

$$\hat{\mathbf{p}} = \arg \min_{\mathbf{p}} \sum_{k=1}^N \sum_{\mathbf{x} \in \Omega_k} \gamma(\mathbf{x}) \mathcal{L}_{\delta_1} \left(s_k(\mathbf{x}) - \mathbb{E}[S_k | D, \phi_m^k](\mathbf{x}) \right) + \beta \|\mathbf{p}_k - \mathbf{p}_{k+1}\|_{h, \delta_3}. \quad (24)$$

However, unlike the methods in [18], [19], [23], the same criteria are being optimized for both the motion estimation and reconstruction steps. Thus, the approach is unified and diffusion sensitivity is explicitly incorporated within the criteria.

The optimization problem is highly nonlinear, and thus good initial estimates of the motion-trajectory and the diffusion volume must be provided to achieve a globally optimal solution to (22). Fig. 3 outlines the key steps in the algorithm. The data is initially loaded into the SLIMMER tool [31] to obtain an initial setup. A nDW volume is then estimated using a

structural motion estimation and reconstruction approach similar to that in [22], and an initial motion-trajectory for the DW slices is estimated with the registration concept of [19]. The next step is an initial robust DTI reconstruction as explained in Section VI-C, which is subsequently used to initialize the proposed DTI reconstruction. This is followed by two cycles of motion estimation and DTI reconstruction as described in Section VI-D.

A. nDW Volume Reconstruction Using Multi Resolution Refinement

To reconstruct a nDW volume we solve the problem in (11) with a slice profile model [(20)] that is similar to the approach in [22]. A combined multi-scale and multi-resolution technique is applied to minimize the chance of falling into a local minima. Slices are down-sampled using an image pyramid scheme, and the problem is initially solved at the lowest resolution. The solution of the previous level is then continuously used to initialize the solution of current level until full resolution is reached. During the lower resolution step, the image reconstruction step is replaced with GWA [14].

B. Initial DW Motion Estimation

An initial estimate of the motion-trajectory is computed from the DW slices using the DW registration concept by Oubel *et al.* [19]. Oubel *et al.* compute the motion-trajectory by matching the DW slices to a common scalar volume which is reconstructed from the DW slices and the current estimate of the motion-trajectory. The reconstructed scalar DW volume is subsequently registered to the nDW volume to bring the estimated motion-trajectories into the same coordinate frame. We apply the same structural registration approach as for the nDW volume estimation with the slight difference that only GWA is used for reconstruction. In this work, we test both NMI as in [19] and MSD to drive the registration. It is our experience that the SSD dissimilarity measure works better in almost all instances. This claim will be supported later in the experimental section, Section VII.

C. Initial DTI Reconstruction

One potential problem with the Rician and the Huber likelihood functions is that they make the optimization problem nonlinear [32]. A robust alternative, which can be applied to compute an initial estimate, D , is the log-linear least-squared-difference (LLLS) likelihood function

$$p(S_k | \phi_k, D) = \prod_{\mathbf{x}} \exp \left(- \left[\log \left(\frac{S_k(\mathbf{x})}{\int_{\mathbf{y}} G(\mathbf{y}) * S_0(\phi_k(\mathbf{x} + \mathbf{y})) d\mathbf{y}} \right) - b_k \int_{\mathbf{y}} G(\mathbf{y}) * D(\phi_k(\mathbf{x} + \mathbf{y}), \tilde{\mathbf{g}}_k) d\mathbf{y} \right]^2 \right). \quad (25)$$

The LLLSD likelihood with the second-order tensor representation in (6) yields a quadratic optimization problem for the reconstruction which is convex and easy to solve. The tensor matrices are subsequently converted to the Log Euclidean representation in (7)—in case of negative-definite tensors the eigenvalues of the tensor are replaced by a small but positive number.

D. Diffusion Sensitive Slice-to-Volume Registration

After the initial estimates have been computed, the algorithm first alternates between slice-to-volume registration with (18) and diffusion volume reconstruction (cycle 1). Finally, it alternates between slice-to-volume registration with (19), and diffusion volume reconstruction (cycle 2). In conjunction with the multi resolution image reconstruction, the motion prior is initially set high at the lowest image resolution to estimate the slowest motion components. As the image resolution is increased, the weight of the motion prior is halved between each level allowing the recovery of faster, discontinuous movements. This procedure forms one cycle of optimization. In practice, the use of the initial estimation and efficient optimization means that the algorithm converges quickly to the solution, taking around two iterations in cycle 1 and a single iteration in cycle 2. We set the motion prior weight at this step to half of the weight of the motion prior weight used to compute the initial motion-trajectory of the DW slices. The alignment problem is such that the coordinate system *ref mm* is floating and may drift in relation to *init mm*, where the initial brain mask was defined. Drifting is eliminated by enforcing the *geometric mean* $\{\phi_m^k\}$ of to be the identity transformation during the optimization. The *ref vox* coordinate frame defines the voxel resolution of the reconstruction by specifying the voxel dimension and size in relation to the *ref mm* coordinate system.

E. Parameter Selection and Algorithm Configuration

There are five parameters $\alpha, \beta, \delta_1, \delta_2, \delta_3$, which have to be chosen prior to the optimization. We set the parameters δ_1 and δ_2 to 20%–30% of the dynamic range of DW slices, while good values for α and β are dependent on the noise level of the acquired data and the number of acquired stacks. We therefore obtain α and β empirically for each dataset. We typically set δ_3 to 0.5 mm and 0.5° or lower.

The algorithm is implemented in C++ using a standard open-source shared memory OpenMP framework for multi-threading as the problem is highly parallelizable. The problem is computationally expensive to solve due to the slice resolution modeling. For a typical clinical study consisting of 5–6 DWI image stacks with 10–20 diffusion sensitivity directions, the implementation described here requires approximately 3 h to compute the entire solution running on a 24-CPU (Intel Xeon 1.87 GHz) machine. The computation time is mostly dependent on the amount of data and the relationship between the slice resolution and reconstruction resolution.

VII. Methods for Experimental Evaluation of Algorithm

A. Evaluation of Slice Registration Accuracy

The aim of this experiment is to examine the fine scale accuracy of the slice alignment process with respect to a known estimate. This is not feasible using actual *in utero* imaging, due to the possible presence of motion in any living fetal study (arising for example from maternal breathing, as well as fetal motion). Here, we made use of a single high quality adult diffusion study as a basis for slice registration accuracy experiments. This was acquired on a 3T Siemens Trio scanner using a single shot EPI sequence with 60 noncollinear diffusion gradient measurement directions with imaging parameters $b = 700 \text{ s/mm}^2$, $TE = 108 \text{ ms}$, TR

= 8000 ms. This has an isotropic voxel resolution of $2.2 \times 2.2 \times 2.2$ mm and high signal to noise in comparison to abdominal fetal studies at 1.5T. This reference data was used to create multiple synthetic multi-slice datasets created to simulate acquisition with known head trajectories and interleaves. Thick slices were formed from this isotropic resolution data using a model of the slice profile of typical fetal scans. A total of nine approximately orthogonal DWI stacks (three axial, three sagittal, and three coronal) were simulated from a given motion-trajectory with 20 noncollinear diffusion gradient directions with $b = 700$ s/mm², a voxel size of $2.5 \times 2.5 \times 5$ mm and an interleave of 2. Rician distributed noise was added to the simulated slices with standard deviation of 15 (where grey and white matter have intensities of around 200 and 400 in the $b = 0$) to form the higher level of noise seen in 1.5T abdominal imaging. To faithfully evaluate the performance of the slice alignment we apply a random walk model to simulate fetal head motion, where a motion-trajectory is generated by sampling a random rigid transformation from a Gaussian distribution for every slice. The *full* motion of the k th slice in time was then computed by composing the random transformation of the slice with the transformations of the $k - 1$ previous slices. Ten artificial multi-stack studies were generated with five different levels of motion (0.5, 0.75, 1.0, 1.25, 1.5 std. in mm/degrees) and two studies for each level. An example of a simulated slice stack is shown in Fig. 4.

From a known trajectory, motion estimation errors are traditionally quantified using the root mean squared (rms) measure. For multi-slice fetal studies we can formulate this measure as

$$\text{rms}[\phi_{\text{ref}}, \phi_{\text{est}}, \phi_{\text{dft}}, \Omega, \gamma] = \sqrt{\frac{\sum_{\mathbf{x} \in \Omega} \gamma(\phi_{\text{ref}}(\mathbf{x})) \|\phi_{\text{est}}(\mathbf{x}) - \phi_{\text{ref}}(\mathbf{x}) - \phi_{\text{dft}}(\mathbf{x})\|_2^2}{\sum_{\mathbf{x} \in \Omega} \gamma(\phi_{\text{gt}}(\mathbf{x}))}} \quad (26)$$

where ϕ_{ref} is the reference or ground truth transformation, ϕ_{est} is the estimated transformation, ϕ_{dft} is a transformation which accounts for global drifting of the solution, Ω is the slice domain where the rms error is computed, and finally γ is a masking approximately separating rigid fetal brain and maternal tissue. Given a set of motion estimates, we can compute the drifting by solving

$\phi_{\text{dft}} = \arg \min_{\phi_{\text{dft}}} \sum_{i=1}^N \text{rms}[\phi_{\text{ref}}^i, \phi_{\text{est}}^i, \phi_{\text{dft}}, \Omega_i, \gamma]$ We evaluate a summary of the accuracy using the mean of the slice rms errors which we denote MSRMS. We also evaluate the mean voxel distance (MVD) defined by

$$\text{MVD} = \frac{\sum_{i=1}^N \sum_{\mathbf{x} \in \Omega_i} \gamma(\phi_{\text{gt}}^i(\mathbf{x})) \|\phi_{\text{est}}^i(\mathbf{x}) - \phi_{\text{ref}}^i(\mathbf{x}) - \phi_{\text{dft}}(\mathbf{x})\|_2}{\sum_{i=1}^N \sum_{\mathbf{x} \in \Omega_i} \gamma(\phi_{\text{gt}}^i(\mathbf{x}))}. \quad (27)$$

The MSRMS measure may be polluted by outlier slices that incorrectly fall outside the mask during the motion estimation. The positioning of these slices is purely determined by the motion-prior, and the positioning of neighboring slices in the mask that drive this prior. The magnitude of the errors for these slices are not directly meaningful as they do not contribute to the reconstruction within the mask, but they do however represent lost data. As such we

exclude these slices from the computation of MSRMS, and separately compute the fraction of lost data

$$LD = \frac{\sum_{i=1}^N H(\sum_{\mathbf{x} \in \Omega_i} \gamma(\phi_{\text{est}}^i(\mathbf{x}))) \sum_{\mathbf{x} \in \Omega_i} \gamma(\phi_{\text{gt}}^i(\mathbf{x}))}{\sum_{i=1}^N \sum_{\mathbf{x} \in \Omega_i} \gamma(\phi_{\text{gt}}^i(\mathbf{x}))} \quad (28)$$

where H is the unit step function. MSRMS and MVD are both sensitive to outliers but MSRMS is more sensitive, as all slices contribute equally to the MSRMS value independently of the intersection with the mask, and specifically because slices with a small mask intersection are more likely to become outliers. Thus a small difference between the values of MSRMS and MVD indicates that there are only a few outliers, and a large difference in the values of MSRMS and MVD indicate that outliers significantly influence the performance.

B. Evaluation of Slice Registration Precision and Capture Range for In Utero Anatomy

The aim of this experiment is to evaluate the capability to recover different scales of slice alignment in the presence of confounding maternal anatomy around the fetal head. As the emphasis is not fine scale accuracy, but anatomic realism, we have used an animal study as a basis for the motion experiments. Diffusion MRI data were acquired from two sedated (1% isoflurane) macaque fetuses *in utero* at gestational ages 85, 110, and 135 days (of a total gestational term of 165 days) using a Siemens 3T Tim Trio scanner at the Oregon National Primate Research Center. These gestational ages correspond approximately to human postmenstrual ages of 21, 27, and 33 weeks, respectively [33]. Each study consisted of nine axial, nine coronal, and nine sagittal single shot EPI sequences with eddy current compensation along 20 noncollinear diffusion gradient directions with $b = 500 \text{ s/mm}^2$, $TE = 93 \text{ ms}$ and $TR = 5000 \text{ ms}$. The slice resolution for all the studies was $1.125 \times 1.125 \times 2.000 \text{ mm}^3$. The primary source of motion was due to maternal breathing. We selected the second time point of the dataset as it was the least influenced by maternal breathing motion and the gyrification pattern at this gestational period reasonably represents the level of anatomical complexity for early human clinical studies. The study was truncated to contain only the first nine of the 27 stacks to create a dataset with clinically realistic extent.

A direct approach to evaluating the precision and capture range of the motion estimation is to examine the slice motion parameters obtained by applying the approach with different starting estimates and evaluating the distributions of solutions. As with the accuracy experiment, a random walk model was applied to produce ten perturbed initialization studies with five different levels of slice perturbation (0.5, 0.75, 1.0, 1.25, 1.5 std. in mm/degrees) and two repetitions. To quantitatively evaluate the precision we evaluated the same measures used in the accuracy evaluation described previously.

C. Evaluation of SR Reconstruction to Isotropic Voxel Resolution

To investigate final image quality and the ability to estimate diffusion at isotropic voxel resolution from the highly anisotropic 3 mm slices, all sets of macaque monkey scans were reconstructed to form a 0.75 mm isotropic DTI volume with the unified approach with MSD

motion initialization. To compute the reconstructions, we used the parameters $\beta = 1e - 2$, $\delta_3 = \infty$, $\delta_1 = \delta_2 = 40$ (approximately 20% of the dynamic range of S_0), and $\alpha = 3e^5$. α is set to a low value since the number of data sets larger. The motion prior δ_3 is set to infinity as the primary subject motion is caused by relatively slow and smooth maternal breathing.

D. Practical Evaluation on Clinically Typical In Utero Human Fetal Data

Because of the possibly more challenging nature of typical clinical data in humans (lower signal level from 1.5T imaging, larger motion), we applied the proposed method to a set of unsedated clinical studies acquired at Seattle Children's hospital using a 1.5T Siemens Avanti scanner at Seattle. Each consisting of 4–5 single shot EPI sequences using eddy current compensation pulse sequences. Subject information and imaging parameters are included in Table I, and selected DWI stacks for each subject are displayed in Fig. 5. The fourth fetus studied was the most neurologically normal with only an initial diagnosis of mild ventriculomegaly. The variation in scan contrast and quality is illustrated in Fig. 5, and arises because of variations in signal strength due to coil and fetal positioning in different clinical studies.

VIII. Experimental Results

A. Results of Slice Registration Accuracy

The results of the accuracy experiments using synthetically corrupted adult imaging data are summarized in Table II, and the best results for each study are underlined. A key observation from Table II is that there is a large and consistent drop in both MSRMS and MVD between the initialization (zeroth iteration) and the first iteration in all cases. This clearly motivates the use of the diffusion sensitive registration. Between the 2–3 iterations, we also observe a large and consistent drop in MSRMS and MVD, which indicates that modeling slice resolution in the slice registration improves the registration. Furthermore, the SSD motion initialization provided equal or superior performance in all studies compared to NMI, which confirms the expectation that making use of our knowledge of the problem and enforcing direct agreement of diffusion measurements provides a more unique alignment criteria. Another interesting observation is that there is no consistent improvement in MSRMS and MVD between 1–2 iterations for NMI, while the SSD studies for larger motions than 0.5 show improvement. The amount of lost data tends to drop between the initialization and the first iteration and then slightly grows. This slight growth can be explained by outliers close to the edge of the mask that are gradually being pushed outside the brain match region to reduce intensity disagreements. For the SSD initializations, the MSRMS values are generally 50% larger than the MVD values. This suggests that the outliers are corrupting the MSRMS and MVD measures to a large extent.

To statistically evaluate whether the improvement between two successive iterations is significant we applied a nonpara-metric paired permutation (PP) test with the null hypothesis, that the two sets of rms errors have the same mean, and the alternative hypothesis, that the latest iteration has a smaller mean. 50 000 permutations were used to compute the p -value for the PP test using the full set of rms errors. The PP test revealed significant (given a 0.05-level) reductions in mean rms between successive iterations for all

SSD studies except for the first study, between the first and second iteration, which had the lowest initial motion level. For the NMI studies, the mean rms was significantly reduced between the 0–1 and 2–3 iterations. The null hypothesis cannot be rejected for the middle motion levels between the 1–2 iterations as they have larger p-values than 0.05. Significant reductions (given a 0.05-level) in the dispersion of the rms error were observed between the 0–1 iterations and the 2–3 iterations in all studies using NMI-based motion initialization. Between the 1–2 iterations significant reductions in rms dispersion were observed only for larger motion levels (> 1.25). A similar pattern for the rms dispersion can be observed for the SSD studies with the exceptions that a p-value of 0.98 for the first study at motion level 1.25 does not support a rejection of the null hypothesis, and that the 1.00 studies have a significant reduction in rms dispersion.

Fig. 4 (right most column) shows three plots of the rms errors against the mask intersection in voxels for the first study for motion levels 0.50, 1.00, and 1.50. The plots largely support the previous observations that both the rms mean and dispersion are significantly reduced between the iterations. Also, slices with a mask intersection less than 1000 voxels have a drastically worse performance than slices with larger mask intersection. This is not surprising because slices close to the surface of the brain have less data to guide the registration, and will also generally have a larger initial rms due to the larger distance from the center of rotation of the head. A conservative estimate from Fig. 4 would be that the mean rms error without outliers would be approximately 50% less those values recorded in Table II. To summarize the results of this experiment, we can conclude that the approach produced motion estimates with sub-voxel accuracy for a slice resolution of $2.5 \times 2.5 \times 5$ mm.

B. Results of Slice Registration Precision and Capture Range of In Utero Anatomy

The results of this experiment using synthetically corrupted macaque fetal data (Table III) show a continuous improvement in precision measured in terms of MSRMS and MVD between successive iterations. Similar to the accuracy experiments, the MSRMS and MVD numbers in Table III clearly favor the SSD initialization compared to NMI initialization. There appears to be a larger drop in precision going between motion levels 1.00 and 1.25 for the SSD studies. This could be due to statistical fluctuation. Our investigation thus far has not provided another answer.

Similarly, we apply an Ansari–Bradley (AB) test to examine whether the dispersion in rms errors is significantly reduced between successive iterations. For the AB test, the null hypothesis is that the rms errors are identically distributed, and the alternative hypothesis is that the rms dispersion for the first iteration is greater than for the later iteration. All PP and AB test statistics on the rms errors favored the alternative hypothesis given a significance level of 0.01 which indicate that each iteration significantly improves the precision. The LD measure tends to increase between the 2–3 iterations, which is likely due to slices close to the boundary with weak matches getting pushed out of the mask. The least amount of lost data in the majority of cases is observed after the second iteration.

Taking a slice resolution of $1.125 \times 1.125 \times 3.000$ mm³ and outlier pollution of the MSRMS and MVD measures into consideration it is reasonable to conclude that the approach has

subvoxel precision for perturbation levels 0.5, 0.75, and 1.00. Also, it is unlikely in practice that fetal motion can approach the perturbation level of 1.00 given the fast slice acquisition speed which for this dataset was approximately four slices per second.

To illustrate the effect of the slice perturbations on the reconstruction and to relate MSRMS to reconstruction quality we show the motion corrected and nonmotion corrected FA maps for the first study with SSD initialization of each motion level in Fig. 6 (Top). The two first rows of images contain axial and sagittal views of the FA maps for the nonmotion corrected DTI reconstructions, and the last two rows contain axial and sagittal views of the FA maps for the motion corrected DTI reconstructions. The first column corresponds to the unperturbed initialization while the last five columns correspond to motion levels of 0.50, 0.75, 1.00, 1.25, and 1.50. From the first column we observe that the FA maps of the nonmotion and motion corrected reconstructions are almost identical that indicates that the subject motion under acquisition was, as assumed, limited. We also observe from the first two rows of FA maps that the perturbations had a major negative impact on the nonmotion corrected DTI reconstructions, and that only the FA map for the smallest perturbation level delineate any useful white matter structures or cortical diffusion anisotropy, with a big fall in quality occurring at a motion level of 0.75. The first four FA maps of the last two rows are similar, and we can conclude that MSRMS values of around 0.6 mm appear to have limited effect on the reconstruction quality. The FA maps for the last two motion levels provide a lower quality reconstruction than the previous motion levels. The cortical anisotropy is especially weakened while the larger primary white matter structures are less negatively affected.

C. Results of SR Reconstruction to Isotropic Voxel Resolution

Fig. 6 (Bottom) shows FA maps color-coded with the direction of the first principal component of the tensor and overlaid on a T2-weighted reconstruction for each time point study. The three rows correspond to the three gestational ages examined, and the first three columns contain selected axial, sagittal and coronal views. The fourth column shows coronal views of the directional FA maps computed from one of the diffusion datasets with one slice plane using an ordinary least-square and voxel-wise DTI reconstruction. The final column shows the Log-Euclidean average of the DTI reconstructions of all stacks after they have been mapped to the same coordinate system neglecting motion. Even with sedation, maps obtained using a single dataset without motion estimation (fourth column in the lower portion of Fig. 6) do not provide information suitable for quantitative analysis of development within white matter or the cerebral cortex. This firstly shows that the quality of the individual fetal data was generally low, even given the sedation. The FA maps of the average DTI reconstructions of all stacks displayed in the fifth column of Fig. 6 (bottom) reveal some key structures including the corpus callosum. In contrast the reconstruction and slice motion estimation provide both improved signal to noise and better spatial resolution than is available from the individual raw datasets or a simple combination of the raw datasets as acquired.

As expected, cortical FA decreases from maximums of around 0.8 to virtually zero over the period spanned by the three time points. The reduction in cortical FA is highlighted at the

top of Fig. 7, which shows the brain surfaces color-coded with the FA values for all three time points for the first monkey. The reduction in cortical FA observed here corresponds closely to the reduction previously reported in analyses of postmortem fetal brains obtained from nonhuman primate subjects obtained at this stage of gestational development [34], [35]. Tractography analyses using Fiber Assignment by Continuous Tracking [36] (FACT) was also performed, to explore the development of the corpus callosum using a manually selected seed region drawn around the midline plane of the corpus callosum in the T2W image at each time point. Results are displayed in the lower part of Fig. 7. In contrast to the case within the cerebral cortex, extracted white matter tracts using FACT become more discernible over time, changing from mostly simple straight tracts through and perpendicular to the corpus callosum (left, 85 days), to the inclusion of the splenium tracts (middle, 110 days), and finally to a white matter tracts resembling that which can be observed in an adult brain (right, 135 days). This can be a result of increasing water diffusion anisotropy with gestational age. Note that these tracts do not represent physical nerve fibers but only trajectories through the data that could run in the directions of actual nerve fibers.

D. Results of Practical Evaluation on Clinically Typical In Utero Human Fetal Data

All subjects were reconstructed to form a DTI volume using the proposed unified slice motion estimation approach to a 1 mm isotropic voxel resolution. To visually evaluate the between slice motion correction, Fig. 8 displays slice stack views of the second subject scan before and after motion estimation, and three selected slice intersections before and after motion estimation. The red line in Fig. 8 indicates the intersection of the slice with the mask that determines the boundary between the fetal head and maternal tissues. Comparing the noncorrected and corrected slice views in Fig. 8, a convincing improvement in alignment is apparent for the selected slices. The green circles highlight areas with an initial volume alignment mismatch that is corrected after slice motion estimation. The final reconstructions are shown in Fig. 9. The DTI reconstruction of case 2 supports the diagnosis of ACC, while the DTI reconstruction of case 1 only partially supports the initial diagnosis of ACC with the identification of a small band of connective diffusion in the region of the corpus callosum. From the directionally color-coded FA map on the first row of Fig. 9, it appears that the first subject has a thinner than normal corpus callosum. Case 3 suffers from severe VM but a close examination of the third row of Fig. 9 reveals that the fetus has clear a corpus callosum, along with other expected tracts. Due to the enlarged ventricles, white matter tracts closely wrap around the ventricles. The DTI reconstruction of the fourth case, with relatively normal brain anatomy (initial diagnosis as mild VM) as shown in Fig. 9 exhibits clear and detailed white matter structure typical of this age, as well as a radial diffusion pattern in the cortical plate (also visible in the other fetuses). Examination of the results using volume and no motion estimation, shown on the right two columns, clearly support the use of slice registration in the case of clinical imaging with, in particular, case 1 and 3 showing delineation of the corpus callosum that was not visible in the reconstructions without slice alignment.

To further evaluate the performance on clinical data we applied a standard ROI seed based FACT tractography to examine the corpus callosum connectivity in a comparison of the mild VM case (4) and the case with confirmed agenesis of the corpus callosum (case 2). For

the mild VM case 4, shown in Fig. 10(a), the improvement provided by slice motion correction in the connectivity measurement is clearly visible from tracts emerging from the ROI placed in the anterior corpus callosum that conform better to that expected in this region. Comparison between the tractography of the whole corpus callosum in the two cases, shown in Fig. 10(b), reveals only a small number of random voxel traces from the corpus callosum extracted in case 2, while a clear whole corpus callosum connectivity model is extracted from case 4.

IX. Discussion

This paper presented a new method for diffusion weighted slice motion estimation and 3-D isotropic DTI reconstruction specifically aimed at the challenging problem of fetal brain imaging. The key novelties of the approach are a unified reconstruction-alignment formulation yielding a diffusion sensitive slice registration model, and the incorporation of point spread function deconvolution into the 3-D image reconstruction. This allows the formation of 3-D datasets with isotropic spatial resolution from multiple slice scattered DWI acquisitions in different anatomical planes.

A. Evaluation of Unified Model Based Diffusion Sensitive Slice Alignment

In the current literature, there exists no gold standard for the validation of data alignment in diffusion MRI. Haselgrove *et al.* [37], Nielsen *et al.* [38], and Rohde *et al.* [39] performed visual comparisons to assess differences, which we have also used in this paper. Visual comparisons however can be biased by the observer and a more objective comparison should be done by comparing quantitative measures computed from the estimated results. A simple quantitative measure of DTI quality is to evaluate the mean FA in the cerebrospinal fluid (CSF), assuming that the FA should be zero if measurements are correctly aligned, or alternatively the standard deviation of the CSF diffusion signal. However, a low or consistent value does not guarantee alignment accuracy of all regions of all slices, since they may not contain CSF, or the CSF may be assigned a consistently incorrect FA value in the case of minimizing its standard deviation. In contrast, Nielsen *et al.* [38] proposed to use the entropy of the FA over the entire brain to evaluate the performance under the assumption that the misalignment of DW slices presents it-self as a noisy measurement of the diffusion signal. However, we and others [19] have found this is not robust to more severe misregistration, and over-fitted tensor maps (e.g. reconstruction with no image prior) may exhibit good entropy values. Because fetal imaging itself may always contain some amount of unknown motion, an alternative is to use simply adult data acquired with and without motion, allowing a comparison of the final motion corrected dataset to the nonmotion data. However, this suffers from the presence of significant geometric susceptibility artifacts which change as the head moves, making the problem of slice correction more complex than the case in most fetal studies.

Here, we instead chose to make use of different carefully created simulations to evaluate slice alignment estimates, using a combination of stationary adult and sedated fetal data to examine accuracy and capture range in the presence of confounding maternal anatomy. These experiments have advantages and disadvantages: A major issue with synthetic data is obviously the fact that it is not possible to model every possible aspect that might that

appear in real data: The compositional nature of the random walk may not capture all types of motion well. The approach though avoids any bias by using a sampling from a lower dimensional parametric model which is itself then used in the registration process. There is however still the possibility for bias in the use of the random walk model to both simulate and recover motion. However, we ensured the use of a consistent model for all methods in the comparison so that no method is favored, leaving only the possibility for over estimating the capability of the approach to deal with real motion.

Our set of focussed experiments looking at accuracy of thick slice alignment, and precision of slice alignment when estimating motion within maternal anatomy have allowed us to compare similarity measures and to meaningfully evaluate accuracy and final reconstruction quality with respect to key real world factors. The accuracy results indicate it is reasonable to conclude that the diffusion sensitive registration (column “3,” Table II) significantly improves the registration accuracy compared to the (state-of-the-art) structural-based registration approach that closely matches that of Oubel *et al.* using NMI (column “0,” Table II). The precision and robustness experiment on sedated macaque fetal data with simulated motion also revealed that the proposed approach outperformed the state-of-the-art in terms of the precision and capture range in the presence of maternal anatomy. The strength of this experiment compared to the accuracy experiment is that it is a validation on *in utero* data, but in contrast it cannot tell us true accuracy as there is no completely stationary estimate available.

B. Resolution Enhancement From Multiple Fetal DWI Acquisitions

For clinical studies with limited numbers of acquisitions and higher noise levels it is possible to increase the sampling resolution (decreased voxel size), but the spatial resolution may not increase. If inadequate numbers of diffusion datasets are available, not all regions of the final DTI reconstruction may support resolution enhancement, as it is clearly dependent on the local density of the scattered diffusion measurements and also the absence of within-slice motion artifacts corrupting those measurements. The deconvolution process is by far the most computationally expensive element, so in clinical datasets there will be a practical trade-off between computation time and reconstruction resolution that depends on the amount of data that was possible to acquire within the clinical imaging time. The practical fetal image reconstruction problem is in many cases ill-posed, which is why an imaging prior is necessary to obtain a reasonable solution among the infinitely many possible solutions. The task of selecting the prior strength is essentially a problem of choosing a suitable balance between avoiding over-fitting to noise and forming an overly smooth estimate. Separately to image noise, treating slice voxels as scattered points in space may itself introduce a form of semi-structured noise: The values from two approximately coinciding thick-slice voxels, from slices having very different orientations, may differ substantially purely because of their slice profiles. In order avoid over-fitting to the additional noise, the motion prior needs to be strengthened, which will result in a smoother reconstruction, effectively resulting in the loss of resolution. Thus, even though it may not be possible to achieve SR it may make sense to still incorporate a form of computationally expensive point spread deconvolution within the reconstruction process.

C. Conclusion

We have presented a unified diffusion sensitive method for 3-D DTI reconstruction from motion scattered multi-slice DWI sequences acquired from moving fetal brain anatomy. The proposed unified approach, that incorporates a robust fitting to negate the effect of within slice motion artifact, outperformed the current state-of-the-art slice alignment method with respect to motion estimation in an accuracy and a precision experiment. The experiments also demonstrated that DW volume motion correction as used by [23] for adults may not be an adequate motion model in fetal imaging, where between slice motion may be a more significant factor. The work presented here is the first approach to support SR reconstruction for *in utero* fetal brain diffusion MRI. Through experiments on the clinical human fetal data and monkey fetal data, we have demonstrated that the approach has practical value when multiple slice planes can be acquired. In the clinical fetal experiment, 3-D DTI reconstruction supported the diagnoses obtained from T2-weighted MRI and provided additional information not available from conventional DTI studies of the fetus. We presented tractography results showing the development of white matter connectivity and cortical diffusion anisotropy maps that motivate the application of the approach to fetal growth modeling in human fetuses. For normal connectivity, slice motion estimation clearly improved the visual accuracy of tract estimation and revealed clear differences in known connectivity abnormalities. However, while the proposed approach promises to provide additional diagnostic information, the variability in image signal to noise level (due for example to coil placement) and number of acquisitions possible in clinical studies must be carefully considered when interpreting results. The recent development of methods to address fetal motion in conventional structural MR led to the first larger scale studies of tissue growth and cortical folding [40] *in utero*. The development of comparable diffusion imaging techniques promises to enable similar advances in the study of developing structural connectivity and tissue microstructure in the developing fetal brain.

Acknowledgments

This work was supported by the National Institutes of Health under Grant R01 NS 055964 (PI: C. Studholme) and Grant R01 NS 061957 (PI: C. Studholme). The work of M. Fogtmann was supported by a postdoctoral grant from The Villum Foundation, Denmark.

References

1. Bui T, Daire JL, Chalard F, Zaccaria I, Alberti C, Elmaleh M, Garel C, Luton D, Blanc N, Sebag G. Microstructural development of human brain assessed in utero by diffusion tensor imaging. *Pediatric Radiol.* 2006; 36:1133–1140.
2. Kim DH, Chung S, Vigneron DB, Barkovich AJ, Glenn OA. Diffusion-weighted imaging of the fetal brain *in vivo*. *Magn Reson Med.* 2008; 59(1):216–220. [PubMed: 18050314]
3. Schneider J, Confort-Gouny S, Le Fur Y, Viout P, Bennathan M, Chapon F, Fogliarini C, Cozzone P, Girard N. Diffusion-weighted imaging in normal fetal brain maturation. *Eur Radiol.* 2007; 17:2422–2429. [PubMed: 17404738]
4. Schneider MM, Berman JI, Baumer FM, Glass HC, JS, Jeremy RJ, Esch M, Biran V, Barkovich AJ, Studholme C, Xu D, Glenn OA. Normative apparent diffusion coefficient values in the developing fetal brain. *Am J Neuroradiol.* 2009; 30(9):1799–1803. [PubMed: 19556350]
5. Baldoli C, Righini A, Parazzini C, Scotti G, Triulzi F. Demonstration of acute ischemic lesions in the fetal brain by diffusion magnetic resonance imaging. *Ann Neurol.* 2002; 52(2):243–246. [PubMed: 12210800]

6. Girard N, Gire C, Sigaudy S, Porcu G, d'Ercole C, Figarella-Branger D, Raybaud C, Confort-Gouny S. MR imaging of acquired fetal brain disorders. *Child's Nervous Syst.* 2003; 19:490–500.
7. Widjaja E, Geibprasert S, Blaser S, Rayner T, Shannon P. Abnormal fetal cerebral laminar organization in cobblestone complex as seen on post-mortem MRI and DTI. *Pediatric Radiol.* 2009; 39:860–864.
8. Berman JI, Hamrick SEG, McQuillen PS, Studholme C, Xu D, Henry RG, Hornberger LK, Glenn OA. *In vivo* MRI of the fetal brain. *Am J Neuroradiol.* 2011; 32(2):21–22.
9. Huang H, Xue R, Zhang J, Ren T, Richards LJ, Yarowsky P, Miller MI, Mori S. Anatomical characterization of human fetal brain development with diffusion tensor magnetic resonance imaging. *J Neurosci.* 2009; 29(13):4263–4273. [PubMed: 19339620]
10. Mori S, Zhang J. Principles of diffusion tensor imaging and its applications to basic neuroscience research. *Neuron.* 2006; 51(5):527–539. [PubMed: 16950152]
11. Zanin E, Ranjeva JP, Confort-Gouny S, Guye M, Denis D, Cozzzone PJ, Girard N. White matter maturation of normal human fetal brain. An *in vivo* diffusion tensor tractography study. *Brain Behav.* 2011; 1(2):95–108. [PubMed: 22399089]
12. Kasprian G, Brugger PC, Weber M, Krssak M, Krampl E, Herold C, Prayer D. In utero tractography of fetal white matter development. *Neuroimage.* 2008; 43(2):213–224. [PubMed: 18694838]
13. Aksoy M, Forman C, Straka M, Skare S, Holdsworth S, Hornegger J, Bammer R. Real-time optical motion correction for diffusion tensor imaging. *Magn Reson Med.* 2011; 66(2):366–378. [PubMed: 21432898]
14. Rousseau F, Glenn OA, Iordanova B, Rodriguez-Carranza CE, Vigneron DB, Barkovich AJ, Studholme C. Registration-based approach for reconstruction of high-resolution in utero fetal MR brain images. *Acad Radiol.* 2006; 13(9):1072–1081. [PubMed: 16935719]
15. Jiang S, Xue H, Glover A, Rutherford M, Rueckert D, Hajnal JV. MRI of moving subjects using multislice snapshot images with volume reconstruction (SVR): Application to fetal, neonatal, adult brain studies. *IEEE Trans Med Imag.* Jul; 2007 26(7):967–980.
16. Studholme C, Hill D, Hawkes D. An overlap invariant entropy measure of 3-D medical image alignment. *Pattern Recognit.* 1999; 32(1):71–86.
17. Kim K, Habas PA, Rousseau F, Glenn OA, Barkovich AJ, Studholme C. Intersection based motion correction of multislice MRI for 3-D in utero fetal brain image formation. *IEEE Trans Med Imag.* Jan; 2010 29(1):146–158.
18. Jiang S, Xue H, Counsell S, Anjari M, Allsop J, Rutherford M, Rueckert D, Hajnal JV. Diffusion tensor imaging (DTI) of the brain in moving subjects: Application to in-utero fetal and ex-utero studies. *Magn Reson Med.* 2009; 62(3):645–655. [PubMed: 19526505]
19. Oubel E, Koob M, Studholme C, Dietemann JL, Rousseau F. Reconstruction of scattered data in fetal diffusion MRI. *Med Image Anal.* 2012; 16:28–37. [PubMed: 21636311]
20. Gholipour A, Estroff JA, Warfield SK. Robust super-resolution volume reconstruction from slice acquisitions: Application to fetal brain MRI. *IEEE Trans Med Imag.* Oct; 2010 29(10):1739–1758.
21. Rousseau F, Kim K, Studholme C, Koob M, Dietemann J. On super-resolution for fetal brain MRI. *Med Image Comput Comput-Assist Intervent.* 2010:355–362.
22. Fogtmann M, Chapman T, Kim K, Seshamani S, Studholme C. A unified approach for motion-estimation and super-resolution reconstruction from structural magnetic resonance on moving subjects. *MICCAI Workshop Perinatal Paediatric Imag.* 2012
23. Scherrer, B.; Gholipour, A.; Warfield, SK. *MICCAI*. Vol. 2. New York: Springer; LNCS; 2011. super-resolution in diffusion-weighted imaging; p. 124–132.
24. Xue Z, Li H, Guo L, Wong ST. A local fast marching-based diffusion tensor image registration algorithm by simultaneously considering spatial deformation and tensor orientation. *Neuroimage.* 2010; 52(1):119–130. [PubMed: 20382233]
25. Park HJ, Kubicki M, Shenton ME, Guimond A, McCarley RW, Maier SE, Kikinis R, Jolesz FA, Westin CF. Spatial normalization of diffusion tensor mri using multiple channels. *Neuroimage.* 2003; 20(4):1995–2009. [PubMed: 14683705]
26. Barmpoutis, A.; Jian, B.; Vemuri, B.; Shepherd, T. Symmetric positive 4th order tensors and their estimation from diffusion weighted mri. In: Karssemeijer, N.; Lelieveldt, B., editors. *Information*

- Processing in Medical Imaging. Vol. 4584. Berlin, Germany: Springer; 2007. p. 308-319. Lecture Notes in Computer Science
27. Basu, S.; Fletcher, T.; Whitaker, R. Medical Image Computing and Computer-Assisted Intervention. Vol. 9. New York: Springer; LNCS; 2006. Rician noise removal in diffusion tensor MRI; p. 117-125.
 28. Joshi, SH.; Marquina, A.; Osher, SJ.; Dinov, I.; Van Horn, JD.; Toga, AW. MRI resolution enhancement using total variation regularization. Proc. 6th IEEE Int. Conf. Symp. Biomed. Imag.: From Nano to Macro; 2009; p. 161-164.
 29. Pearson K. The problem of the random walk. Nature. 1905; 72(1865):294.
 30. Liu DC, Nocedal J. On the limited memory BFGs method for large scale optimization. Math Program. Dec; 1989 45(3):503–528.
 31. Kim K, Habas PA, Rajagopalan V, Scott JA, Corbett-Detig JM, Rousseau F, Barkovich AJ, Glenn OA, Studholme C. Slimmer: Slice mri motion estimation and reconstruction tool for studies of fetal anatomy. Proc SPIE. 2011; 7962:79624D.
 32. Jones DK, Cercignani M. Twenty-five pitfalls in the analysis of diffusion MRI data. NMR Biomed. 2010; 23(7):803–820. [PubMed: 20886566]
 33. Clancy B, Kersh B, Hyde J, Darlington R, Anand K, Finlay B. Web-based method for translating neurodevelopment from laboratory species to humans. Neuroinformatics. 2007; 5(1):79–94. [PubMed: 17426354]
 34. Kroenke CD, Bretthorst GL, Inder TE, Neil JJ. Diffusion mr imaging characteristics of the developing primate brain. Neuroimage. 2005; 25(4):1205–1213. [PubMed: 15850738]
 35. Kroenke CD, Van Essen DC, Inder TE, Rees S, Bretthorst GL, Neil JJ. Microstructural changes of the baboon cerebral cortex during gestational development reflected in magnetic resonance imaging diffusion anisotropy. J Neurosci. 2007; 27(46):12 506–12 515.
 36. Mori S, van Zijl P. Fiber tracking: Principles and strategies—A technical review. NMR Biomed. 2002; 15(7–8):468–480. [PubMed: 12489096]
 37. Haselgrove JC, Moore JR. Correction for distortion of echo-planar images used to calculate the apparent diffusion coefficient. Magn Reson Med. 2005; 36(6):960–964. [PubMed: 8946363]
 38. Nielsen JF, Ghugre NR, Panigrahy A. Affine and polynomial mutual information coregistration for artifact elimination in diffusion tensor imaging of newborns. Magn Reson Imag. 2004; 22(9): 1319–1323.
 39. Rohde G, Barnett A, Basser P, Marengo S, Pierpaoli C. Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI. Magn Reson Med. 2003; 51(1):103–114. [PubMed: 14705050]
 40. Studholme C. Mapping fetal brain development in utero using magnetic resonance imaging: The big bang of brain mapping. Ann Rev Bio Eng. 2011; 13(1):345–368.

Fig. 1.

Illustration of motion induced scattering during a multi-slice DWI acquisition. Figure (a) shows a stack of acquired slices in the scanner coordinate system which are positioned regularly. But, due to rigid target motion during imaging, the slices do not corresponding the original underlying anatomy at the start of the acquisition. Figure (c) contains a distribution density plot of the diffusion measurement directions in the scanner coordinate system. The slices must be rigidly mapped to a common anatomical coordinate system to enable volumetric reconstruction, as in Figure (b). The intended diffusion measurement directions in Figure (c) must also be reoriented using the rotational component of the rigid motion transformation for each diffusion weighted slice, causing the data to be scattered in the spherical (orientation) domain. Figure (d) illustrates the resulting distribution of the diffusion measurement directions in the anatomical coordinate system, here generated from a synthetic motion-trajectory representative of a fetal head motion.

Fig. 2.

Coordinate systems and transformations. Diagram shows the mappings between a slice and the reference coordinate system, where the diffusion volume is reconstructed, as well as the mappings to and from an intermediate initial setup coordinate system. Both the slice and the reference coordinate system have a corresponding voxel coordinate system controlling the resolution.

Fig. 3.

Overview of pipeline: The data is loaded into the SLIMMER tool to obtain an initial volume alignment for each DWI slice stack refined from the planned DICOM coordinates. nDW volume is then estimated using the structural motion estimation and reconstruction approach in [22], and an initial motion-trajectory is estimated using the registration approach of [19]. The next step is an initial robust reconstruction as explained in Section VI-C, followed by the first DTI reconstruction with the proposed approach. This is followed by two cycles of slice motion estimation and diffusion reconstruction as described in Section VI-D.

Fig. 4.

Left three columns: Examples of simulated motion corrupted multi slice nDW data formed from a high quality stationary adult study on a 3T system. The three rows show a motion simulated stacks with the motion levels 0.5, 1.0, and 1.5, respectively and an intensity noise level simulating 1.5T data. Right column: The slice rms as a function of the slice-to-mask intersection given the ground truth for three studies with three different motion level. The color of the markers indicate the iteration label of plotted rms error. Red, green, yellow, and blue correspond to the zeroth, first, second, and third iteration of the unified approach, respectively. There is a clear pattern of rms reduction between each iteration in all studies which justifies each iteration level. The plots also reveal that there is a drop in performance for slices with a mask intersection size of less than approximately a 1000 voxels.

Fig. 5.

Selected slice stacks for the four clinical cases showing examples of minimum motion.

Displays are approximately axial and sagittal views created using nearest neighbor interpolation to highlight the underlying voxel size. The upper two rows show nDW slices and the bottom two rows show DW slices. This also highlights the variation in image quality which is dependent on the positioning of the coil and the fetus during the clinical study, and the low level of signal available in the slices motivating the combination of data from different views.

Fig. 6.

Top: FA maps computed from DTI reconstructions of the 110 day study, starting from different simulated motion perturbation levels. The first two rows show axial and sagittal views of the FA maps resulting from the DTI reconstructions without motion estimation while the next two rows show the FA maps resulting from the DTI reconstruction with motion estimation. First column corresponds to the unperturbed SLIMMER initialization, and the remaining five columns correspond to the reconstructions at the different motion simulation levels. From the first two rows it can be observed that the quality of the reconstruction deteriorates with the increased motion level. The only nonmotion-corrected reconstruction with perturbation of the initialization, which shows well-defined white matter structures and cortical diffusion anisotropy, is made from the lowest perturbation level. It is observed that the quality of the motion-corrected reconstructions with perturbations levels

1.00 are comparable, while the reconstructions with motion levels >1.00 exhibit a drop in quality either due to missing data or slice alignment errors. Bottom: FA maps color-coded using the tensor direction for the reconstructions for the macaque fetus overlaid upon structural T2-weighted reconstructions. The FA values are scaled linearly between 0.05 and 0.35. The rows of the figure correspond to the gestational ages 85, 110, and/or 135 days, the first column contains an axial view, the second column contains a sagittal view, and the third column contains a coronal view. The next column displays the coronal view of FA maps of

the reconstructions computed DTI fitting to the first DWI stack assuming no motion highlighting the lower signal to noise, resolution and image artifacts. The final column shows volume only motion estimation and reconstruction from all the data assuming no motion.

Fig. 7.

Cortical plate FA pattern (top) and FACT based ROI tractography of the corpus callosum (bottom) for all three gestational ages constructed using the 0.75 mm resolution reconstruction from all the acquired diffusion weighted slice stacks. Extracted tracts are color-coded with the local tract direction (green anterior–posterior, blue superior–inferior, red medial–lateral).

Fig. 8.

An example of the amount of between slice motion recovered in a typical human clinical study (subject 2) in terms of slice positioning before and after slice alignment. Slices shown are from one diffusion weighted component in each of the three orthogonal diffusion weighted acquisitions using an interleave of 2. Right column shows an enlarged view of the improved alignment of corresponding structures in the diffusion weighted slices from the different diffusion sets collected.

Fig. 9.

FA maps color-coded using the primary tensor direction for all four human clinical studies overlaid on structural T2-weighted reconstructions (all data reconstructed to 1 mm cubic voxels). Left three columns are the iterative reconstruction using full slice motion estimation. The next column is the reconstruction using volume based motion correction [23], the right most column shows the reconstruction with no DW motion estimation from the initial SLIMMER setup. Yellow arrows indicate improved delineation of the small corpus callosum in case 1 and 3 with enlarged ventricles using slice motion estimation. In case 2 the yellow arrows highlight the absence of any left-right diffusion (red) in the region of the corpus callosum.

Fig. 10.

Comparison of connectivity in 1.5T human studies using tractography in 3-D data derived from typical clinical DTI scans shown with planes from underlying T2W MRI reconstructions (all data reconstructed to 1 mm cubic voxels). (a) Comparison of tractography from a right anterior ROI (yellow) in the normal corpus callosum of case 4. Results show improved delineation of white matter connectivity (from the subjects right to the subjects left) after correction for between slice motion, with longer connections, more clustered tract shape and FA values (color coding). (b) Comparison of tractography from the Corpus Callosum in a 3-D reconstruction after motion correction in case 2 with agenesis of the corpus callosum (right), and case 4 with normal corpus callosum (left).

TABLE I
Acquisition Information for the Four Clinical *IN UTERO* Fetal DWI Studies Analyzed.

Case	TE (ms)	TR (ms)	b (s/mm ²)	#dir	#stacks (ax-sag-cor)	Diagnosis	Age (weeks)
1	96	5800	500	10	2-2-1	ACC and VM	32 $\frac{3}{7}$
2	96	6100	600	10	2-2-1	ACC and VM	26 $\frac{4}{7}$
3	96	4100	700	10	2-2-0	Severe VM	22
4	96	4100	700	10	2-2-2	Mild VM	26

ACC = agenesis of Corpus Callosum, VM = Ventriculomegaly

TABLE II
Accuracy of Slice Alignment Estimates From Synthetically Perturbed Stationary Adult Imaging Data

Study		Initialization		NMI				SSD			
Motion	ID	Iter		0	1	2	3	0	1	2	3
		Mtc	Mtc								
0.5	1	MSRMS		0.72	0.50	0.49	<u>0.30</u>	0.96	0.52	0.51	<u>0.30</u>
		MVD		0.58	0.37	0.37	<u>0.20</u>	0.81	0.39	0.39	<u>0.20</u>
		LD/%		0.04	0.03	0.03	0.03	<u>0.00</u>	0.01	0.01	0.01
		PP/p		NA	< 0.01	< 0.01	< 0.01	NA	< 0.01	0.80	< 0.01
		AB/p		NA	< 0.01	0.21	< 0.01	NA	< 0.01	0.46	< 0.01
	2	MSRMS		0.72	0.52	0.52	0.33	0.49	0.52	0.52	<u>0.31</u>
		MVD		0.59	0.40	0.40	0.25	0.38	0.38	0.39	<u>0.21</u>
		LD/%		0.04	0.03	0.03	0.03	0.07	<u>0.00</u>	<u>0.00</u>	<u>0.00</u>
		PP/p		NA	< 0.01	< 0.01	< 0.01	NA	< 0.01	< 0.01	< 0.01
		AB/p		NA	< 0.01	0.45	< 0.01	NA	< 0.01	0.49	< 0.01
0.75	1	MSRMS		1.08	0.80	0.80	0.62	0.83	0.74	0.73	<u>0.44</u>
		MVD		0.88	0.63	0.64	0.48	0.66	0.54	0.53	<u>0.28</u>
		LD/\$		0.07	0.08	0.08	0.08	0.20	<u>0.03</u>	<u>0.03</u>	<u>0.03</u>
		PP/p		NA	< 0.01	> 0.99	< 0.01	NA	< 0.01	< 0.01	< 0.01
		AB/p		NA	< 0.01	0.06	< 0.01	NA	< 0.01	0.03	< 0.01
	2	MSRMS		1.08	0.97	0.98	0.86	1.05	0.74	0.72	<u>0.45</u>
		MVD		0.87	0.83	0.85	0.76	0.93	0.54	0.52	<u>0.29</u>
		LD/%		0.11	0.07	0.07	0.07	0.08	<u>0.01</u>	<u>0.01</u>	0.02
		PP/p		NA	< 0.01	> 0.99	< 0.01	NA	< 0.01	< 0.01	< 0.01
		AB/p		NA	< 0.01	0.65	< 0.01	NA	< 0.01	0.39	< 0.01
1.0	1	MSRMS		1.52	1.36	1.36	1.18	2.06	1.02	0.98	<u>0.62</u>
		MVD		1.22	1.14	1.16	1.03	1.70	1.02	0.98	<u>0.38</u>
		LD/%		0.13	0.09	0.09	0.09	<u>0.03</u>	<u>0.03</u>	0.04	0.06

Study		Initialization		NMI					SSD				
Motion	ID	Iter		0	1	2	3	0	1	2	3		
		Mtc	Mtc										
1.25		PP/p	NA	< 0.01	0.79	< 0.01	NA	< 0.01	< 0.01	< 0.01	< 0.01		
			AB/p	NA	< 0.01	0.29	< 0.01	NA	< 0.01	< 0.01	< 0.01		
		2	MSRMS	1.51	1.47	1.48	1.34	1.44	1.01	0.98	<u>0.62</u>		
			MVD	1.22	1.29	1.32	1.22	1.29	0.70	0.68	<u>0.37</u>		
			LD/%	0.17	0.11	0.11	0.11	0.13	<u>0.09</u>	0.10	0.10		
			PP/p	NA	< 0.01	> 0.99	< 0.01	NA	< 0.01	< 0.01	< 0.01		
	AB/p	NA	< 0.01	0.15	< 0.01	NA	< 0.01	< 0.01	< 0.01				
		1	MSRMS	2.15	1.63	1.59	1.34	2.75	1.36	1.21	<u>0.79</u>		
	MVD		1.73	1.32	1.33	1.13	2.25	0.93	0.87	<u>0.49</u>			
	LD/%		0.36	0.30	0.30	0.30	<u>0.09</u>	0.14	0.13	0.14			
	PP/p		NA	< 0.01	< 0.01	< 0.01	NA	< 0.01	< 0.01	< 0.01			
	AB/p		NA	< 0.01	0.06	< 0.01	NA	0.98	< 0.01	< 0.01			
2			MSRMS	2.05	2.20	2.23	2.10	2.33	1.35	1.23	<u>0.79</u>		
	MVD	1.66	1.95	1.99	1.94	1.99	0.93	0.87	<u>0.47</u>				
	LD/%	0.18	0.13	0.14	0.14	0.26	<u>0.10</u>	0.11	0.12				
	PP/p	NA	< 0.01	> 0.99	< 0.01	NA	< 0.01	< 0.01	< 0.01				
	AB/p	NA	< 0.01	0.02	< 0.01	NA	< 0.01	< 0.01	< 0.01				
		1.50	1	MSRMS	2.97	2.12	2.00	1.73	3.55	1.84	1.60	<u>1.08</u>	
MVD	2.41			1.63	1.58	1.39	2.92	1.27	1.13	<u>0.65</u>			
LD/%	0.38			0.36	0.35	0.41	<u>0.10</u>	0.19	0.22	0.26			
PP/p	NA			< 0.01	< 0.01	< 0.01	NA	< 0.01	< 0.01	< 0.01			
AB/p	NA			< 0.01	< 0.01	< 0.01	NA	< 0.01	< 0.01	< 0.01			
	2			MSRMS	2.92	2.55	2.44	2.23	2.24	1.82	1.55	<u>1.02</u>	
MVD			2.37	2.13	2.08	1.94	1.88	1.25	1.08	<u>0.60</u>			
LD/%			0.42	0.47	0.47	0.48	0.66	<u>0.26</u>	0.28	0.30			
PP/p			NA	< 0.01	< 0.01	< 0.01	NA	< 0.01	< 0.01	< 0.01			

Study		Initialization	NMI					SSD				
Motion	ID	<div>Iter Mtc</div>	0	1	2	3	0	1	2	3	AB/p	
			NA	< 0.01	< 0.01	< 0.01	< 0.01	NA	< 0.01	< 0.01	< 0.01	< 0.01

Precision/Robustness to Motion Level Investigated Using Synthetically Perturbed *In Utero* Sedated Fetal Macaque Imaging

TABLE III

Motion	Initialization		NMI				SSD				
	ID	Iter	0	1	2	3	0	1	2	3	
		Mtc									
0.5	1	MSRMS	0.74	0.46	0.42	0.36	0.84	0.42	0.34	<u>0.29</u>	
		MVD	0.55	0.31	0.29	0.25	0.66	0.35	0.28	<u>0.25</u>	
		LD/%	0.24	0.16	0.11	0.09	0.08	<u>0.01</u>	0.04	0.06	
	2	MSRMS	0.73	0.45	0.41	0.44	0.86	0.39	0.29	<u>0.25</u>	
		MVD	0.55	0.31	0.29	0.25	0.67	0.32	0.24	<u>0.21</u>	
		LD/%	0.18	0.08	0.08	0.07	<u>0.01</u>	<u>0.01</u>	<u>0.01</u>	<u>0.01</u>	
0.75	1	MSRMS	1.29	0.89	0.72	0.64	1.39	0.65	0.50	<u>0.43</u>	
		MVD	0.97	0.63	0.53	0.45	1.10	0.50	0.41	<u>0.36</u>	
		LD/\$	0.67	0.37	0.31	0.29	0.25	0.02	<u>0.01</u>	<u>0.01</u>	
	2	MSRMS	1.22	0.74	0.64	0.60	1.38	0.61	0.45	<u>0.40</u>	
		MVD	0.93	0.50	0.46	0.43	1.10	0.49	0.37	<u>0.34</u>	
		LD/%	0.49	0.33	0.30	0.29	0.30	0.04	<u>0.02</u>	0.11	
1.0	1	MSRMS	2.16	1.41	1.17	1.03	2.20	1.13	0.94	<u>0.74</u>	
		MVD	1.68	1.01	0.86	0.74	1.66	0.77	0.60	<u>0.52</u>	
		LD/%	1.11	0.90	0.82	0.89	0.67	0.14	<u>0.10</u>	0.14	
	2	MSRMS	1.95	1.19	0.98	0.86	2.04	0.95	0.72	<u>0.63</u>	
		MVD	1.22	1.29	1.32	1.22	1.65	0.77	0.60	<u>0.52</u>	
		LD/%	0.82	0.62	0.59	0.73	0.46	<u>0.06</u>	<u>0.06</u>	0.08	
1.25	1	MSRMS	3.04	2.14	1.75	1.58	3.05	1.76	1.40	<u>1.19</u>	
		MVD	2.48	1.63	1.33	1.19	2.52	1.48	1.18	<u>1.01</u>	
		LD/%	0.18	0.14	0.12	0.13	1.29	0.33	<u>0.18</u>	0.20	
	2	MSRMS	2.83	1.91	1.58	1.40	2.77	1.51	1.29	<u>1.01</u>	

Motion	Initialization		NMI				SSD			
	ID	$\frac{\text{Iter}}{\text{Mtc}}$	0	1	2	3	0	1	2	3
1,50	1	MVD	2.26	1.38	1.15	1.01	2.25	1.22	1.07	<u>0.84</u>
		LD %	1.82	1.39	1.20	1.31	1.07	<u>0.12</u>	0.27	0.37
		MSRMS	3.82	2.75	2.35	2.20	3.60	2.18	1.77	<u>1.40</u>
		MVD	3.22	2.16	1.84	1.70	2.98	1.81	1.49	<u>1.14</u>
		LD %	1.86	1.76	1.49	2.02	1.04	0.27	<u>0.18</u>	0.23
		MSRMS	4.01	2.49	2.08	1.80	3.87	2.38	1.92	<u>1.51</u>
	2	MVD	3.40	1.97	1.63	1.40	3.21	1.99	1.63	<u>1.26</u>
		LD %	2.10	1.66	1.57	1.01	1.76	0.48	<u>0.32</u>	0.36