

HAITCH: A Framework for Distortion and Motion Correction in Fetal Multi-Shell Diffusion-Weighted MRI

Haykel Snoussi, Davood Karimi, Onur Afacan, Mustafa Utkur, Ali Gholipour

Boston Children’s Hospital, and Harvard Medical School, Boston, MA 02115 USA.

July 1, 2024

Abstract

Diffusion magnetic resonance imaging (dMRI) is pivotal for probing the microstructure of the rapidly-developing fetal brain. However, fetal motion during scans and its interaction with magnetic field inhomogeneities result in artifacts and data scattering across spatial and angular domains. The effects of those artifacts are more pronounced in high-angular resolution fetal dMRI, where signal-to-noise ratio is very low. Those effects lead to biased estimates and compromise the consistency and reliability of dMRI analysis. This work presents HAITCH, the first and the only publicly available tool to correct and reconstruct multi-shell high-angular resolution fetal dMRI data. HAITCH offers several technical advances that include a blip-reversed dual-echo acquisition for dynamic distortion correction, advanced motion correction for *model-free* and robust reconstruction, optimized multi-shell design for enhanced information capture and increased tolerance to motion, and outlier detection for improved reconstruction fidelity. The framework is open-source, flexible, and can be used to process any type of fetal dMRI data including single-echo or single-shell acquisitions, but is most effective when used with multi-shell multi-echo fetal dMRI data that cannot be processed with any of the existing tools. Validation experiments on real fetal dMRI scans demonstrate significant improvements and accurate correction across diverse fetal ages and motion levels. HAITCH successfully removes artifacts and reconstructs high-fidelity fetal dMRI data suitable for advanced diffusion modeling, including fiber orientation distribution function estimation. These advancements pave the way for more reliable analysis of the fetal brain microstructure and tractography under challenging imaging conditions.

Keywords: Fetal Imaging, Diffusion MRI, Motion Correction, Dynamic Distortion Correction

1 Introduction

Fetal diffusion magnetic resonance imaging (dMRI) provides detailed insights into the development of brain connectivity and microstructure during the prenatal period [1–13]. This powerful technique relies on a mathematical model that delves into the microscopic behavior of water molecules within brain tissue. This model begins by calculating the probability $P(\mathbf{R}_\tau|\mathbf{R}_0, \tau)$, which represents the likelihood of a water molecule moving from position \mathbf{R}_0 to \mathbf{R}_τ over time τ . Since determining this probability for a single molecule is impossible, we calculate the Ensemble Average Propagator (EAP), $P(\mathbf{R})$, which represents the averaged probability across all molecules within a voxel. Building upon Stejskal and Tanner’s work [14] on the pulsed gradient spin-echo sequence, we can express the normalized dMRI signal, $E(\mathbf{q})$, as the Fourier transform of the EAP. When the diffusion-weighted gradient duration δ is sufficiently shorter than the time between two gradient pulses Δ , $E(\mathbf{q})$ can be expressed as:

$$E(\mathbf{q}) = \int_{\mathbf{R} \in \mathbb{R}^3} P(\mathbf{R}) \exp(-2\pi i \mathbf{q} \cdot \mathbf{R}) d\mathbf{R} \quad (1)$$

where \mathbf{q} is a 3D-vector representing the effective gradient direction ($\mathbf{q} = q\mathbf{u}$ where \mathbf{u} is a 3D unit vector). A standard dMRI scan consists of acquiring a reference insensitive to diffusion (b-value = 0) and a set of diffusion-weighted images (b-value > 0) in non-collinear directions. By employing a suitable diffusion model, this setup allows the estimation of the local diffusion properties and the brain’s structural connectivity network. The diffusion model estimation assumes that the gathered

dMRI signals originate from the same physical point. However, achieving this condition in fetal scans is challenging due to several factors such as maternal breathing and unpredictable fetal motion [3, 15]. These issues cause misalignment in dMRI sequences and disrupt the assumption of a consistent relationship between image space and anatomy. Moreover, the inherent limitations of echo planar imaging (EPI) used in dMRI acquisition, such as geometric distortion and susceptibility to spin history effects, distort the dMRI signals [2, 16, 17]. The problem is magnified in the context of fetal imaging due to the complex interaction between fetal motion and the susceptibility-induced inhomogeneities of the main magnetic field (B_0), combined with the small size of the fetal brain and the low signal-to-noise ratio (SNR). Inconsistencies originating from these factors are a major impediment to accurately characterizing fetal brain connectivity and hinder the overall reproducibility of *in-utero* dMRI studies [1, 18].

To address these challenges, several fetal-specific retrospective motion correction methods have been developed. The initial approach of motion correction for fetal dMRI, introduced by [19], extended slice-to-volume reconstruction (SVR) techniques originally developed for *in-utero* structural MRI. By assuming that local diffusion properties can be represented by a rank-2 tensor model, each slice was registered to the simulated volume. A diffusion tensor matrix reconstruction was then employed to integrate the realigned slices into a regular grid to produce the final reconstructed volume. [20] proposed a groupwise registration method that aligns diffusion-weighted images collectively before utilizing a derived image for registration with the T2-weighted (T2w) reference image. Affine transformation matrices were applied to realign the original sequences, addressing both motion and eddy-current effects. Then, a dual radial basis function-based interpolation was used to reconstruct a consistent image. [21] presented a super-resolution reconstruction of the diffusion tensors, with some similarities in the registration concept to [20] and several key distinctions from [22]. Notable differences include the implementation of a unified reconstruction-alignment formulation, yielding a diffusion-sensitive slice registration model. Additionally, Fogtmann *et al.* incorporated point spread function deconvolution into the 3D image reconstruction process. This allowed the creation of 3D datasets with isotropic spatial resolution from multiple scattered slices acquired in different anatomical planes. [23] proposed motion correction and reconstruction through a dynamic model of fetal head motion; where motion parameters are estimated through a Kalman filtering approach. Additionally, they reconstructed diffusion tensors using a weighted least squares fit. This approach notably contributed to the creation of the first spatio-temporal diffusion tensor fetal atlas, as detailed in [6]. [3] advanced the field by correcting motion using super-resolution reconstruction with spherical harmonics (SH). This approach is particularly effective in identifying crossing fibers in the fetal brain. Building upon their previous work [24], this technique notably includes intensity correction within the reconstruction process.

The above-mentioned techniques have predominantly employed data representations grounded in diffusion tensor imaging and single-shell SH. However, these approaches have limitations in capturing the complexity of the developing brain. Tensor model struggles to represent crossing fibers which is a prevalent feature in fetal brain white matter. A *model-free* signal representation is crucial for such motion correction methodologies to guarantee a broad range of imaging analyses. Similarly, [25] highlight the insufficiency of single-shell SH in capturing the full spectrum of diffusion information needed for neonatal imaging, which is highly similar to fetal imaging. Consequently, these methodologies either restrict the type of diffusion information extracted or reduce the information content and impose constraints on the variety of input data that can be processed. This significantly hinders the scope of subsequent analytical endeavors aimed at comprehensively characterizing the developing fetal brain microstructure. Moreover, the effect of local geometric distortions and their interaction with fetal and maternal motion has been either ignored or not adequately addressed in those studies. While conventional methods for EPI distortion correction rely on static field mapping [26], which is only appropriate for motion-free data, methods that rely on dMRI volumes with reversed phase encoding are not effective for the continuous and large motion that typically affects fetal dMRI [27, 28]. This warrants further investigation of dMRI acquisition and processing methods to address motion and geometric distortions.

In this work, we introduce the High Angular resolution diffusion Imaging reconsTruction and Correction approach (HAITCH), a novel framework that tackles these identified challenges by integrating optimized acquisition and reconstruction strategies to mitigate the combined effects of fetal motion, geometric distortion, and their interaction in fetal dMRI. To this end, our objective was to develop a robust methodology and a toolbox for acquiring and reconstructing fetal dMRI data that ensures high fidelity and suitability for detailed *in-vivo* and *in-utero* analysis of the fetal brain microstructure development.

To achieve this, HAITCH focuses on three core technical contributions, represented by the bold-bordered boxes in Fig. 1: *i*) Optimized multi-shell sampling scheme for improved fetal dMRI, *ii*) Dual-Echo sequence enables dynamic distortion correction, *iii*) Advanced motion correction and model-free reconstruction techniques. We validate the accuracy of our framework using real fetal dMRI scans acquired at Boston Children’s Hospital. The implementation of the HAITCH framework is publicly available as the first module of the Fetal and Neonatal Development Imaging (FEDI) toolbox: <https://fedi.readthedocs.io>.

2 Material and methods

HAITCH employs a multi-stage approach to address the challenges associated with fetal dMRI acquisition. Initially, a specialized multi-shell high angular resolution diffusion imaging (HARDI) sampling scheme is designed to increase the overall dataset’s tolerance to motion. Our recommended acquisition involves a modified dual-echo EPI sequence that allows dynamic correction of time-varying geometric distortions. Following acquisition, the data undergoes preprocessing including denoising, Gibbs ringing correction, and Rician bias correction. Geometric distortions are then addressed through non-static field map estimation. Subsequently, B1 bias field correction and fetal brain segmentation are performed. Motion correction is achieved through an iterative refinement process. A flowchart summarizing the various processing stages of HAITCH is shown in Fig. 1. This approach progressively improves the data by repeatedly updating slice weights, transform coefficients, and motion parameters, which will be described in detail in the sections that follow. Through these iterative updates, the self-consistency of the data is enhanced, leading to progressively improved motion-corrected images. In this section, we first explain the multi-shell sampling scheme, then we describe the signal representation that makes the foundation of our modeling, sampling, image reconstruction, outlier detection, slice weighting, basis updating, and motion parameters estimation. At the end, we explain the dynamic distortion correction component based on our recommended multi-echo EPI sequence. The details of the experiments conducted, pre-processing, implementation aspects of our methods, and post-processing are explained in Section 3. Section 4 will then present the results obtained.

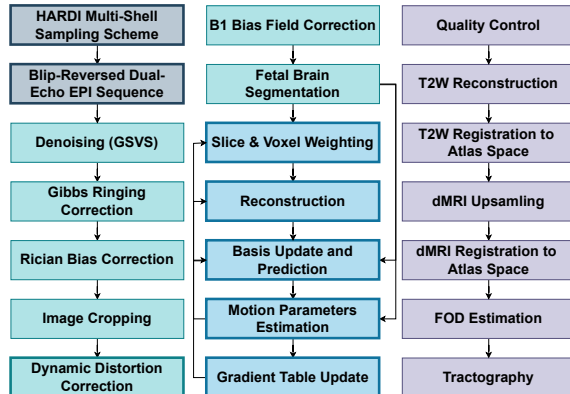


Figure 1: HAITCH Framework Processing Stages. This flowchart illustrates the key steps involved in the HAITCH Framework. Boxes with bold borders refer to our main contributions. Grey boxes outline the acquisition stages, including the sampling scheme and the dual-echo sequence. Green boxes detail the pre-processing steps, distortion correction, B1 bias field correction, and fetal brain segmentation. Blue boxes refer to the iterative approach of motion correction and image reconstruction that includes slice & voxel weighting, *model-free* reconstruction, signal basis update, signal prediction, registration, and updating the gradient table. Purple boxes showcase the post-processing steps, encompassing T2-weighted image reconstruction, spatial normalization for atlas registration, and streamlined tractography. Refer to Section 2 for details on the sampling scheme, dual-echo sequence, motion correction, and dynamic distortion correction. Section 3 discusses pre-processing, post-processing, and implementation specifics. Boxes with bold borders refer to our main contributions.

2.1 A Multi-Shell HARDI Sampling Scheme for the Fetal Brain

dmMRI with multi-shell HARDI provides rich information about the tissue microstructure by varying b-values and directional sampling. Utilizing an extension of the electrostatic repulsion cost function [29,30], we have customized a sequence with HARDI, interlaced multi-shell, and an incremental sampling scheme for uniform angular coverage. This sampling scheme ensures a balance between the global angular distribution of all diffusion gradient directions and the angular distribution on each shell. The angular density of coverage for each shell increases uniformly as fetal dmMRI acquisition proceeds. This optimization of the angular coverage and the temporal ordering aims to mitigate the effects of both sudden and slow fetal motion during acquisition. Acquired slices and volumes with close gradient directions are separated in acquisition time, thereby increasing the motion tolerance of the entire dataset. With this acquisition scheme, any part of the acquired data will contain uniformly distributed orientations, which is not possible with conventional schemes.

Furthermore, to maximize the information content of the data within our multi-shell HARDI acquisition scheme, we determined the per-shell sampling density and the associated b-values based on the detectable number of SH, as detailed below. Previous studies have shown that the angular frequency in neonates is significantly lower compared to adults [25, 31]. [25] demonstrated in a neonatal study that SH terms beyond order 8 are indistinguishable from noise for typical SNR values. SH terms of order 6 were small, suggesting a minimum requirement of 28 directions at high b-values. At lower b-values, only terms of orders 2 or 4 could be detected, implying minimum requirements of 6 and 15 directions, respectively. In practice, acquiring more than the minimum diffusion-weighted directions is highly recommended to account for potential imperfections in the diffusion-encoding gradient directions and ensure sufficient effective SNR for utilizing SH terms with order 6. Additionally, [25] determined that $[0, 400 \text{ s/mm}^2, 1800 \text{ s/mm}^2]$ and $[0, 400 \text{ s/mm}^2, 1000 \text{ s/mm}^2, 2600 \text{ s/mm}^2]$ b-value combinations as the optimal for 2-shells and 3-shells sampling schemes in their neonatal imaging study. Their corresponding number of directions expressed as a percentage of the total number of volumes acquired are as follows: [12%, 29%, 59%] and [6%, 19%, 28%, 47%], respectively.

Considering the potential differences and similarities between fetal and neonatal imaging, the dual-echo nature of our sequence, and achievable b-values due to inherent low SNR in fetal dmMRI [32], we propose a 2-shell HARDI scheme summarized in Table 1. This scheme incorporates 11 interleaved $b = 0$ images, and diffusion-weighted images at $b = 400 \text{ s/mm}^2$ and $b = 900 \text{ s/mm}^2$ with 28 and 56 directions per shell, respectively.

b-value (s/mm ²)	0	400	900
Number of Directions	11	28	56
Percentage of Total Volumes	12%	29%	59%

Table 1: Summary of the scheme: b-values and directions per shell

2.2 Signal Representation

2.2.1 Diffusion Signal Modeling

To effectively capture the orientation-dependent nature of dmMRI contrast, we employ the 3D Simple Harmonic Oscillator Reconstruction and Estimation (SHORE) technique [33, 34]. SHORE serves as a basis function representation that accurately models both the radial and angular properties of the dmMRI signal by fitting a linear combination of orthogonal basis functions. SHORE allows us to represent the diffusion signal $E(\mathbf{q}\mathbf{u})$ as a truncated linear combination of orthonormal basis functions $\Phi_{nlm}(q, \mathbf{u})$:

$$E(\mathbf{q}\mathbf{u}) = \sum_{l=0, \text{even}}^{N_{\max}} \sum_{n=l}^{(N_{\max}+l)/2} \sum_{m=-l}^l \mathbf{c}_{nlm} \Phi_{nlm}(q, \mathbf{u}) \quad (2)$$

where the $\mathbf{c}_{nlm} = \langle E, \Phi_{nlm} \rangle$ are the transform coefficients that capture the contribution of each basis function, $\Phi_{nlm}(\mathbf{q}, \mathbf{u}) = X_{nl}(q, \zeta) Y_l^m(\mathbf{u})$ can be decomposed into (i) $X_{nl}(q, \zeta)$, the radial basis, is a generalized Laguerre polynomial modulated by exponential decay and scaling factor ζ , capturing the radial behavior and ensuring orthonormality. n represents the radial order. (ii) $Y_l^m(\mathbf{u})$, the angular basis, is a SH function of order ℓ and degree m , modeling the angular properties, N_{\max} is the maximal order of the functions in the truncated series. The sparsity inherent in the SHORE model, characterized by the concentration of significant information within a few coefficients, is crucial for effective

diffusion signal modeling [34]. Furthermore, the ability of SHORE to accurately detect multiple diffusion directions has been demonstrated in previous studies [35]. This representation translates the fetal dMRI scan into a compact 4D representation of SHORE coefficients associated with each voxel in the image grid. With this diffusion signal representation and motion parameters, we can predict the expected dMRI contrast using a forward model. This prediction allows us to compare it with the motion-corrupted acquired data. Subsequently, we formulate an inverse problem that optimizes the similarity between the predicted and acquired data while simultaneously estimating the fetal motion traces.

2.2.2 Forward Modeling

Fetal motion during dMRI acquisition disrupts the signal, introduces artifacts, and degrades the acquired data consistency. By assuming rigid motion, we can model this disruption using a forward model. This model allows us to predict the expected dMRI signal in the presence of motion. This involves transformations in the q -space domain to account for changes in spatial and angular coordinates due to motion. It essentially simulates how diffusion would appear under the influence of motion, providing a mathematical relationship between the acquired dMRI data, the underlying true signal, and the motion parameters. For a given set of fetal motion parameters $\boldsymbol{\mu}$, the predicted signal $\hat{\mathbf{E}}$ could be expressed as:

$$\hat{\mathbf{E}}(\boldsymbol{\mu}, \mathbf{q}) = \mathbf{C}(\mathbf{q}) \mathbf{M}(\boldsymbol{\mu}) \mathbf{D}(\boldsymbol{\mu}) \mathbf{A} \mathbf{c} \quad (3)$$

where \mathbf{A} represents noise, \mathbf{D} accounts for EPI distortion, $\mathbf{M}(\boldsymbol{\mu})$ represents the effects of rigid motion parameters on the signal, and $\mathbf{C}(\mathbf{q})$ encodes the SHORE basis for q -space. This equation essentially enables generating predictions for the signal in a scattered collection, each with their own motion state and dMRI encoding.

2.3 Inverse Problem and Reconstruction

Given the scattered data acquired during the dMRI scan, we aim to find the optimal reconstruction coefficients \mathbf{c}^* and rigid motion parameters $\boldsymbol{\mu}^*$ that maximize the similarity between the acquired data and their prediction. We achieve this by solving an inverse problem. This involves finding the combination of reconstruction coefficients and motion parameters that best explain the acquired data. Mathematically, this is formulated as a large, sparse least-squares problem with an ℓ_2 regularization:

$$\mathbf{c}^{(k)} = \arg \min_{\mathbf{c} \in \mathbb{R}^{n_c}} \left\| W \left(\hat{\mathbf{E}}^{(k)}(\boldsymbol{\mu}^{(k)}, \mathbf{q}) - \mathbf{E} \right) \right\|_{\ell_2} + \lambda_l \|Lc\|_{\ell_2} + \lambda_n \|Nc\|_{\ell_2} \quad (4)$$

where $\hat{\mathbf{E}}^{(k)}$ is the predicted signal of all scattered data in the current iteration k using the coefficients $\mathbf{c}^{(k)}$, $W = \text{diag}(\dots, \sqrt{w_s}, \dots)$, w_s is a slice-wise or voxel-wise weight for outlier rejection. The weighting matrix W assigns higher importance to reliable slices or voxels and reduces the influence of outliers. N and $L \in \mathbb{R}^{n_q \times n_c}$ are two diagonal matrices with $\text{diag}(N) = n(n+1)$ and $\text{diag}(L) = l(l+1)$, respectively. These matrices penalize the high frequencies of the radial and angular parts, respectively. The constants λ_n and λ_l are weights of the penalty terms.

Equation 4 minimizes a cost function that measures a weighted difference between the predicted signal $\hat{\mathbf{E}}^{(k)}$ and the acquired signal \mathbf{E} . We formulate the solution of this least squares problem as an iterative process that leverages a data-driven signal representation. We focus on optimizing the target image for maximum similarity to the acquired scattered data, while simultaneously identifying the unknown fetal motion traces that correspond to the fetal brain position. We utilize surrounding directional information to compensate for potential data loss and reconstruct scattered data across spatial and angular domains. The optimization process consists of repeatedly alternating between slice weighting, a reconstruction step optimizing the SHORE coefficients \mathbf{c} given current motion parameters $\boldsymbol{\mu}$, and a registration step updating $\boldsymbol{\mu}$ for the current \mathbf{c} throughout each epoch of our framework.

2.4 Outlier Detection and Weighting

Fetal dMRI is inherently susceptible to outliers and imaging irregularities, which can negatively impact image quality and the subsequent processing steps. Outliers are signal intensities that significantly deviate from the expected signal based on other measurements. Fetal motion and maternal breathing during diffusion encoding can induce severe signal dropouts. Moreover, rapid slice acquisition times combined with fetal motion can lead to spin-history artifacts, manifesting as signal

hyper-intensities due to overlapping slice excitation. Hence, detecting and appropriately weighting outliers are essential to mitigate the impact of unreliable slices on the reconstruction. To examine the likelihood of a slice S_i or voxel V_i with gradient g being an outlier, we employ three complementary methods.

2.4.1 Modified Z-Score

We employ a modified z -score and median absolute deviation (MAD) to detect outliers within each b-value shell. MAD focuses on the median, which is less swayed by outliers, and is a better choice than the standard deviation for outlier detection in datasets that might have outliers. Given a fetal dMRI dataset, let $\mathbf{E}[\hat{S}_{i,b,g}]$ represent the diffusion signal intensity average for the i -th slice with b-value b and gradient direction g , and let $\mathbf{E}[\widetilde{S}_{i,b}]$ present the median of those average across the same b-value b . We compute the modified z -score for slice i with gradient direction g as follows:

$$Z_{i,g} = \frac{|\mathbf{E}[\hat{S}_{i,b,g}] - \mathbf{E}[\widetilde{S}_{i,b}]|}{1.4826 \cdot \text{MAD}(\mathbf{E}[\hat{S}_{i,b,g}] - \mathbf{E}[\widetilde{S}_{i,b}])} \quad (5)$$

Subsequently, lower η_l and upper η_u thresholds for the modified z -scores are applied to assign a weight w_s to each slice as following:

$$w_s = 1 - \frac{Z_{i,g} - \eta_l}{\eta_u - \eta_l} \quad (6)$$

These weights are incorporated within the initialization epoch to influence the initial predicted signal $\hat{\mathbf{E}}^{(1)}$.

2.4.2 Bayesian Gaussian Mixture Modeling

The second method utilizes a probabilistic criterion to identify outliers based on the intensity difference between the acquired dMRI data and their corresponding predictions [36]. This approach identifies and excludes outlier slices exhibiting large residuals compared to other slices with similar b-values. We first calculate the root-mean-squared error (RMSE), denoted by ϵ , between the acquired data and the provided signal predictions. The RMSE values for inlier and outlier slices are expected to follow different probability distributions. Within each shell, a two-component Bayesian Gaussian Mixture Model (GMM) is employed to model the log-transformed RMSE values, thereby delineating inlier and outlier slices. The GMM assigns a probability, considered as w_s , to each slice belonging to the inlier component.

2.4.3 Voxel-Wise Weighting

The previous two methods focused on slice-level weighting to address signal dropout caused by bulk motion during acquisition. However, local variations within slices can arise due to factors like spin history and physiological motion. To down-weight these localized outliers, we employ a voxel-wise weighting applied in the final processing epoch. This method builds upon the voxel-wise standardized residuals with the modified z -score which are then converted to weights for each voxel as follows:

$$w_v = \frac{1}{(z_v^2 + 1)^2} \text{ where } z_v = \frac{E_v - \hat{E}_v}{1.4826 \cdot \text{MAD}(E_v - \hat{E}_v)} \quad (7)$$

Here, E_v and \hat{E}_v represent the acquired and predicted signal intensity for a particular voxel, respectively.

2.5 Registration

In this step, we seek to compute and update the motion parameters by optimizing the spatial alignment between the predicted signal and the acquired scattered data. We achieve this by performing rigid image registration directly between the predicted and the acquired dMRI contrasts at each time point. This approach offers several advantages over conventional registration methods that often use the b0 image (unweighted image) as the reference for aligning the acquired dMRI data. Here, the predicted dMRI contrasts at each time point serve as the reference for registration with the acquired dMRI data. This way of registration is more accurate than the conventional registration to b0 images, as the fixed (predicted) and moving (acquired) images have the same diffusion sensitization.

2.6 Dynamic Distortion Correction

Fetal dMRI often relies on EPI for its efficient data acquisition. However, EPI inherits a limitation: the presence of local geometric distortions. These distortions are a result of the inhomogeneities of the main magnetic field B_0 , which are caused by the differences in magnetic susceptibility of tissue and air/gas in air/gas-tissue interfaces. A common approach to correct geometric distortions in adult dMRI involves acquiring scans with a single phase encoding direction and adding a $b=0$ image with reversed phase encoding. A field map, assumed to be static, is then generated from the data and used for correcting distortions [27]. While effective for subjects with minimal motion, this method is inadequate for fetal imaging because fetal motion affects field inhomogeneities. The nearly-continuous fetal movements as well as gas movement in the maternal intestine, and maternal respiratory motion, can cause significant changes in B_0 inhomogeneity. These changes would render static field map correction techniques inaccurate in fetal dMRI [2, 16, 37, 38].

To address this challenge, we propose using a modified spin-echo EPI sequence incorporating a second readout (echo). This second echo generates data that have the same resolution, field of view, and bandwidth of the first echo data but with the order of k -space traversal reversed in the phase encoding direction. Fig. 2 illustrates the pulse sequence diagram. A 180° refocusing pulse and spoiler gradients were employed to generate the second spin-echo. The frequency encoding direction is reversed between the two echoes to ensure consistent temporal spacing between k -space points. Notably, the dual-echo sequence maintains the SNR efficiency comparable to the conventional EPI scan despite the extended readout train. The temporal proximity between the two echoes (< 50 ms) ensures closely matched and nominally identical motion states for the two acquired slices. Importantly, the first and second echo-time data exhibit geometric distortions in opposite directions at each acquisition time point. Fig. 2 further illustrates this concept, where susceptibility-induced stretching in the first echo corresponds to signal pile-up in the second echo. This enables the application of the blip-reversed approach in dynamic environments such as fetal and neonatal imaging. This key advantage allows us to estimate a dynamic field map, thereby correcting for geometric distortions in the presence of motion. Implementation details are discussed in Section 3.

3 Implementation and Experiments

3.1 Fetal Datasets

To examine the efficacy of our proposed framework, we conducted 36 dMRI studies on 14 pregnant volunteers, scanned at a gestational age between 24 and 36 weeks. The studies were approved by the institutional review board, and written informed consent was obtained from all participants. Each scan session involved repeated single-shot fast spin echo T2w acquisitions, which were used to reconstruct a reference anatomical volume with an isotropic resolution of 0.8 mm^3 using SVR [39, 40].

Each scan session includes also a dMRI scan acquired in three different ways for testing, (i) *Dual-Echo data*: Twenty-four dMRI scans with two echoes were acquired from 13 subjects. Echo times were 50 ms for TE1 and between 85 ms to 90 ms for TE2. These scans included various single and double-shell b-values (250, 400, 500, 600, 750, 800, 900, and 1000 s/mm^2) across 6-direction sets (13, 15, 20, 22, 26, 27, 44, 80 and 93 directions). (ii) *Multi-Shell HARDI data*: Nine multi-shell HARDI dMRI scans were acquired from five volunteers. These scans were acquired following our optimized scheme described in subsection 2.1. (iii) *Dual-Echo Multi-Shell HARDI data*: An additional three scans from two volunteers possess also the sampling scheme described in subsection 2.1.

All scans are acquired using a 3T Siemens Prisma scanner with a 30-element body coil. Scans were performed with parallel imaging with GRAPPA factor 2, minimum repetition times (TR) between 2.5 and 7 seconds, isotropic resolution settings of 2.0, 2.3, or 2.5 mm^3 , a field of view of 256 or 300 mm depending on the age and size of the fetus and the maternal body, and between 30-40 slices to cover the fetal brain.

3.2 Pre-processing

dMRI data is susceptible to various artifacts arising from the limitations of the MRI hardware and techniques. To address these artifacts and ensure reliable estimation of diffusion parameters, we have established the following steps, as represented by green boxes in Fig. 1:

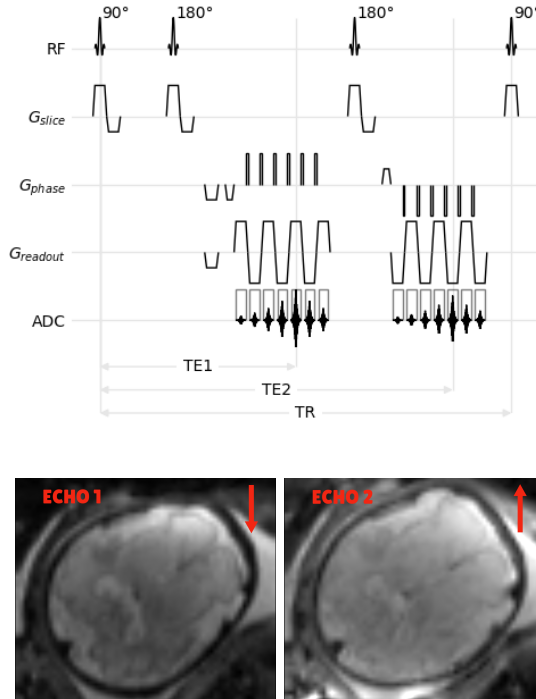


Figure 2: (Top): Schematic diagram of the modified spin-echo EPI sequence. This sequence acquires two echoes (TE1 and TE2) with reversed phase encoding directions. Despite the extended readout, the dual-echo sequence maintains signal-to-noise ratio (SNR) efficiency comparable to a conventional EPI scan. Temporal proximity between the echoes (≤ 50 ms) ensures minimal motion artifacts between acquisitions. (Bottom): Illustration of the opposing susceptibility-induced distortions in the two echoes for an example. The first echo (TE1) exhibits stretching (arrow), while the second echo (TE2) shows signal pile-up in the corresponding location. This key feature allows for dynamic field map estimation and correction of geometric distortions even in the presence of fetal motion. Details of the implementation are discussed in Section 3.

3.2.1 Denoising

The data is initially denoised using the generalized singular value shrinkage (GSVS) method for noise estimation and reduction, as detailed in [41]. This method extends the random matrix theory-based approach of the Marchenko-Pastur Principal Component Analysis (MPPCA), introduced in [42, 43]. GSVS moves beyond the limitations of MPPCA, including the homoscedastic and uncorrelated noise assumptions and the hard separability of signal and noise. Instead, it utilizes the findings of [44] for optimal singular value shrinkage under a generalized Marchenko-Pastur law. This makes GSVS particularly suitable for motion-degraded and low SNR data, including fetal dMRI data with high motion artifacts. GSVS effectively removes only thermal noise without compromising anatomical features in approximately 2 minutes.

3.2.2 Gibbs Ringing Artifact Correction

High-contrast boundaries, such as the edges between cerebrospinal fluid and gray matter or white matter in the fetal brain, may produce image artifacts known as Gibbs ringing. These artifacts arise from inadequate sampling of high frequencies, and may significantly impact the diffusion signal. To remove these artifacts, we employ a method that modulates the truncation in k-space as a convolution with a sinc-function in image space and interpolates the image through local subvoxel shifts [45].

3.2.3 Rician Bias Correction

dMRI images exhibit Rician or non-central χ noise, especially in low-SNR settings like multi-shell HARDI fetal dMRI. Unlike Gaussian noise, Rician noise is non-additive and intensity-dependent, leading to Rician bias and reduced image contrast. To address this, we adapt the methodology introduced in [46] with GSVS replacing MPPCA for Rician bias correction. GSVS denoising provides

an unbiased estimate of the noise standard deviation σ at each voxel based on data acquired with low b-values [41, 42]. We then estimate the true signal intensity E_c through the following equation:

$$E_c^2 = \langle M \rangle^2 - (\xi(\theta) - 2) \sigma^2 \quad (8)$$

Here, $\langle M \rangle$ represents the expected value of the measured magnitude signal intensity, and $\xi(\theta)$ is a correction factor defined by $\theta = \frac{\eta}{\sigma}$, which is equivalent to the SNR [47]. For SNR exceeding 2 dB, $\xi(\theta)$ approaches 1, allowing for accurate signal intensity estimation.

3.3 Dynamic Distortion Correction

Both the first echo time dMRI(TE1) scan and the second echo time dMRI(TE2) scan were used to correct the distortion following the reversed gradient polarity technique. To ascertain the sequence’s efficiency and effectiveness, distortion correction was approached via two dynamic and one static correction methods: (i) Slice-wise correction: field map is estimated for each slice using the reversed echo time slice at every time point for low and high b-values images using the method introduced in Voss *et al.* [48]. (ii) Volume-wise correction: field map is estimated for each volume using the reversed echo time volume at every time point for both low and high b-values images using FSL TOPUP [27]. (iii) Static map correction: the classic way of distortion correction with FSL TOPUP, employing the first b=0 image of TE1 and the last b=0 image of TE2 to calculate the static field map, which is then applied to the entire TE1 dataset. These three correction methods were applied reciprocally between dMRI(TE1) and dMRI(TE2), with outcomes evaluated via the Structural Similarity Index Measure (SSIM) and Peak signal-to-noise ratio (PSNR) metrics. The method with the highest SSIM and PSNR was chosen for subsequent processing.

3.4 B1 Bias Field Correction

Fetal dMRI data acquired with high-density coils suffers from signal variations due to B1 field inhomogeneity. This is particularly problematic in multi-shell acquisitions where complex anatomy and varying diffusion weightings magnify these effects. To address this, we employ the N4ITK algorithm [49] to estimate the corrected signal $E_c(\mathbf{q})$ as follows:

$$E_c(\mathbf{q}) = E_{DC}(\mathbf{q}) \cdot \frac{\hat{E}_{DC(N4)}(\mathbf{b} = \mathbf{0})}{\beta \cdot \hat{E}_{DC}(\mathbf{b} = \mathbf{0})} \quad (9)$$

where $E_{DC}(\mathbf{q})$ is the distortion corrected dMRI signal, β is the bias field estimated using the N4ITK from the masked average of all b=0 images, $\hat{E}_{DC}(\mathbf{b} = \mathbf{0})$ and $\hat{E}_{DC(N4)}(\mathbf{b} = \mathbf{0})$ are the estimated average intensity of b=0 images before and after applying the N4ITK bias correction respectively.

3.5 Fetal Brain Segmentation

Fetal dMRI scans encompass not only the fetal brain but also surrounding maternal structures. Efficient background removal is essential for enhancing the reliability of subsequent processing steps. However, fetal brain extraction is challenging due to factors like variable fetal head position, movement during scans, and the diverse appearance of the developing fetal brain with adjacent anatomy. To address these complexities, we leverage a deep learning-based method built upon a fully convolutional neural network architecture similar to U-Net++ [50]. The model was trained following the established strategy recommended by nnU-Net [51] and has been extensively validated on manually labeled data. This approach effectively removes background structures and performs image cropping to reduce processing time.

3.6 Motion Correction and Reconstruction

The HAITCH framework employs an iterative motion correction and reconstruction process that spans five epochs. Each epoch alternates between outlier reweighting, SHORE coefficients computation, basis updating, signal prediction, registration, and gradient table rotation steps.

The process begins with slice weighting using the modified Z-score method, which does not require a pre-existing predicted signal. The reconstruction step then estimates the coefficients \mathbf{c} by minimizing the cost function of Eq. 4. The computation of these coefficients involves the use of the SHORE basis of order 6 defined by Merlet and Deriche [34]. It is important to note that this basis differs from the one available in the Dipy library [52], resulting in more number of basis functions

(72 in Merlet & Deriche vs. 50 in Dipy for radial order 6). The scale parameter ζ in Eq. 2 is fixed to a value of 700, and the regularization parameters λ_l and λ_n in Eq. 4 are set to 10^{-8} as recommended by [34]. Using these coefficients and the pre-defined sampling scheme, we predict the dMRI signal which, in turn, is used as a reference for the registration step. Rigid registration is performed using ANTs [53] to compute motion parameters $\boldsymbol{\mu}$ from the acquired data. This involves volume-to-volume registration, where the fixed and moving images share the same diffusion contrast. A global correlation metric is utilized to guide the registration process. Based on the estimated motion parameters, the gradient directions are subsequently rotated individually. A new epoch commences with slice weighting using the GMM method, which compares the registered data with the predicted signal. The reconstruction step then updates the SHORE basis, estimates new coefficients, and predicts a new dMRI signal in the initial basis functions for the next registration step. Therefore, motion parameters $\boldsymbol{\mu}$ are updated and refined over four epochs. The final epoch utilizes these refined parameters to obtain final registered data, followed by a voxel-wise outlier detection. Subsequently, SHORE coefficients are computed in the updated basis, and the process concludes by predicting the final dMRI signal in the original sampling scheme.

3.7 Spatial Normalization to Atlas Space

Following the motion correction and reconstruction steps, the HAITCH framework performs spatial normalization to a common reference space. This process involves two stages of registration. First, the motion-corrected and reconstructed dMRI data of the subject is registered to a reconstructed T2w image of the same subject in the world coordinates. The T2w image is obtained using a super-resolution SVR algorithm, either SVRTK [39] or NiftyMIC [40]. Second, the T2w image is registered to the publicly available CRL fetal brain atlas [54]. This atlas provides a standardized anatomical reference for the fetal brain. Finally, the two transformation matrices obtained from these separate registrations are combined into a single transformation, which is then applied to the dMRI data to map it into the atlas space while upsampling it to an isotropic voxel resolution of 1.25 mm. These upsampled dMRI data are used in subsequent post-processing steps including fiber tractography.

3.8 Diffusion Model Estimation and Tractography

White matter tractography was conducted utilizing the MRtrix3 toolbox [55]. Initially, white matter, gray matter, and cerebrospinal fluid response functions were estimated using the unsupervised method. Subsequently, multi-shell multi-tissue constrained spherical deconvolution was performed to estimate fiber orientation distribution (FOD) maps [56]. Streamline tractography was then conducted using the iFOD2 algorithm, a probabilistic tracking method that utilizes second-order integration of the FOD maps [57]. Tracking parameters were carefully chosen to optimize the tractography and generate accurate white matter pathways from the fetal FOD maps. These included an angle threshold of 45° , a cutoff value of 0.01, a specified number of streamlines (100,000), and an FOD power raised to 6 for enhanced specificity. Additionally, minimum and maximum track lengths were set to 10 mm and 120 mm, respectively, to target relevant white matter tracts.

3.9 HAITCH: Open-Source Framework and Flexibility

As mentioned earlier, HAITCH framework is publicly available as the first module of our FEDI toolbox, accessible at <https://fedi.readthedocs.io>, and <https://github.com/FEDIToolbox>. This toolbox includes both the HAITCH implementation and our optimized sampling scheme for fetal dMRI, making these resources readily available to the research community. HAITCH demonstrates its versatility by effectively handling various fetal dMRI data types, including single-shell, multi-shell, or single-echo acquisitions. It further empowers users with a range of options for customization. These include additional outlier detection methods and alternative registration methods and metrics, allowing researchers to tailor the framework to their specific needs. For dynamic distortion correction, besides the methods mentioned previously, users can easily switch to alternative approaches like Block-Matching [58] or EPIC [59] methods.

To take advantage of the full capacity of HAITCH, we recommend using a blip-reversed dual-echo EPI sequence. We have proactively shared our in-house built dual-echo spin echo EPI sequence (developed by co-authors OA and MO) on the Siemens R2C platform. This ensures that researchers will have all the necessary tools readily available for successful HAITCH implementation and application in their fetal brain imaging studies.




Table 2: Comparison of Fetal dMRI Distortion Correction Techniques. This figure illustrates the impact of different distortion correction techniques on fetal dMRI data. The top row displays dMRI (TE1) and the bottom row shows dMRI (TE2) images. Left and right panels represent axial and coronal views, respectively. Each column showcases the results for a different dataset or correction technique: Raw Data: Uncorrected image exhibiting distortions due to susceptibility variations. Slice-wise Correction: Distortion correction is applied to individual slices independently. Volume-wise Correction: Distortion correction is simultaneously applied to the entire volume (HAITCH approach). Static Correction: Correction based on a static field map, potentially inaccurate due to fetal motion.

4 Results

4.1 Distortion Correction: Dynamic Field Map Estimation

Field map estimation utilizing both dynamic and static echo reversal correction approaches was conducted on 27 dMRI scans. The effectiveness of each correction method was evaluated visually and quantitatively. Fig. 2 presents an illustrative example of distortion correction achieved for each method. Overall, visual inspection shows that volume-wise dynamic correction is superior in recovering anatomical brain shape compared to the raw data and other correction techniques. Corrected dMRI (TE1) and dMRI (TE2) appear to have greater visual similarity than those processed with other methods. To quantify these observations, Fig. 3 illustrates the SSIM and PSNR statistics for both $b=0$ and $b > 0$ images across all slices for all scans. The analysis reveals that both dynamic field map estimation methods have significantly enhanced the SSIM and PSNR values compared to the raw data (uncorrected). The volume-wise approach yielded the highest SSIM and PSNR, indicating superior performance in correcting geometric distortions. Conversely, static field map-based correction resulted in lower SSIM and approximately similar PSNR compared to raw data. This highlights its ineffectiveness in addressing motion-induced field inhomogeneities.

4.2 Evaluation of the Motion Correction

Fig. 4 showcases the effectiveness of the motion correction stage of our framework by comparing motion-corrupted fetal dMRI data (first and third columns) with their corresponding motion-corrected reconstructions (second and fourth columns) for two subjects *A* and *B*. The raw data exhibits significant motion artifacts, including signal dropouts and inconsistencies in head orientation across multiple volumes. In contrast, the motion-corrected images in Fig. 4 (second and fourth columns) show dramatic improvements. Signal intensity is successfully recovered, revealing clear visualization of the fetal brain anatomy across all three orthogonal views (axial, coronal, and sagittal). This visual comparison emphasizes HAITCH’s capability to effectively mitigate motion artifacts, and leads to improved data quality suitable for further analysis.

To gain deeper insights into the motion correction process, Fig. 5 depicts the estimated motion parameters μ for subject *B*. The observed translations range from -12 mm to 20 mm, and rotations range from -24 to 10 degrees. The motion profile reveals a slowly varying baseline started by one extended period and punctuated by three shorter bursts of rapid motion. These transient motion events correspond to regions with low slice weights in the lower panel of Fig. 5, as identified by Z-score. This alignment highlights the efficacy of the HAITCH motion correction stage, as it successfully accounts for inter-volume head motion inconsistencies, corrects for intra-volume motion artifacts, and ultimately recovers signal intensity and improves slice uniformity.

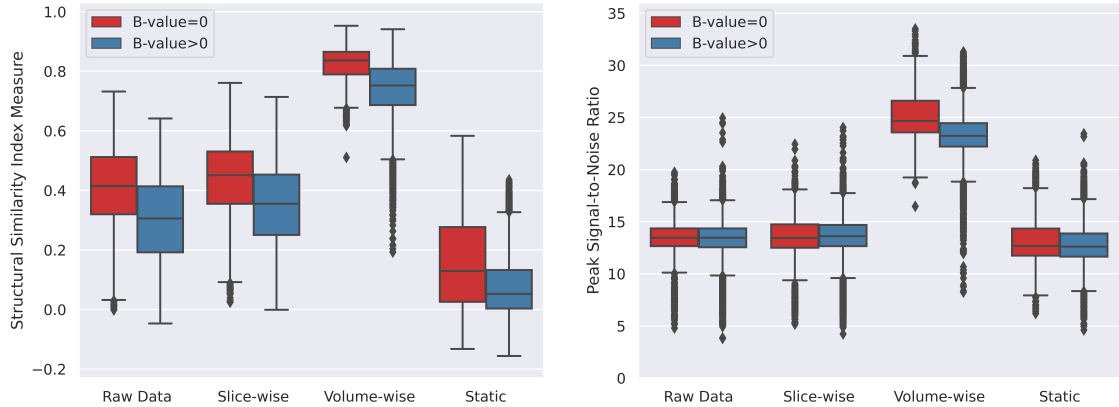


Figure 3: Quantitative Evaluation of Distortion Correction Techniques using SSIM and PSNR. Boxplots summarize the distribution of SSIM (left) and PSNR (right) values across 27 subjects. Results are presented for raw data and data processed with: Slice-wise correction, Volume-wise correction (HAITCH approach), and Static field map correction obtained from the three echo reversal correction methods. Red boxes represent $b=0$ images, while blue boxes display diffusion-weighted images ($b > 0$). Higher SSIM and PSNR values indicate better image quality and reduced distortion. As evident from the plots, the volume-wise dynamic correction method achieved the best results across both metrics and image types.




Figure 4: Example fetal dMRI scans before and after motion correction. The left two columns display axial, coronal, and sagittal views of the raw data (pre-processed up to B1 bias field correction step) (Subject A) and corresponding motion-corrected data, respectively. Each row represents a specific volume (index: 5, 8, 15, 18, 27, 30, 31 from top to bottom). The right two columns show similar data for Subject B. Each row represents a specific volume (index: 7, 37, 38, 50, 58, 60, 87 from top to bottom).

Fig. 6 showcases HAITCH performance even in challenging scenarios, such as low SNR data exemplified by subject C. The figure displays three volumes with close temporal proximity and various brain orientations. Despite this, our framework successfully recovers anatomical details. The well-defined masked brain in the right panel of Fig. 6 also demonstrates the robustness of our deep

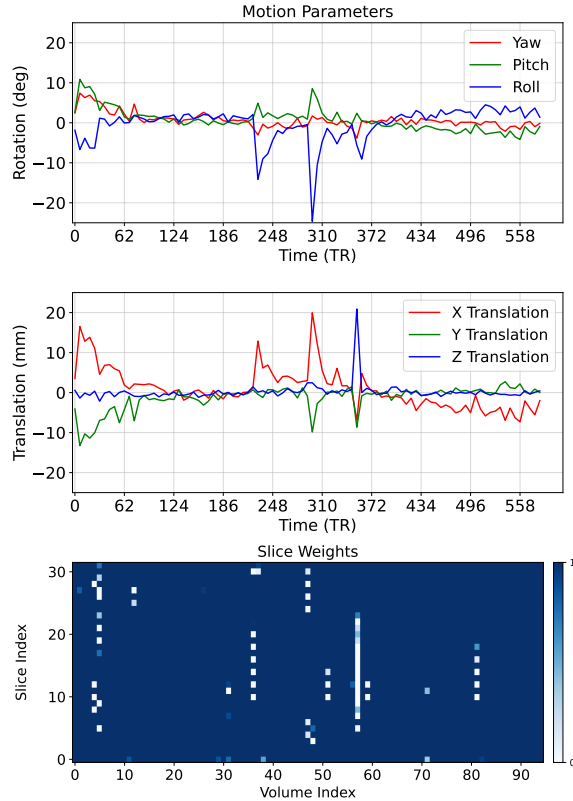


Figure 5: Estimated motion parameters over time and slice weights for Subject B: The top two panels show the estimated motion parameters: translation (mm) and rotation (Euler angles). The bottom panel displays the slice weights calculated using the modified Z -score method. Peaks in motion parameters coincide with low slice weights (outliers).

learning model in fetal brain segmentation. Note that segmentation can be performed before or after motion correction depending on the registration stage which can be sometimes better when keeping the background.

Furthermore, we employed quantitative analysis to complement the visual assessment. Fig. 7 presents the distribution of SSIM values for raw and motion-corrected data across slices between the first and last volumes (b -value = 900 s/mm^2) for eight subjects. The initial and final volumes represent a significant temporal gap during the scan, making them particularly susceptible to intra-scan motion. Our analysis revealed a significant increase in SSIM values after motion correction, which indicates enhanced data integrity and reduced artifacts. This quantitative finding reinforces the visual observations of improved anatomical detail and strengthens the overall validation of the HAITCH framework.




Figure 6: Fetal dMRI with Motion Correction in Challenging Scenario (Subject C). The left column shows raw data from three volumes with close temporal proximity (index: 33, 46, 49 from top to bottom) with significant outliers, signal dropouts, and various brain orientations. The right column displays the corresponding masked reconstructed volumes with improved anatomical details.

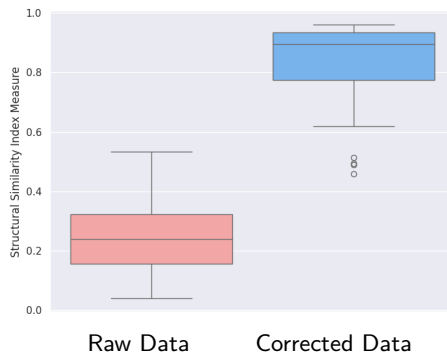


Figure 7: Distribution of SSIM Values for Raw and Motion-Corrected dMRI. Boxplots compare the SSIM distribution across slices of the first and last volume ($b\text{-value} = 900 \text{ s/m}^2$) for eight subjects. Raw data (red) exhibits lower SSIM compared to motion-corrected data (blue), indicating improved data integrity after motion correction.

4.3 Anatomical Accuracy of FOD and Tractography

While visual inspection of the dMRI data itself provides valuable insights, an additional assessment of the framework’s effectiveness in motion correction can be achieved by analyzing and comparing the derived white matter pathways with structural scan. Fig. 8 demonstrates good alignment of the FODs with the developing white matter structures, as visualized by the tractography streamlines overlaid on the T2w image. This qualitative finding suggests the potential of the framework for generating reliable dMRI data suitable for studying fetal brain development.




Figure 8: Fetal Brain Anatomy Visualized with HAITCH Output. Left panel displays the fiber orientation distribution (FOD) overlaid on a T2-weighted image (background) of a 35-week fetus. Tractography streamlines, colored according to diffusion directions, are shown in the right panel.

5 Discussion

Fetal dMRI faces unique challenges due to inherent motion and susceptibility-induced distortions. These factors lead to scattered data in the q-space domain and result in inconsistencies within and across volumes regarding the diffusion encoding direction and the corresponding diffusion-weighted contrast. Additionally, fetal motion and maternal breathing cause significant fluctuations in the B0 inhomogeneity during scans and render static field map correction techniques ineffective. These limitations necessitate dynamic field map estimation and multi-dimensional q-space reconstruction for reliable analysis of the fetal brain microstructure. HAITCH addresses these challenges by incorporating several strategies. First, the recommended multi-shell HARDI sampling scheme in HAITCH enriches the information content of the acquired data while enhancing its tolerance to fetal motion. Second, HAITCH leverages a dual-echo EPI sequence for dynamic field map estimation. Classic static field map correction shows inaccurate results due to motion occurring between the two unweighted ($b=0$) scans. Third, HAITCH builds on a data-driven representation that spans both the angular and the radial q-space domain, using SH alongside generalized Laguerre polynomials. The volume-level registration proved to be effective across hundreds of processed datasets. This feature

is attributed to the effectiveness of the new slice and voxel-level weighting and reconstruction steps in HAITCH to handle motion-corrupted data. The reconstructed fetal dMRI data obtained from HAITCH allows fitting various diffusion models, which paves the way for subsequent processing steps and enables a deeper understanding of tissue microstructure and white matter connectivity within the developing brain.

HAITCH is the first and the only publicly available tool, to-date, to process multi-shell and/or multi-echo fetal dMRI data. While validating the accuracy of HAITCH, it was not feasible to quantitatively compare it with the existing methods due to their limitations - limitations that HAITCH sought to address. The SVRTK [3] toolkit includes modules for fetal dMRI processing, but is restricted to single-shell data and cannot readily process multi-shell or multi-echo dMRI data. The Spherical Harmonics And Radial Decomposition (SHARD) [36], which was originally developed for neonatal dMRI data, requires an interleaved phase encoding scheme, which is based on a specific sequence described in [60]. SHARD utilizes an invariant B0 field map, neglecting dynamic field variations. An extended version of SHARD for fetal dMRI, explained by [61], utilizes the spin and field echo (SAFE) sequence for phase-based dynamic B0 estimation following [62]. While promising, this method relies on MRI sequences that may not be accessible to many researchers, and involves complex phase subtraction, which can be challenging. [62] reported limited accuracy of phase-based B0 estimation in the presence of motion exceeding 8.1 degrees, which may not be sufficient for all fetal MRI applications. [2] proposed a slice-level diffusion encoding and interleaved double spin-echo sequence for distortion and motion correction. We did not have access to this sequence and did not re-implement it on our scanner platform. Future work warrants a detailed comparison between various sequences and their effectiveness in the presence of fetal motion and geometric and intensity distortions that vary by fetal and maternal motion.

The information gained by fetal diffusion MRI cannot be obtained by any other imaging modality. Therefore, there has been significant interest in studying the development of the fetal brain microstructure and structural connectivity using dMRI. Early *ex-vivo* studies [63–68] have shed light on the capacity of dMRI to study the microstructure of the fetal brain. These studies were compared to histology and have also been used to indirectly validate *in-vivo* studies, e.g. [13]. However, postmortem fetal brain samples are scarce, and may not be available in sufficient numbers to study any specific disorder. *In-vivo* fetal dMRI is challenging due to fetal motion, geometric distortion, and the extremely low SNR available from the small fetal brain anatomy. Despite these impediments, significant progress has been made in the field: various motion-robust fetal dMRI techniques have been proposed [3, 20, 21, 23], atlases have been built and released [6, 11, 12, 69, 70], and several studies characterized normal and abnormal development of the fetal brain based on dMRI [5, 7–10, 71]. Almost all of these works, however, were based on the diffusion tensor model, which is known to have significant limitations in depicting complex fiber structures and connections. Advanced diffusion models that are needed to characterize complex brain connections and study subtle differences between groups of normal and abnormal fetuses, require multi-shell high-angular resolution dMRI, which cannot be processed reliably with any of the existing tool for fetal dMRI processing. By providing an open-source multi-stage framework for multi-shell high-angular resolution fetal dMRI, HAITCH fills in this critical gap.

6 Conclusion

We have developed HAITCH, a novel framework and an open-source public toolkit for acquiring and reconstructing high-quality fetal diffusion-weighted MRI. HAITCH utilizes a dual-echo EPI sequence which enables the dynamic field map estimation and reduces motion-related distortions. Our framework employs an advanced scheme that enriches the informational depth of the data and improves its motion tolerance. For motion correction and reconstruction, HAITCH leverages neighborhood directional information and a sophisticated data-driven *model-free* representation in conjunction with slice weighting and registration, effectively addressing signal dropouts and scattered data. Rigorous validation experiments on fetal dMRI scans have demonstrated HAITCH’s ability to significantly improve the integrity and reliability of fetal brain diffusion MRI data, paving the way for more accurate analyses of fetal brain development. The enhanced data quality facilitated by HAITCH has the potential to unlock new insights into the complex processes of brain development *in-utero*.

Acknowledgment

This research was supported in part by the National Institute of Biomedical Imaging and Bioengineering, the National Institute of Neurological Disorders and Stroke, and Eunice Kennedy Shriver National Institute of Child Health and Human Development of the National Institutes of Health (NIH) under award numbers R01NS106030, R01EB031849, R01EB032366, R01HD109395, R01HD110772, R01NS128281, and R01NS121657; in part by the Office of the Director of the NIH under award number S10OD025111; and in part by the National Science Foundation (NSF) under grant number 212306. This research was also partly supported by an award from NVIDIA Corporation and utilized NVIDIA RTX A6000 and RTX A5000 GPUs. The content of this publication is solely the responsibility of the authors and does not necessarily represent the official views of the NIH, NSF, or NVIDIA.

Code Availability

The code is publicly available in <https://github.com/bchimagine> and <https://github.com/FEDIToolbox>.

References

- [1] András Jakab, Ruth Tuura, Christian Kellenberger, and Ianina Scheer. In utero diffusion tensor imaging of the fetal brain: a reproducibility study. *NeuroImage: Clinical*, 15:601–612, 2017.
- [2] Jana Hutter, Daan J Christiaens, Torben Schneider, Lucilio Cordero-Grande, Paddy J Slator, Maria Deprez, Anthony N Price, J-Donald Tournier, Mary Rutherford, and Joseph V Hajnal. Slice-level diffusion encoding for motion and distortion correction. *Medical image analysis*, 48:214–229, 2018.
- [3] Maria Deprez, Anthony Price, Daan Christiaens, Georgia Lockwood Estrin, Lucilio Cordero-Grande, Jana Hutter, Alessandro Daducci, Jacques-Donald Tournier, Mary Rutherford, Serena J Counsell, et al. Higher order spherical harmonics reconstruction of fetal diffusion mri with intensity correction. *IEEE transactions on medical imaging*, 39(4):1104–1113, 2019.
- [4] David Hunt, Manjiri Dighe, Christopher Gatenby, and Colin Studholme. Challenges and opportunities in connectome construction and quantification in the developing human fetal brain. *Topics in Magnetic Resonance Imaging*, 28(5):265–273, 2019.
- [5] András Jakab. Developmental pathoconnectomics and advanced fetal mri. *Topics in Magnetic Resonance Imaging*, 28(5):275–284, 2019.
- [6] Shadab Khan, Lana Vasung, Bahram Marami, Caitlin K Rollins, Onur Afacan, Cynthia M Ortinau, Edward Yang, Simon K Warfield, and Ali Gholipour. Fetal brain growth portrayed by a spatiotemporal diffusion tensor mri atlas computed from in utero images. *Neuroimage*, 185:593–608, 2019.
- [7] Camilo Jaimes, Fedel Machado-Rivas, Onur Afacan, Shadab Khan, Bahram Marami, Cynthia M Ortinau, Caitlin K Rollins, Clemente Velasco-Annis, Simon K Warfield, and Ali Gholipour. In vivo characterization of emerging white matter microstructure in the fetal brain in the third trimester. *Human Brain Mapping*, 41(12):3177–3185, 2020.
- [8] Fedel Machado-Rivas, Onur Afacan, Shadab Khan, Bahram Marami, Clemente Velasco-Annis, Hart Lidov, Simon K Warfield, Ali Gholipour, and Camilo Jaimes. Spatiotemporal changes in diffusivity and anisotropy in fetal brain tractography. *Human Brain Mapping*, 42(17):5771–5784, 2021.
- [9] Camilo Calixto, Fedel Machado-Rivas, Maria C Cortes-Albornoz, Davood Karimi, Clemente Velasco-Annis, Onur Afacan, Simon K Warfield, Ali Gholipour, and Camilo Jaimes. Characterizing microstructural development in the fetal brain using diffusion mri from 23 to 36 weeks of gestation. *Cerebral Cortex*, page bhad409, 2023.
- [10] Camilo Calixto, Fedel Machado-Rivas, Davood Karimi, Clemente Velasco-Annis, Maria Camila Cortes-Albornoz, Onur Afacan, Simon K Warfield, Ali Gholipour, and Camilo Jaimes. Population atlas analysis of emerging brain structural connections in the human fetus. *Journal of Magnetic Resonance Imaging*, 2023.

- [11] Camilo Calixto, Athena Taymourash, Davood Karimi, Haykel Snoussi, Clemente Velasco-Annis, Camilo Jaimes, and Ali Gholipour. Advances in fetal brain imaging. *Magnetic Resonance Imaging Clinics*, 2024.
- [12] Camilo Calixto Nunez, Matheus Soldatelli, Camilo Jaimes Cobos, Simon K Warfield, Ali Gholipour, and Davood Karimi. A detailed spatio-temporal atlas of the white matter tracts for the fetal brain. *bioRxiv*, pages 2024–04, 2024.
- [13] Hamza Kebiri, Ali Gholipour, Rizhong Lin, Lana Vasung, Camilo Calixto, Željka Krsnik, Davood Karimi, and Meritxell Bach Cuadra. Deep learning microstructure estimation of developing brains from diffusion mri: a newborn and fetal study. *Medical Image Analysis*, 95:103186, 2024.
- [14] Edward O Stejskal and John E Tanner. Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *The journal of chemical physics*, 42(1):288–292, 1965.
- [15] Ali Gholipour, Judith A Estroff, Carol E Barnewolt, Richard L Robertson, P Ellen Grant, Borjan Gagoski, Simon K Warfield, Onur Afacan, Susan A Connolly, Jeffrey J Neil, et al. Fetal mri: a technical update with educational aspirations. *Concepts in Magnetic Resonance Part A*, 43(6):237–266, 2014.
- [16] Onur Afacan, Judy A Estroff, Edward Yang, Carol E Barnewolt, Susan A Connolly, Richard B Parad, Robert V Mulkern, Simon K Warfield, and Ali Gholipour. Fetal echoplanar imaging: promises and challenges. *Topics in Magnetic Resonance Imaging*, 28(5):245–254, 2019.
- [17] Jesper LR Andersson. Diffusion mri artifact correction. In *Advances in Magnetic Resonance Technology and Applications*, volume 4, pages 123–146. Elsevier, 2021.
- [18] Jiaxin Xiao, Cong Sun, Ruike Chen, Zhiyong Zhao, Guangbin Wang, and Dan Wu. Reproducibility of diffusion mri-based tractography in the fetal brain. *Journal of Magnetic Resonance Imaging*, 2024.
- [19] Shuzhou Jiang, Hui Xue, Serena Counsell, Mustafa Anjari, Joanna Allsop, Mary Rutherford, Daniel Rueckert, and Joseph V Hajnal. Diffusion tensor imaging (dti) of the brain in moving subjects: application to in-utero fetal and ex-utero studies. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 62(3):645–655, 2009.
- [20] Estanislao Oubel, Mériam Koob, Colin Studholme, Jean-Louis Dietemann, and François Rousseau. Reconstruction of scattered data in fetal diffusion mri. *Medical image analysis*, 16(1):28–37, 2012.
- [21] Mads Fogtman, Sharmishta Seshamani, Christopher Kroenke, Xi Cheng, Teresa Chapman, Jakob Wilm, Francois Rousseau, and Colin Studholme. A unified approach to diffusion direction sensitive slice registration and 3-d dti reconstruction from moving fetal brain anatomy. *IEEE transactions on medical imaging*, 33(2):272–289, 2013.
- [22] Benoit Scherrer, Ali Gholipour, and Simon K Warfield. Super-resolution reconstruction to increase the spatial resolution of diffusion weighted images from orthogonal anisotropic acquisitions. *Medical image analysis*, 16(7):1465–1476, 2012.
- [23] Bahram Marami, Seyed Sadegh Mohseni Salehi, Onur Afacan, Benoit Scherrer, Caitlin K Rollins, Edward Yang, Judy A Estroff, Simon K Warfield, and Ali Gholipour. Temporal slice registration and robust diffusion-tensor reconstruction for improved fetal brain structural connectivity analysis. *NeuroImage*, 156:475–488, 2017.
- [24] Maria Kuklisova-Murgasova, Gerardine Quaghebeur, Mary A Rutherford, Joseph V Hajnal, and Julia A Schnabel. Reconstruction of fetal brain mri with intensity matching and complete outlier removal. *Medical image analysis*, 16(8):1550–1564, 2012.
- [25] Jacques-Donald Tournier, Daan Christiaens, Jana Hutter, Anthony N Price, Lucilio Cordero-Grande, Emer Hughes, Matteo Bastiani, Stamatios N Sotiropoulos, Stephen M Smith, Daniel Rueckert, et al. A data-driven approach to optimising the encoding for multi-shell diffusion mri with application to neonatal imaging. *NMR in Biomedicine*, 33(9):e4348, 2020.

- [26] Peter Jezzard and Robert S Balaban. Correction for geometric distortion in echo planar images from b0 field variations. *Magnetic resonance in medicine*, 34(1):65–73, 1995.
- [27] Jesper LR Andersson, Stefan Skare, and John Ashburner. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *Neuroimage*, 20(2):870–888, 2003.
- [28] Jesper LR Andersson, Mark S Graham, Ivana Drobnjak, Hui Zhang, Nicola Filippini, and Matteo Bastiani. Towards a comprehensive framework for movement and distortion correction of diffusion mr images: Within volume movement. *Neuroimage*, 152:450–466, 2017.
- [29] Jessica Dubois, C Poupon, F Lethimonnier, and D Le Bihan. Optimized diffusion gradient orientation schemes for corrupted clinical dti data sets. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 19:134–143, 2006.
- [30] Emmanuel Caruyer, Christophe Lenglet, Guillermo Sapiro, and Rachid Deriche. Design of multishell sampling schemes with uniform coverage in diffusion mri. *Magnetic resonance in medicine*, 69(6):1534–1540, 2013.
- [31] J-Donald Tournier, Fernando Calamante, and Alan Connelly. Determination of the appropriate b value and number of gradient directions for high-angular-resolution diffusion-weighted imaging. *NMR in Biomedicine*, 26(12):1775–1786, 2013.
- [32] Davood Karimi, Camilo Jaimes, Fedel Machado-Rivas, Lana Vasung, Shadab Khan, Simon K Warfield, and Ali Gholipour. Deep learning-based parameter estimation in fetal diffusion-weighted mri. *Neuroimage*, 243:118482, 2021.
- [33] Evren Özarıslan, Cheng Guan Koay, Timothy M Shepherd, Michal E Komlos, M Okan İrfanođlu, Carlo Pierpaoli, and Peter J Basser. Mean apparent propagator (map) mri: a novel diffusion imaging method for mapping tissue microstructure. *NeuroImage*, 78:16–32, 2013.
- [34] Sylvain L Merlet and Rachid Deriche. Continuous diffusion signal, eap and odf estimation via compressive sensing in diffusion mri. *Medical image analysis*, 17(5):556–572, 2013.
- [35] Rutger Henri Jacques Fick, Mauro Zucchelli, Gabriel Girard, Maxime Descoteaux, Gloria Menegaz, and Rachid Deriche. Using 3d-shore and map-mri to obtain both tractography and microstructural contrast from a clinical dmri acquisition. In *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, pages 436–439. IEEE, 2015.
- [36] Daan Christiaens et al. Scattered slice shard reconstruction for motion correction in multi-shell diffusion mri. *Neuroimage*, 225:117437, 2021.
- [37] Maria Kuklisova-Murgasova, Georgia Lockwood Estrin, Rita G Nunes, Shaihan J Malik, Mary A Rutherford, Daniel Rueckert, and Joseph V Hajnal. Distortion correction in fetal epi using non-rigid registration with a laplacian constraint. *IEEE transactions on medical imaging*, 37(1):12–19, 2017.
- [38] Onur Afacan, W Scott Hoge, Tess E Wallace, Ali Gholipour, Sila Kurugol, and Simon K Warfield. Simultaneous motion and distortion correction using dual-echo diffusion-weighted mri. *Journal of Neuroimaging*, 30(3):276–285, 2020.
- [39] Alena Uus, Tong Zhang, Laurence H Jackson, Thomas A Roberts, Mary A Rutherford, Joseph V Hajnal, and Maria Deprez. Deformable slice-to-volume registration for motion correction of fetal body and placenta mri. *IEEE transactions on medical imaging*, 39(9):2750–2759, 2020.
- [40] Michael Ebner, Guotai Wang, Wenqi Li, Michael Aertsen, Premal A Patel, Rosalind Aughwane, Andrew Melbourne, Tom Doel, Steven Dymarkowski, Paolo De Coppi, et al. An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain mri. *NeuroImage*, 206:116324, 2020.
- [41] Lucilio Cordero-Grande, Daan Christiaens, Jana Hutter, Anthony N Price, and Jo V Hajnal. Complex diffusion-weighted image estimation via matrix recovery under general noise models. *Neuroimage*, 200:391–404, 2019.
- [42] Jelle Veraart, Els Fieremans, and Dmitry S Novikov. Diffusion mri noise mapping using random matrix theory. *Magnetic resonance in medicine*, 76(5):1582–1593, 2016b.

- [43] Jelle Veraart, Dmitry S Novikov, Daan Christiaens, Benjamin Ades-Aron, Jan Sijbers, and Els Fieremans. Denoising of diffusion mri using random matrix theory. *Neuroimage*, 142:394–406, 2016c.
- [44] Florent Benaych-Georges and Raj Rao Nadakuditi. The singular values and vectors of low rank perturbations of large rectangular random matrices. *Journal of Multivariate Analysis*, 111:120–135, 2012.
- [45] Elias Kellner, Bibek Dhital, Valerij G Kiselev, and Marco Reisert. Gibbs-ringing artifact removal based on local subvoxel-shifts. *Magnetic resonance in medicine*, 76(5):1574–1581, 2016.
- [46] Benjamin Ades-Aron, Jelle Veraart, Peter Kochunov, Stephen McGuire, Paul Sherman, Elias Kellner, Dmitry S Novikov, and Els Fieremans. Evaluation of the accuracy and precision of the diffusion parameter estimation with gibbs and noise removal pipeline. *Neuroimage*, 183:532–543, 2018.
- [47] Cheng Guan Koay and Peter J Basser. Analytically exact correction scheme for signal extraction from noisy magnitude mr signals. *Journal of magnetic resonance*, 179(2):317–322, 2006.
- [48] Henning U Voss, Richard Watts, Aziz M Uluğ, and Doug Ballon. Fiber tracking in the cervical spine and inferior brain regions with reversed gradient diffusion tensor imaging. *Magnetic resonance imaging*, 24(3):231–239, 2006.
- [49] Nicholas J Tustison, Brian B Avants, Philip A Cook, Yuanjie Zheng, Alexander Egan, Paul A Yushkevich, and James C Gee. N4itk: improved n3 bias correction. *IEEE transactions on medical imaging*, 29(6):1310–1320, 2010.
- [50] Davood Karimi, Haoran Dou, Simon K Warfield, and Ali Gholipour. Deep learning with noisy labels: Exploring techniques and remedies in medical image analysis. *Medical image analysis*, 65:101759, 2020.
- [51] Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2):203–211, 2021.
- [52] Eleftherios Garyfallidis, Matthew Brett, Bagrat Amirbekian, Ariel Rokem, Stefan Van Der Walt, Maxime Descoteaux, Ian Nimmo-Smith, and Dipy Contributors. Dipy, a library for the analysis of diffusion mri data. *Frontiers in neuroinformatics*, 8:8, 2014.
- [53] Brian B Avants, Nick Tustison, Gang Song, et al. Advanced normalization tools (ants). *Insight j*, 2(365):1–35, 2009.
- [54] Ali Gholipour, Caitlin K Rollins, Clemente Velasco-Annis, Abdelhakim Ouaalalm, Alireza Akhondi-Asl, Onur Afacan, Cynthia M Ortinau, Sean Clancy, Catherine Limperopoulos, Edward Yang, et al. A normative spatiotemporal mri atlas of the fetal brain for automatic segmentation and analysis of early brain growth. *Scientific reports*, 7(1):476, 2017.
- [55] J-Donald Tournier, Robert Smith, David Raffelt, Rami Tabbara, Thijs Dhollander, Maximilian Pietsch, Daan Christiaens, Ben Jeurissen, Chun-Hung Yeh, and Alan Connelly. Mrtrix3: A fast, flexible and open software framework for medical image processing and visualisation. *Neuroimage*, 202:116137, 2019.
- [56] Ben Jeurissen, Jacques-Donald Tournier, Thijs Dhollander, Alan Connelly, and Jan Sijbers. Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion mri data. *NeuroImage*, 103:411–426, 2014.
- [57] J Donald Tournier, Fernando Calamante, Alan Connelly, et al. Improved probabilistic streamlines tractography by 2nd order integration over fibre orientation distributions. In *Proceedings of the international society for magnetic resonance in medicine*, volume 1670. John Wiley & Sons, Inc New Jersey, USA, 2010.
- [58] Renaud Hedouin, Olivier Commowick, Elise Banner, Benoit Scherrer, Maxime Taquet, Simon K Warfield, and Christian Barillot. Block-matching distortion correction of echo-planar images with opposite phase encoding directions. *IEEE Trans. Med. Imaging*, 36(5):1106–1115, 2017.

- [59] Dominic Holland, Joshua M Kuperman, and Anders M Dale. Efficient correction of inhomogeneous static magnetic field-induced distortion in echo planar imaging. *Neuroimage*, 50(1):175–183, 2010.
- [60] Chitresh Bhushan, Anand A Joshi, Richard M Leahy, and Justin P Haldar. Improved b0-distortion correction in diffusion mri using interlaced q-space sampling and constrained reconstruction. *Magnetic Resonance in Medicine*, 72(5):1218–1232, 2014.
- [61] Lucilio Cordero-Grande, Anthony Price, Giulio Ferrazzi, Jana Hutter, Daan Christiaens, Emer Hughes, and Joseph V Hajnal. Spin and field echo (safe) dynamic field correction in 3t fetal epi. In *Proceedings of the 26th annual meeting of ISMRM, City*, page 208, 2018.
- [62] Barbara Dymerska, Benedikt A Poser, Markus Barth, Siegfried Trattnig, and Simon D Robinson. A method for the dynamic correction of b0-related distortions in single-echo epi at 7 t. *Neuroimage*, 168:321–331, 2018.
- [63] Hao Huang, Rong Xue, Jiangyang Zhang, Tianbo Ren, Linda J Richards, Paul Yarowsky, Michael I Miller, and Susumu Mori. Anatomical characterization of human fetal brain development with diffusion tensor magnetic resonance imaging. *Journal of Neuroscience*, 29(13):4263–4273, 2009.
- [64] Lana Vasung, Hao Huang, Nataša Jovanov-Milošević, Mihovil Pletikos, Susumu Mori, and Ivica Kostović. Development of axonal pathways in the human fetal fronto-limbic brain: histochemical characterization and diffusion tensor imaging. *Journal of anatomy*, 217(4):400–417, 2010.
- [65] Emi Takahashi, Rebecca D Folkerth, Albert M Galaburda, and Patricia E Grant. Emerging cerebral connectivity in the human fetal brain: an mr tractography study. *Cerebral cortex*, 22(2):455–464, 2012.
- [66] Emi Takahashi, Emiko Hayashi, Jeremy D Schmahmann, and P Ellen Grant. Development of cerebellar connectivity in human fetal brains revealed by high angular resolution diffusion tractography. *Neuroimage*, 96:326–333, 2014.
- [67] Hao Huang and Lana Vasung. Gaining insight of fetal brain development with diffusion mri and histology. *International Journal of Developmental Neuroscience*, 32:11–22, 2014.
- [68] Lana Vasung, Marina Raguz, Ivica Kostovic, and Emi Takahashi. Spatiotemporal relationship of brain pathways during human fetal development using high-angular resolution diffusion mr imaging and histology. *Frontiers in neuroscience*, 11:259415, 2017.
- [69] Ruike Chen, Cong Sun, Tingting Liu, Yuhao Liao, Junyan Wang, Yi Sun, Yi Zhang, Guangbin Wang, and Dan Wu. Deciphering the developmental order and microstructural patterns of early white matter pathways in a diffusion mri based fetal brain atlas. *NeuroImage*, 264:119700, 2022.
- [70] A Uus, V Kyriakopoulou, L Cordero Grande, D Christiaens, M Pietsch, A Price, S Wilson, P Patkee, S Karolis, A Schuh, et al. Multi-channel spatio-temporal mri atlas of the normal fetal brain development from the developing human connectome project. *G-Node doi*, 10, 2023.
- [71] András Jakab, Gregor Kasprian, Ernst Schwartz, Gerlinde Maria Gruber, Christian Mitter, Daniela Prayer, Veronika Schöpf, and Georg Langs. Disrupted developmental organization of the structural connectome in fetuses with corpus callosum agenesis. *Neuroimage*, 111:277–288, 2015.