



Through-Plane Super-Resolution With Autoencoders in Diffusion Magnetic Resonance Imaging of the Developing Human Brain

Hamza Kebiri^{1,2*}, Erick J. Canales-Rodríguez³, Hélène Lajous^{1,2}, Priscille de Dumast^{1,2}, Gabriel Girard^{1,2,3}, Yasser Alemán-Gómez¹, Mériam Koob¹, András Jakab^{4,5} and Meritxell Bach Cuadra^{1,2,3}

¹ Department of Radiology, Lausanne University Hospital and University of Lausanne, Lausanne, Switzerland, ² CIBM Center for Biomedical Imaging, Lausanne, Switzerland, ³ Signal Processing Laboratory 5 (LTS5), Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, ⁴ Center for MR Research University Children's Hospital Zurich, Zurich, Switzerland, ⁵ Neuroscience Center Zurich, University of Zurich, Zurich, Switzerland

OPEN ACCESS

Edited by:

Daan Christiaens,
KU Leuven, Belgium

Reviewed by:

Maryam Afzali,
Cardiff University, United Kingdom
Arghya Pal,
Harvard Medical School,
United States

*Correspondence:

Hamza Kebiri
hamza.kebiri@unil.ch

Specialty section:

This article was submitted to
Applied Neuroimaging,
a section of the journal
Frontiers in Neurology

Received: 02 December 2021

Accepted: 28 March 2022

Published: 02 May 2022

Citation:

Kebiri H, Canales-Rodríguez EJ, Lajous H, de Dumast P, Girard G, Alemán-Gómez Y, Koob M, Jakab A and Bach Cuadra M (2022) Through-Plane Super-Resolution With Autoencoders in Diffusion Magnetic Resonance Imaging of the Developing Human Brain. *Front. Neurol.* 13:827816. doi: 10.3389/fneur.2022.827816

Fetal brain diffusion magnetic resonance images (MRI) are often acquired with a lower through-plane than in-plane resolution. This anisotropy is often overcome by classical upsampling methods such as linear or cubic interpolation. In this work, we employ an unsupervised learning algorithm using an autoencoder neural network for single-image through-plane super-resolution by leveraging a large amount of data. Our framework, which can also be used for slice outliers replacement, overperformed conventional interpolations quantitatively and qualitatively on pre-term newborns of the developing Human Connectome Project. The evaluation was performed on both the original diffusion-weighted signal and the estimated diffusion tensor maps. A byproduct of our autoencoder was its ability to act as a denoiser. The network was able to generalize fetal data with different levels of motions and we qualitatively showed its consistency, hence supporting the relevance of pre-term datasets to improve the processing of fetal brain images.

Keywords: unsupervised learning, autoencoders, super-resolution, diffusion-weighted imaging, magnetic resonance imaging (MRI), pre-term neonates, fetuses, brain

1. INTRODUCTION

The formation and maturation of white matter are at their highest rate during the fetal stage of human brain development. To have more insight into this critical period, *in utero* brain imaging techniques offer a unique opportunity. Diffusion weighted-magnetic resonance imaging (DW-MRI) is a well-established tool to reconstruct *in vivo* and non-invasively the white matter tracts in the brain (1, 2). Fetal DW-MRI, in particular, could characterize early developmental trajectories in brain connections and microstructure (3–6). Hence, fetal DW-MRI has been of significant interest for the past years where studies (7–9) have provided analysis using diffusion tensor imaging (DTI) by computing diffusion scalar maps such as fractional anisotropy (FA) or mean diffusivity (MD), using a limited number of gradient directions. A recent study focused on reconstructing fiber Orientation Distribution Functions (fODF) (10) using higher quality datasets and rich information including several gradient directions (32 and 80), higher b-values (750 and

1,000 s/mm^2), and signal-to-noise ratio (SNR) (3 Tesla magnetic field strength). Additionally, the datasets were acquired in a controlled and uniform research setting with healthy volunteers, which can hardly be reproduced in the clinical environment.

Albeit promising results, acquiring high-quality data remains the main obstacle in the field of fetal brain imaging. First, unpredictable and uncontrollable fetal motion is a major challenge. To overcome this problem, fast echo-planar imaging (EPI) sequences are typically used to freeze intra-slice motion. However, intra- and inter-volume motion still have to be addressed in the post-processing steps using sophisticated slice-to-volume registration (SVR) (11–13). Moreover, EPI sequences generate severe non-linear distortions that need adapted distortion correction algorithms (14). Additionally, the resulting images display low SNR due to at least three factors: the inherently small size of the fetal brain, the surrounding maternal structures and amniotic fluid, and the increased distance to the coils. In order to compensate for the low SNR in EPI sequences, series with thick voxels (i.e., low through-plane resolution) are often acquired. Finally, to shorten the acquisition time, small b-values ($b = 400 - 700s/mm^2$) and a low number of gradient directions (10–15) (8, 9) are commonly used in fetal imaging, which in turn will result in a low angular resolution.

Clinical protocols typically acquire several anisotropic orthogonal series of 2D thick slices to cope with high motion and low SNR. Then, super-resolution reconstruction techniques that have been originally developed for structural T2-weighted images (15–20) by combining different 3D low resolution volumes have also been successfully applied in 4D fetal functional (21) and diffusion MRI contexts (10, 12). Still, despite these two pioneer works, super-resolution DW-MRI from multiple volumes has been barely explored *in vivo*. In fact, the limited scanning time to minimize maternal discomfort hampers the acquisition of several orthogonal series, resulting in a trade-off between the number of gradient directions and orthogonal series. Thus, DW-MRI fetal brain protocols are not standardized from one center to another (**Supplementary Table S1**) and more experiments have to be conducted in this area to design optimal sequences (22, 23). Sequence-based super-resolution methods that were applied in adult brains (24–27) could also be explored and adapted to fetal brains such as in Ning et al. (24) that acquire same orientation shifted low-resolution images in the slice encoding direction and in a non-overlapping gradient scheme to reconstruct one high-resolution volume using compressed sensing. The term *super-resolution* is used by both the image processing and the MR sequence development communities, though in a slightly different way. While the former works mainly on image space and the latter works on *k-space*, both aim at increasing the image resolution at different stages either using multiple volumes or single volumes.

In fact, fetal DW-MRI resolution enhancement could also benefit from *single image super-resolution* approaches, i.e., either within each DW-MRI 3D volume separately or using the whole 4D volume including all diffusion measurements. It has indeed been demonstrated that a linear or cubic interpolation of the raw signal enhances the resulting scalar maps and tractography (28). In practice, this is typically performed either at the signal

level or at DTI scalar maps (29). We believe that single volume and multiple volumes super-resolution can also be performed together, i.e., where the output of the former is given as the input of the latter. This aggregation could potentially lead to a better motion correction and hence to a more accurate final high-resolution volume.

Several studies have proposed single image super-resolution enhancement methods for DW-MRI but, to the best of our knowledge, none of them has been applied neither to anisotropic datasets nor to the developing brain. In Coupé et al. (30), the authors utilized a non-local patch-based approach in an inverse problem paradigm to improve the resolution of adult brain DW-MRI volumes using a non diffusion weighted image ($b = 0s/mm^2$) prior. Although this approach yielded competitive results, it was built upon a sophisticated pipeline which made it not extensively used. The first machine learning study (31) have used shallow learning algorithms to learn the mapping between diffusion tensor maps of a downsampled high-resolution image and the maps of the original image. Recently, deep learning models which can implicitly learn relevant features from training data were used to perform single image super-resolution with a convolutional neural network (32, 33) and a customized U-Net (34, 35). Both approaches produced promising results in a *supervised* learning scheme. Supervision needs however large high quality datasets that are scarce for the perinatal brain for the reasons enumerated above.

The specific challenge of fetal DW-MRI is 3–5 mm acquired slice thickness, with only a few repetitions available. Hence, our main objective is to focus on through-plane DW-MRI resolution enhancement. This would be valuable not only for native anisotropic volumes but also for outlier slice recovery. In fact, motion-corrupted slices in DW-MRI is either discarded, which results in a loss of information, or replaced using interpolation (36–38). We approached this problem from an image synthesis point of view using *unsupervised* learning networks such as autoencoders (AEs), as demonstrated in cardiac T2-weighted MRI (39) and recent works in DW-MRI (40). Here, we present a framework with autoencoders that are neural networks learning in an *unsupervised* way to encode efficient data representations and can behave as generative models if this representation is structured enough. By accurately encoding DW-MRI slices in a low-dimensional latent space, we were able to successfully generate new slices that accurately correspond to in-between “missing” slices. In contrast to the above referred *supervised* learning approaches, this method is scale agnostic, i.e., the enhancement scale factor can be set *a posteriori* to the network training.

Realistically enhancing the through-plane resolution would potentially help the clinicians to better assess whether the anterior and posterior commissures are present in cases with complete agenesis of the corpus callosum (6). It can reduce partial volume effects and thus contribute to the depiction of more accurate white matter properties in the developing brain.

In this work, we present the first unsupervised through-plane resolution enhancement for perinatal brain DW-MRI. We leverage the high-quality dataset of the developing Human Connectome Project (dHCP) where we train and quantitatively

TABLE 1 | Pre-term newborns and fetal attributes of the processed data that were used in our experiments.

	GW	Resolution (mm^3)	#directions	b-value (s/mm^2)	Scanner	#subjects
Pre-terms	[29.3,37.0]	1.17x1.17x1.5	88	0 and 1,000	Philips 3T	31
Fetal						
Sub-1	35	1x1x5	15	0 and 700	GE 1.5T	–
Sub-2	29	1x1x4	15	0 and 700	GE 1.5T	–
Sub-3	24	1x1x5	15	0 and 700	GE 1.5T	–
Sub-4	23	1x1x4	15	0 and 700	GE 1.5T	–
Sub-5	24	1x1x5	15	0 and 700	GE 1.5T	–
Sub-6	27	1x1x5	15	0 and 700	GE 1.5T	–

The distribution of gestational ages is shown in **Supplementary Figure S1**.

validate pre-term newborns that are anatomically close to fetal subjects. We finally demonstrate the performance of our approach in fetal brains.

2. MATERIALS AND METHODS

2.1. Data

2.1.1. Pre-term dHCP Data

We selected all the 31 pre-term newborns of 37 gestational weeks (GW) or less at the time of scan (range: [29.3,37.0], mean: 35.5, median: 35.7) from the dHCP dataset (41) (subject IDs in **Supplementary Table S2**). Acquisitions were performed using a 3T Philips Achieva scanner (32-channel neonatal head-coil and 70 mT/m gradients) with a monopolar spin-echo EPI Stejskal-Tanner sequence ($\Delta = 42.5$ ms, $\delta = 14$ ms, TR = 3,800 ms, TE = 90,ms, echo spacing = 0.81ms, EPI factor = 83) and a multiband factor of 4, resulting in an acquisition time of 19:20 min. In a field of view of $150 \times 150 \times 102$ mm^3 , 64 interleaved slices were acquired with an in-plane resolution of 1.5 mm, a slice thickness of 3 mm, and a slice overlap of 1.5 mm. An isotropic volume of 1.5 mm^3 was obtained after super-resolution. The dataset was acquired with a multi-shell sequence using four b-values ($b \in \{0, 400, 1,000, \text{and } 2,600\} s/mm^2$) with 300 volumes but we have only extracted the 88 volumes corresponding to $b = 1,000 s/mm^2$ (b1000) as a compromise of high contrast-to-noise ratio (CNR), i.e., b1000 has a higher CNR than b400 and b2600 (42), and proximity to the $b = 700 s/mm^2$ that is typically used in clinical settings for fetal DW-MRI. The main attributes of the pre-term data are summarized in **Table 1**. Brain masks and region/tissue labels segmented using a pipeline based on the Draw-EM algorithm (43, 44) were available in the corresponding anatomical dataset. All the images were already corrected (42) for inter-slice motion and distortion (susceptibility, eddy currents and motion). After pre-processing, the final image resolution and FOV were, respectively, $1.17 \times 1.17 \times 1.5$ mm^3 and $128 \times 128 \times 64$ mm^3 .

2.1.2. Fetal Data

Fetal acquisitions were performed at 1.5T (MR450, GE Healthcare, Milwaukee, WI, USA) in the University Children's Hospital Zürich (KISPI) using a single-shot EPI sequence (TE = 63 ms, TR = 2200 ms) and 15 gradient directions at $b = 700 s/mm^2$ (b700). The acquisition time was approximately 1.3

min per 4D volume. The in-plane resolution was 1×1 mm^2 , the slice thickness was 4–5 mm, and the field of view $256 \times 256 \times 14$ – 22 voxels. Three axial series and a coronal one were acquired for each subject. Brain masks were manually generated for the b_0 ($b = 0 s/mm^2$) of each acquisition and automatically propagated to the diffusion-weighted volumes. Between 8 and 18, T2-weighted images were also acquired for each subject where corresponding brain masks were automatically generated using an in-house deep learning based method using transfer learning from Salehi et al. (45). Manual refinements were needed for a few cases at the brain boundaries.

2.1.3. Fetal Data Processing

We selected three subjects with high quality imaging and without motion artifacts (24, 29, and 35 GW) and three subjects with a varying degree of motion (23, 24, and 27 GW). **Supplementary Figure S1** shows the distribution of gestational age of both 31 pre-term newborns and the 6 fetal subjects used in this study. A DW-MRI volume of a motion-free case (Sub-2, 29 GW) and a pre-term of equivalent age are illustrated in **Figure 1**. By performing quality control, we discarded highly corrupted volumes due to motion resulting in severe signal drops in two moving subjects and very low SNR volumes in one motion-free subject. **Table 2** presents the different characteristics of each subject as well as its corresponding discarded volumes. The coronal volume was not used to avoid any interpolation confounding factor while co-registering different orientations. All the subjects were pre-processed for noise, bias field inhomogeneities, and distortions using the Nipype framework (46). The denoising was performed using a Principal Component Analysis based method (47), followed by an N4 bias-field inhomogeneity correction (48). Distortion was corrected using an in-house implementation of a state-of-the-art algorithm for the fetal brain (14) consisting in rigid registration (49) of a structural T2-weighted image to the b_0 image, followed by a non-linear registration (49) in the phase-encoding direction of the b_0 to the same T2-weighted image. The transformation was then applied to the diffusion-weighted volumes. A block matching algorithm for symmetric global registration was also performed for two subjects (sub-4, sub-6) with motion [NiftyReg, (50)]. The b_0 image of the first axial series was selected as a reference to which we subsequently registered the remaining volumes, i.e.,

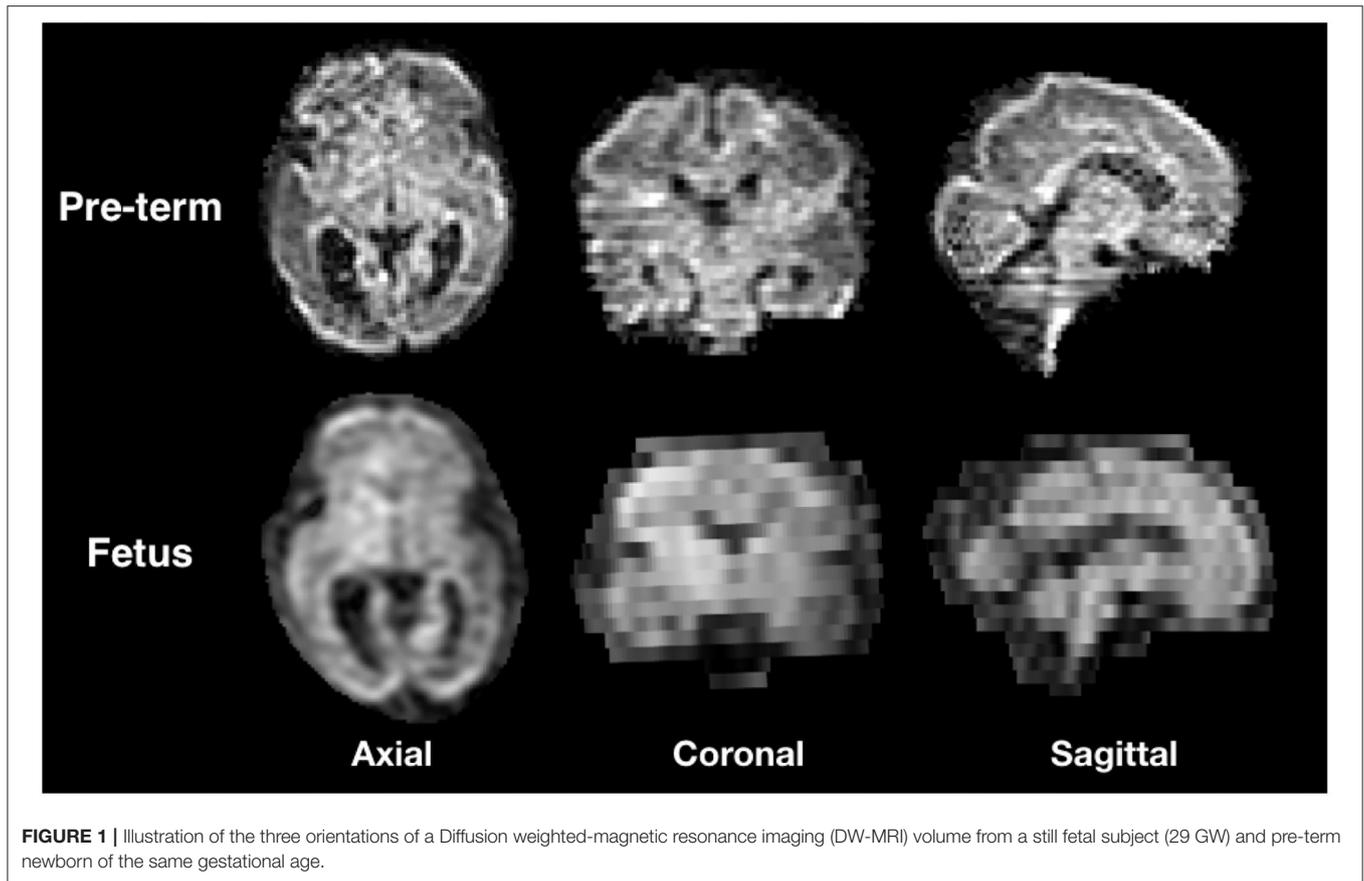


FIGURE 1 | Illustration of the three orientations of a Diffusion weighted-magnetic resonance imaging (DW-MRI) volume from a still fetal subject (29 GW) and pre-term newborn of the same gestational age.

TABLE 2 | Fetal motion level and discarded directions for each 4D volume.

	Motion level	Discarded			
		Axial-1	Axial-2	Axial-3	Coronal
Sub-1	No	–	–	–	All
Sub-2	No	–	–	–	All
Sub-3	No	All except b0	–	All except b0	All
Sub-4	High	Vol 2,3,4,5,6,7	Vol 11,14	Vol 14	All
Sub-5	Average	–	Vol 6,7,15	Vol 11	All
Sub-6	Low	Vol 11	Vol 11	Vol 11,14	All

the non b0 images from the first axial and all volumes from the two others. Gradient directions were rotated accordingly. **Supplementary Figure S2** shows an example of a DWI volume (from sub-4) of original, pre-processing, and motion correction.

2.2. Model

2.2.1. Architecture

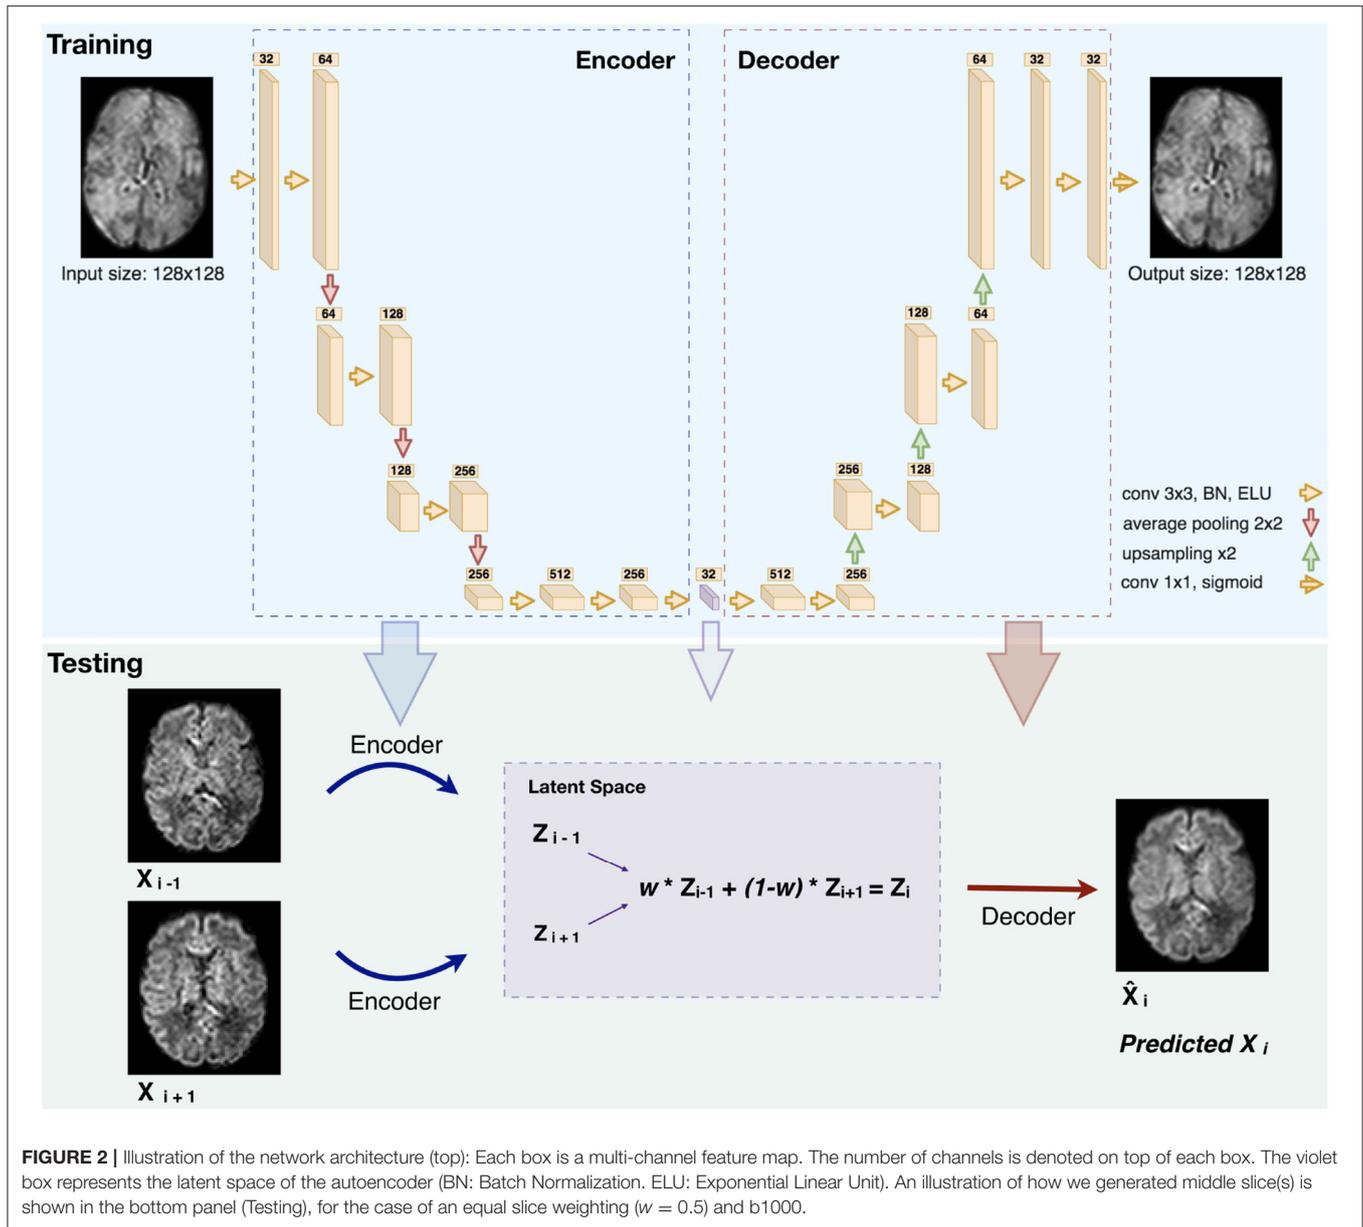
Our network architecture, similarly to Sander et al. (39), is composed of four blocks in the encoder and four in the decoder (**Figure 2**). Each block in the encoder consists of two layers made of 3 x 3 convolutions followed by a batch normalization (51) and an Exponential Linear Unit non-linearity.

The number of feature maps is doubled from 32 after each layer and the resulting feature maps are average-pooled. We

further added two layers of two 3 x 3 convolutions in which the feature maps of the last layer were used as the latent-space of the autoencoder. The decoder uses the same architecture as the encoder but by conversely halving the number of feature maps and upsampling after each block using nearest-neighbor interpolation. At the final layer a 1 x 1 convolution using the sigmoid function is applied to output the predicted image. The number of network parameters is 6,098,689.

2.2.2. Training and Optimization

We have trained our network solely on b0 images (15 per subject), using an 8-fold nested cross validation where we trained and validated on 27 subjects and tested on four. The proportion of the validation data was set to 15% of the training set. The



training/validation set contains 25,920 slices of a 128 x 128 field of view, totaling 424,673,280 voxels. Our network was trained in an unsupervised manner by feeding normalized 2D axial slices that are encoded as feature maps in the latent space. The number of feature maps, and hence the dimensionality of the latent space, was optimized (optimal value to 32) using Keras-Tuner (52). The batch size and the learning rate were additionally optimized and set to 32 and 5e-5, respectively. The network that was initialized using (53) was trained for 200 epochs to minimize the mean squared error loss between the predicted and the ground truth image. We have utilized for this aim the Adam optimizer (54) with the default parameters $\beta_1 = 0.5$, $\beta_2 = 0.999$, and the network corresponding to the epoch with the minimal validation loss was then selected. The implementation was performed in

the framework of TensorFlow 2.4.1 (55) and an Nvidia GeForce RTX 2080 GPU was deployed for training. Network code and checkpoint examples can be found in our Github repository¹.

2.2.3. Inference

The network trained on b0 images was used for the inference of b0 and b1000 volumes. Two slices were encoded in the latent space and their N “in-between” slice(s) ($N = 1, 2$ in our experiments) were predicted using weighted averages of the latent codes of the two slices. The weights for $N = 1$ and $N = 2$ were set proportionally to their distance to the neighboring

¹www.github.com/Medical-Image-Analysis-Laboratory/Perinatal_SR_Auto-encoder

original slices [as performed in Sander et al. (39)], i.e., an equal weighting for $N = 1$ and $\{\frac{1}{3}, \frac{2}{3}\}$, $\{\frac{2}{3}, \frac{1}{3}\}$ for $N = 2$. Performing a grid search on ten weights (0.1–0.9 with a step of 0.1) confirmed the optimality of the previous choice. An example of pre-term b1000 data for a weight of 0.5 is shown in **Figure 2** (Testing). Similarly, the same b0 network was also used to enhance the through-plane resolution of fetal b0 and b700 volumes. Finally, since the network outputs were normalized between 0 and 1, histogram normalization to the weighted average of the input images was performed.

2.3. Experiments and Evaluation

2.3.1. Pre-term Newborns

Our network was separately tested on b0 images and the 88 volumes of b1000 using an 8-fold cross validation where 7-folds contain four subjects and 1-fold contains three subjects. We removed N intermediate slices ($N = [1, 2]$) from the testing set volumes in alternating order and used the (weighted) average latent space feature maps of the to-be adjacent slices to encode the N missing slice(s) using the autoencoder (**Figure 2**, Testing). The resulting latent representation was then decoded to predict the N slices in the voxel space, which were compared to the previously removed N slices, i.e., the ground truth (GT). The same N slices were also generated using three baseline approaches: trilinear, tricubic, and B-spline of 5th order interpolations [using Tournier et al. (56) and Avants et al. (49)] for comparison. We denote them, respectively, for removing one or two slices: Linear-1, Cubic-1, Spline-1 and Linear-2, Cubic-2, Spline-2.

Latent space exploratory analysis - In order to have an intuitive idea of the latent space representation, we have compared the latent space representations between different gradient directions of all possible pairs from the 88 volumes of the b1000 4D volume. As two volumes with closely aligned gradient directions are more similar than two volumes with orthogonal directions, we aimed to check whether this property is globally preserved in the latent encoding of our input images.

Robustness to noise - We have added different low levels of Rician noise (57) to the original signal as follows: for each pixel with a current intensity S_{clean} , the new intensity $S_{noisy} = \sqrt{(S_{clean} + GN_1)^2 + GN_2^2}$, where GN_1 and GN_2 are random numbers sampled from a Gaussian distribution with zero mean and a SD of $S_{clean}(b = 0)/SNR_{out}$ and SNR_{out} is the desired SNR we aim to simulate. Three SNRs of {27, 25, 23} and {20, 16, 13} were simulated for b0 and b1000, respectively. We have used higher noise levels for b1000 to better simulate the inherently lower SNR in this configuration.

Scalar maps - By merging the b0 and b1000 using the autoencoder enhancement, we reconstructed FA, MD, axial diffusivity (AD), and radial diffusivity (RD) from DTI using Dipy (58) separately for AE-1 or AE-2, i.e., where we, respectively, remove one ($N = 1$) or two slices ($N = 2$). We further subdivided the computation in specific brain regions (cortical gray matter, white matter, corpus callosum, and brainstem as provided by the dHCP). Region labels were upsampled and manually refined to match the super-resolved/interpolated volumes. We performed similar computation of the diffusion

maps generated using the trilinear, tricubic, and B-spline interpolated signals.

2.3.2. Fetal

For each subject and each 3D volume (b0 or DW-MRI), we generated one or two middle slices using the autoencoder, hence synthetically enhancing the resolution from $1 \times 1 \times 4\text{--}5 \text{ mm}^3$ to a simulated resolution of $1 \times 1 \times 2\text{--}2.5 \text{ mm}^3$ and $1 \times 1 \times 1.33\text{--}1.67 \text{ mm}^3$, respectively. We then generated whole-brain DTI maps (FA, MD, AD, and RD) and showed the colored FA. Splenium and genu structures of the corpus callosum were additionally segmented on FA maps for subjects in which these structures were visible. The mean FA and MD were reported for these regions for original and autoencoder enhanced volumes.

2.3.3. Quantitative Evaluation

Raw diffusion signal - We computed the voxel-wise error between the raw signal synthesized by the autoencoder and the GT using the mean squared error (MSE) and the peak SNR (PSNR). We compared the autoencoder performance with the three baseline approaches: trilinear, tricubic, and B-spline of 5th order interpolations.

Latent space exploratory analysis - We have computed the average squared Euclidian voxel-wise distance between slices of all 3D b1000 volume pairs. This was performed both at the input space and at the latent space representation. The images were flattened from 2D to one-dimensional vectors and compared as follows:

$$d(\vec{u}, \vec{v}) = \|\vec{u} - \vec{v}\| = \sqrt{(u_1 - v_1)^2 + (u_2 - v_2)^2 + \dots + (u_n - v_n)^2} \quad (1)$$

Where \vec{u} and \vec{v} are the vectors to be compared for all the n corresponding pixels. The final distance between each two 3D volumes is the average distance of all 2D distance computed in 1.

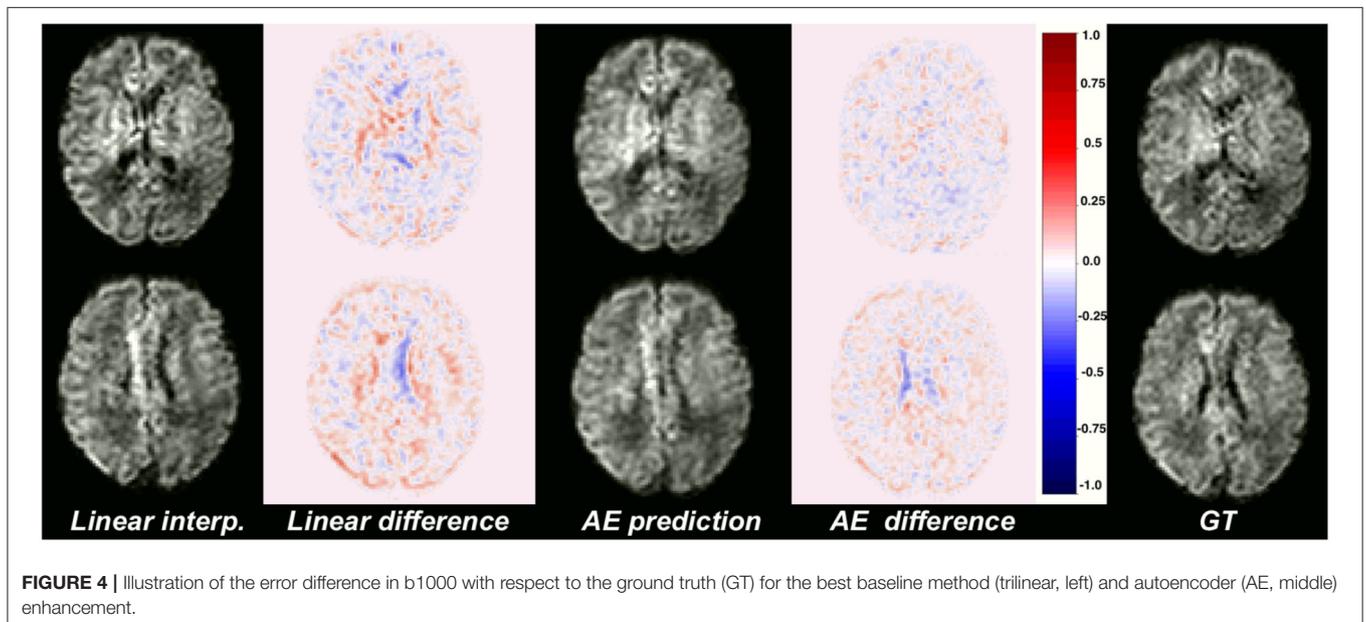
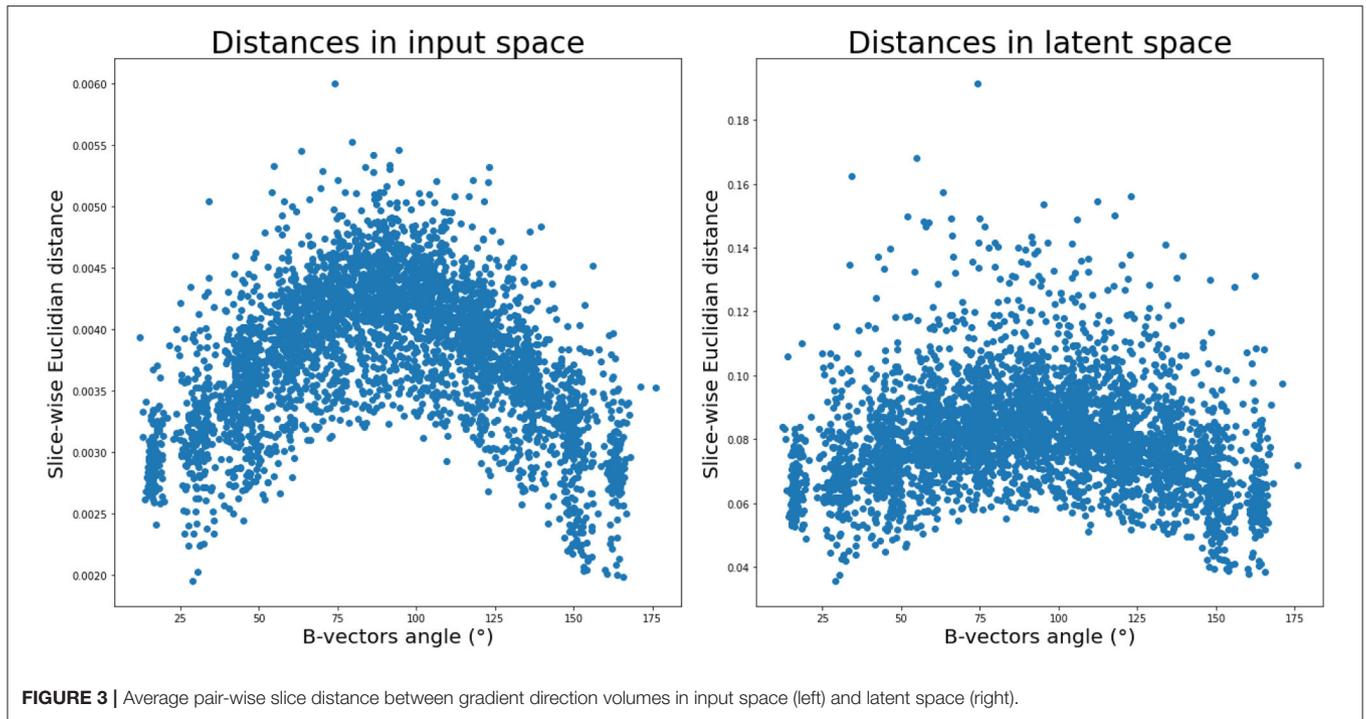
Robustness to noise - We computed with respect to the GT signal, the error of the signal with noise, and the output of the autoencoder using the signal with noise as input. We compared the results using MSE separately for b0 and b1000.

Scalar maps - We computed the voxel-wise error between the diffusion tensor maps reconstructed with the GT and the one by merging the b0 and b1000 using the autoencoder enhancement. We computed the error separately using either AE-1 or AE-2. We used the MSE and the PSNR as metrics and the same diffusion maps generated using the trilinear, tricubic, and B-spline interpolated signal as a baseline. Moreover, we qualitatively compare colored FA generated using the best baseline method, autoencoder, and the GT.

3. RESULTS

3.1. Pre-term Newborns

First, we inspected the latent space and how the 88 DW-MRI volumes are encoded with respect to each other. We can notice in **Figure 3** (right panel) that as two b-vectors' angle approaches orthogonality (90°), the difference between the latent representations of their corresponding volumes increases. On



the contrary, the difference decreases the more the angle tends toward 0° or 180° . Although the pattern is more pronounced in the input space (Figure 3, left panel), this trend is a fulfilled necessary condition to the generation for coherent representations of the input data by our network.

Moreover, our network that was exclusively trained on b0 images was able to generalize to b1000. In fact, the signal similarity between b0 and DW images was also used in Coupé et al. (30) in an inverse problem paradigm in which a b0 prior was incorporated to reconstruct b700 volumes.

Figure 4 illustrates qualitative results and absolute errors for $N = 1$ with respect to the GT (right) between the best interpolation baseline (trilinear, left) and the autoencoder enhancement (middle) for b1000. We overall saw from these representative examples, higher absolute intensities in the Linear-1 configuration than in the AE-1. However, ventricles are less visible when using an autoencoder. We hypothesize this is because of their higher intensity in b0 images on which the network was trained.

The average MSE with respect to the original DW-MRI signal within the whole brain is shown in **Figure 5** for both the autoencoder enhanced volume and the baseline methods

(trilinear, tricubic, and B-spline), for the configurations where one (Method-1) or two (Method-2) slices were removed. The first observation was the expected higher error for the configuration

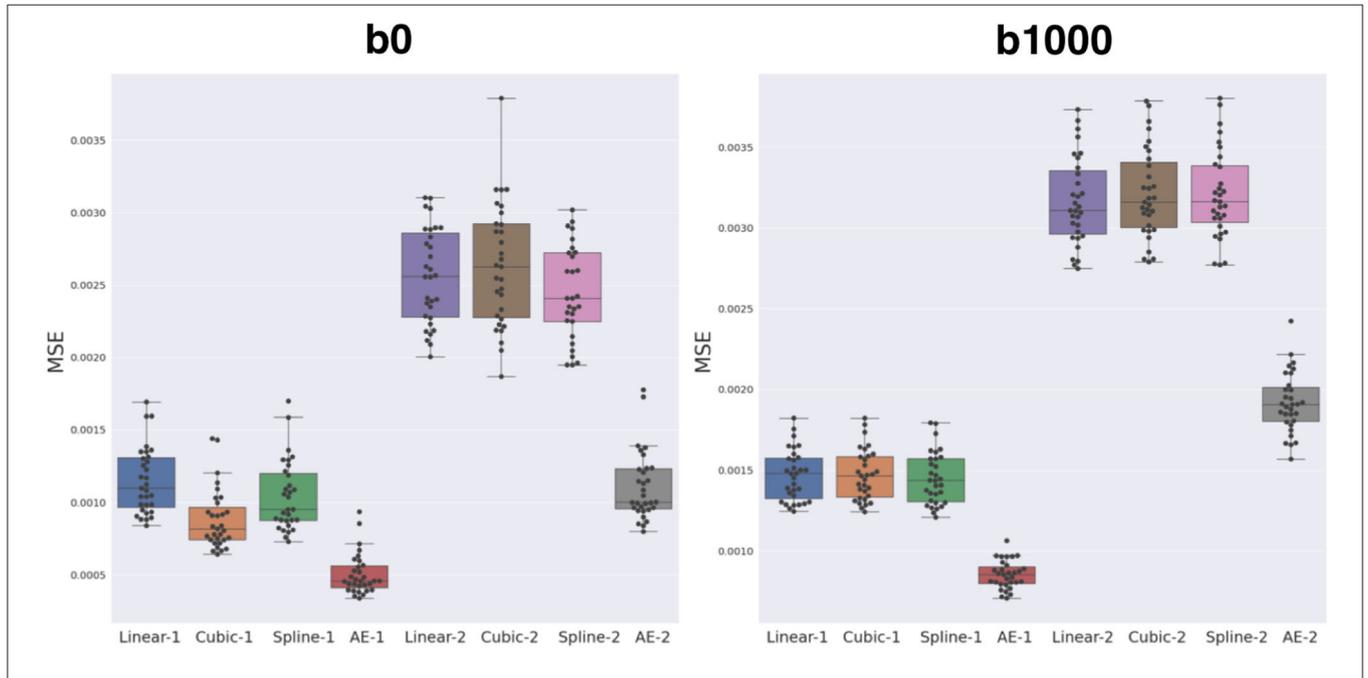


FIGURE 5 | Mean squared error (MSE) between the three baseline methods (linear, cubic, and B-spline 5th order) and autoencoder (AE) enhancement both for b0 (left) and b1000 (right). Two configurations were assessed: either $N = 1$, i.e., removing one slice and interpolating/synthesizing it (Linear-1, Cubic-1, Spline-1, AE-1) or $N = 2$, i.e., the same approach with two slices (Linear-2, Cubic-2, Spline-2, and AE-2). The autoencoder has a significantly lower MSE when compared to each respective best baseline method (paired Wilcoxon signed-rank test $p < 1.24e-09$).

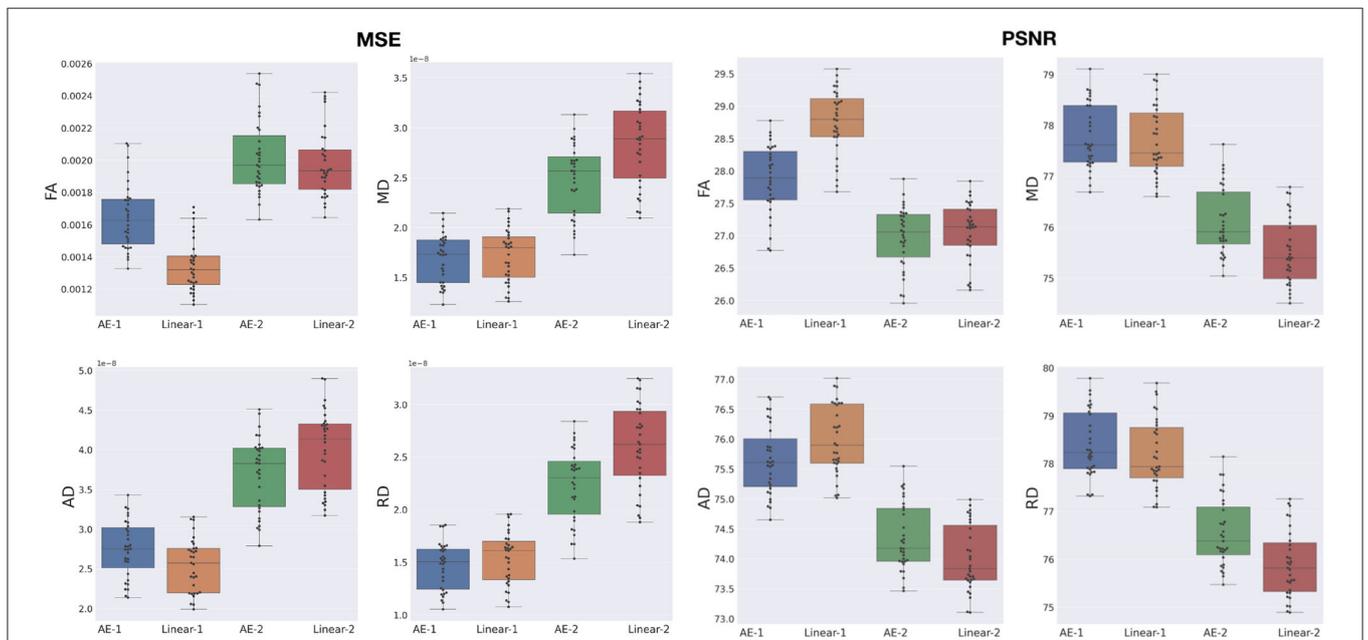


FIGURE 6 | MSE and peak signal-to-noise ratio (PSNR) between the best baseline (Linear) and the autoencoder (AE) enhancement for whole-brain diffusion tensor maps, when removing and synthesizing/interpolating one or two slices. FA, Fractional Anisotropy; MD, Mean Diffusivity; AD, Axial Diffusivity; RD, Radial Diffusivity.

where two slices are removed ($N = 2$), independently of the method used. Additionally, the autoencoder enhancement clearly outperformed the baseline methods in all configurations (paired Wilcoxon signed-rank test $p < 1.24e-09$). Particularly, the more slices we remove, the higher the gap between the baseline interpolation methods and the autoencoder enhancement. For b0, the MSE gain was around 0.0005 for $N = 1$ and 0.0015 for $N = 2$ between the autoencoder and the average baseline method (Spline-1 v.s. AE-1 and Linear-2 v.s. AE-2). For b1000, the gain

between AE-1 and Cubic-1 was 0.0007 and 0.0015 between AE-2 and Cubic-2.

The overperformance of the autoencoder is also shown overall in the DTI maps, where MD, AD, and RD were better approximated when compared to the best baseline method (linear interpolation), particularly in the configuration where two slices were removed (Figure 6). However, the FA showed the opposite trend, especially for the configuration, where one slice was removed (AE-1 v.s. Linear-1). However, FA

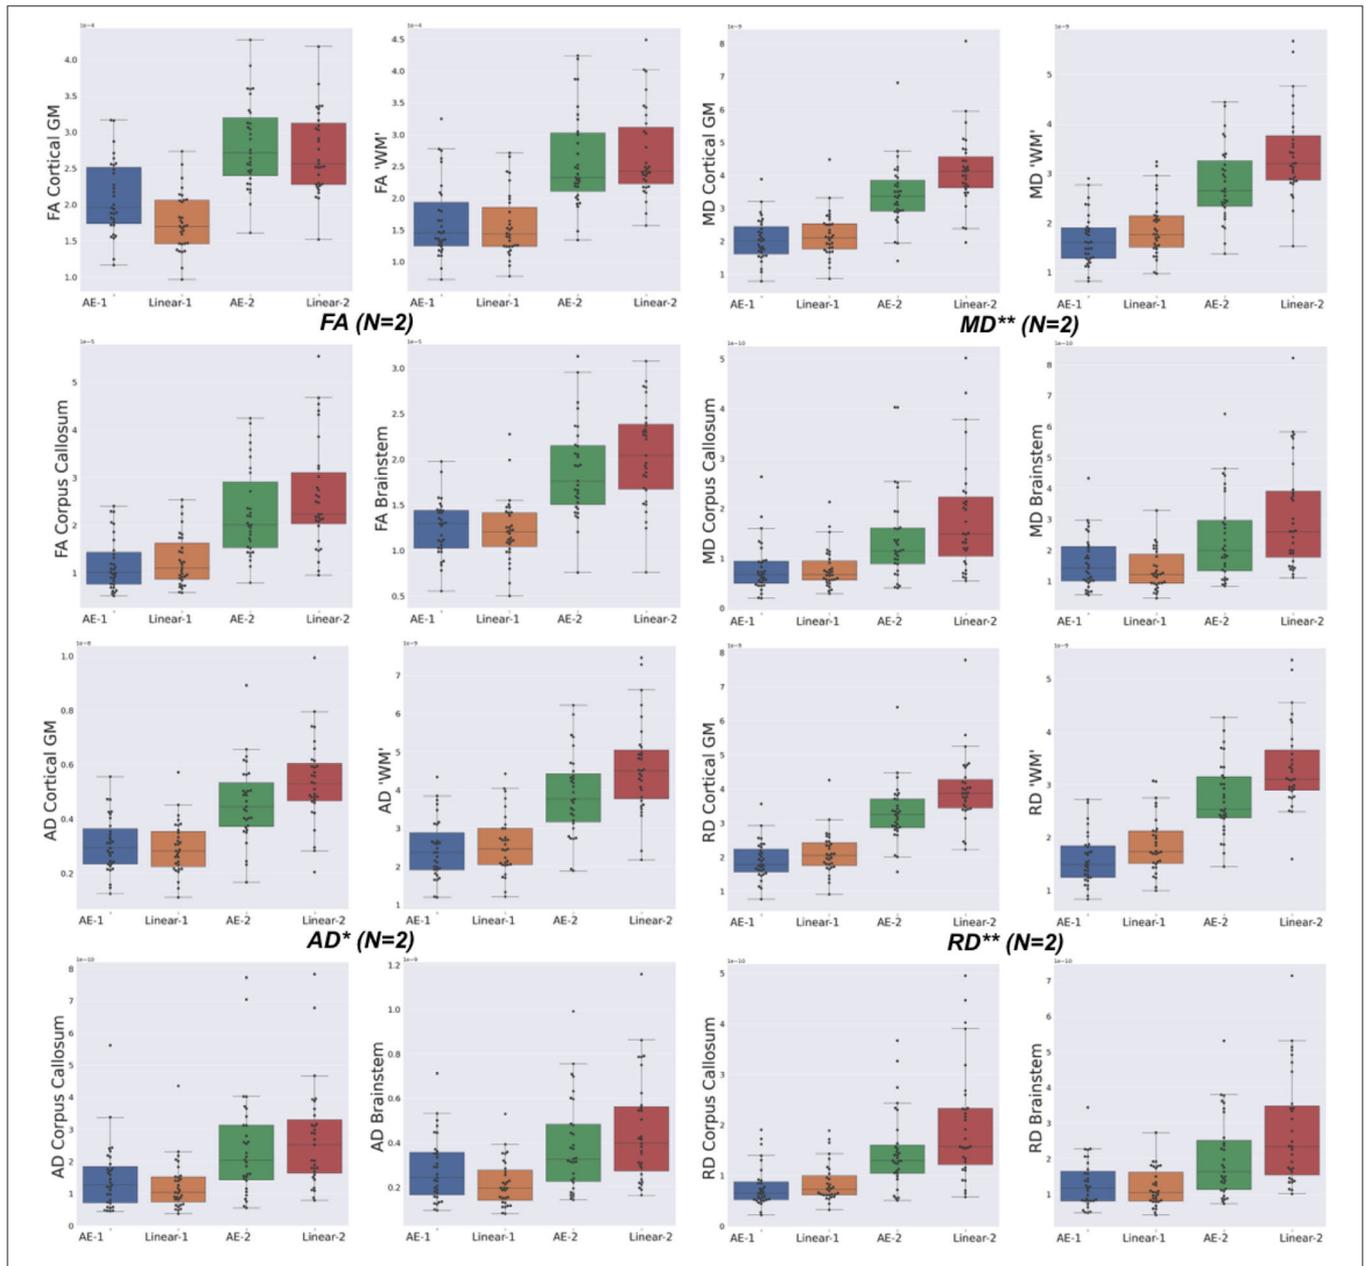
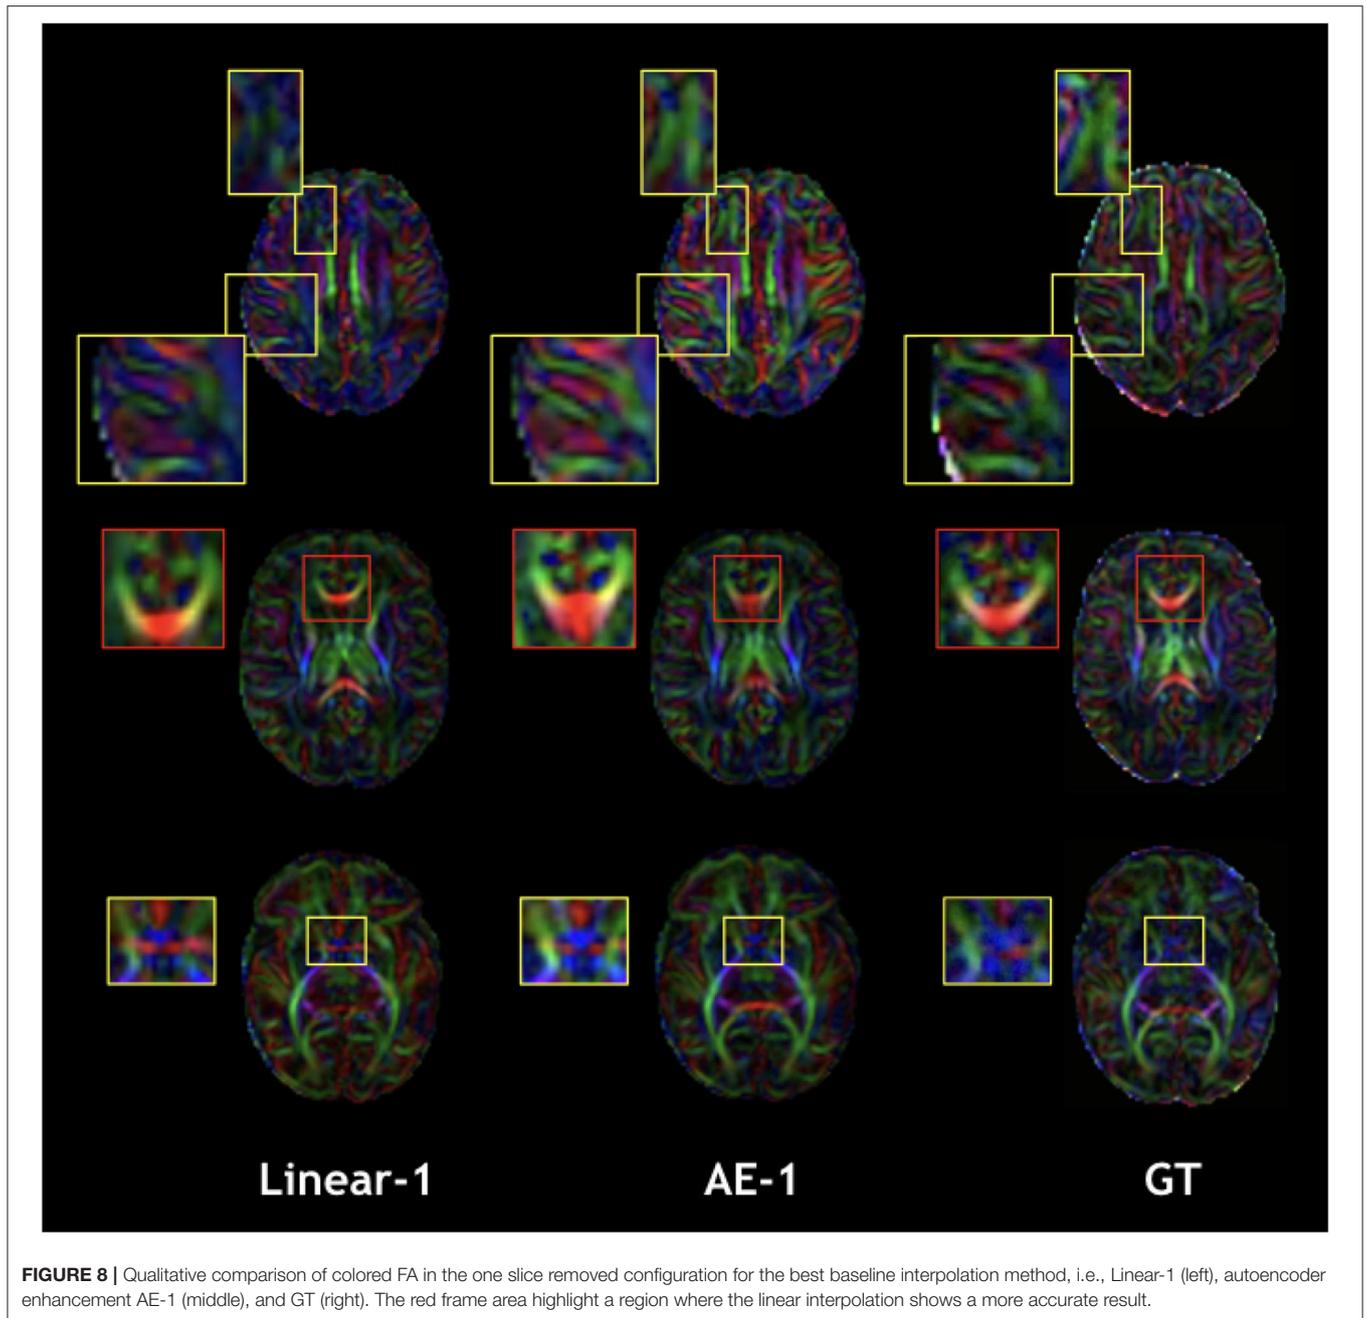


FIGURE 7 | Mean squared error (MSE) with respect to the GT of the best baseline method (Linear) and the autoencoder (AE) enhancement in the different brain structures [Cortical Gray Matter (GM), White Matter (WM), Brainstem, and Corpus Callosum] for each diffusion tensor map (FA, MD, AD, and RD) for one slice removal ($N = 1$) and two slices removal ($N = 2$). Comparing the DTI maps of the merged brain region labels, we found that the AE-2 significantly outperforms other conventional methods for MD, RD, and AD. (Paired Wilcoxon signed-rank test: $**p < 0.0018$ and $*p < 0.017$).



for white matter-like structures (“WM”, corpus callosum, and brainstem) showed higher performance with the autoencoder as depicted for each structure in **Figure 7**. In fact, by plotting colored FA for these two configurations, we observed that the autoencoder generates tracts that were consistent with the GT. For instance, autoencoder enhancement showed higher frequency details around the superficial WM area (**Figure 8**, top row) and removed artifacts between the internal capsules better than the linear method (**Figure 8**, bottom row). However, in some cases, the baseline method better depicted tracts such as in the corpus callosum (**Figure 8**, middle row).

ODFs generated using spherical harmonics order 8 are also depicted in **Supplementary Figure S3** where the autoencoder enhanced data show little qualitative differences with the GT ODFs. **Figure 9** shows similar comparisons for MD in different brain regions between the baseline method (Linear), the autoencoder, and the GT. Overall, quantitatively, for structures in the case where two slices were removed, the autoencoder enhancement outperformed the best baseline method in 15 out of 16 configurations (**Figure 7**). However, it is not always the case when one slice is removed, such as in the AD of the brainstem.

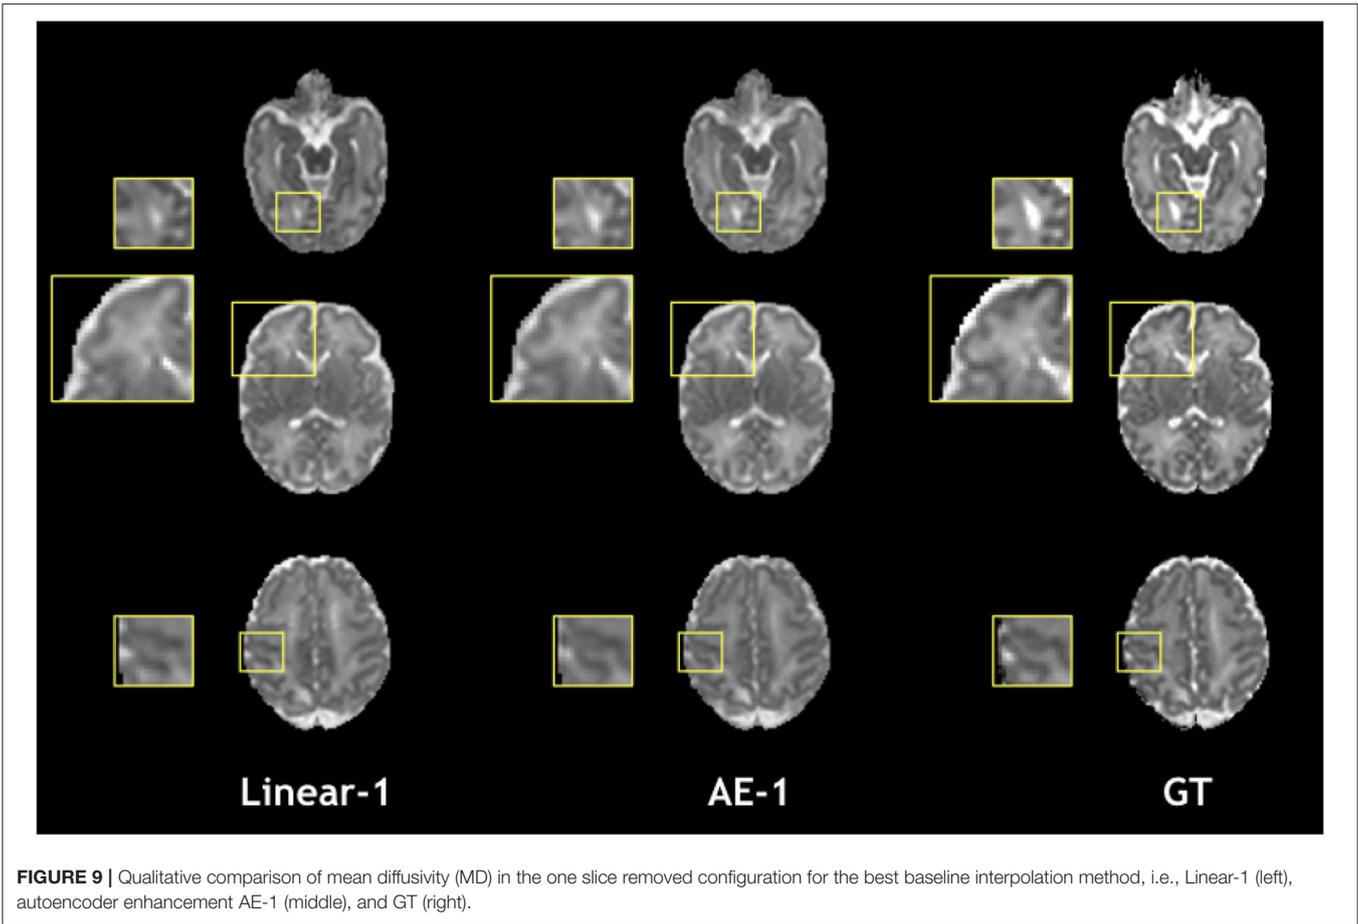


FIGURE 9 | Qualitative comparison of mean diffusivity (MD) in the one slice removed configuration for the best baseline interpolation method, i.e., Linear-1 (left), autoencoder enhancement AE-1 (middle), and GT (right).

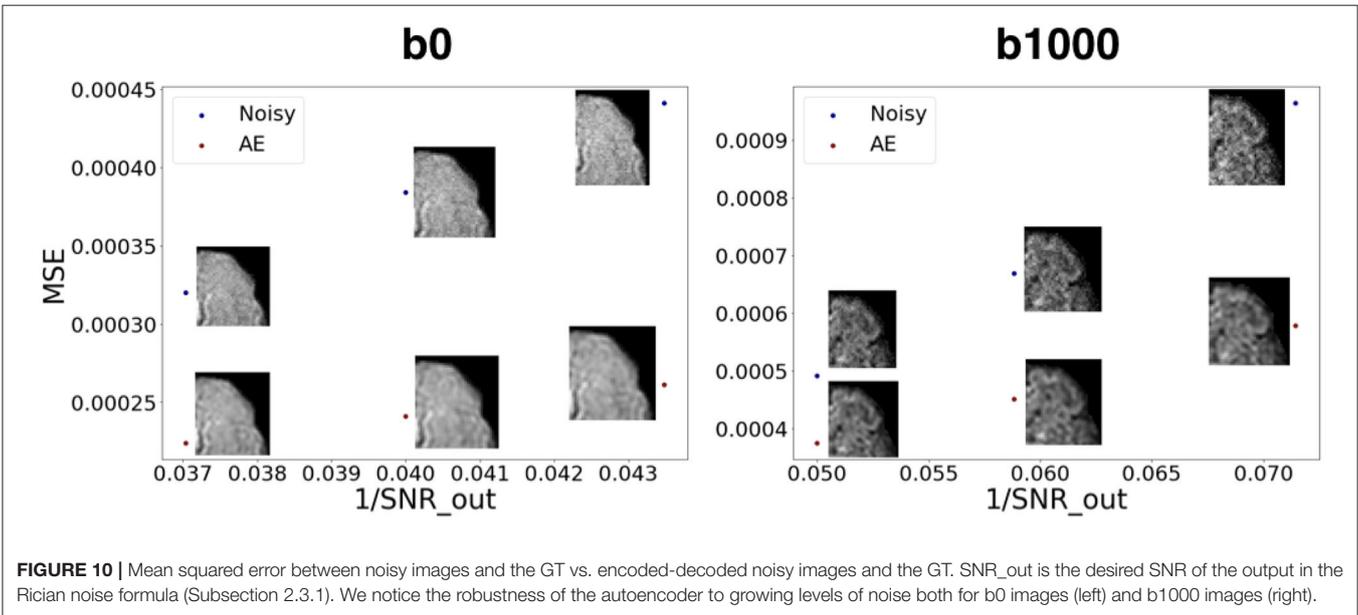


FIGURE 10 | Mean squared error between noisy images and the GT vs. encoded-decoded noisy images and the GT. SNR_{out} is the desired SNR of the output in the Rician noise formula (Subsection 2.3.1). We notice the robustness of the autoencoder to growing levels of noise both for b0 images (left) and b1000 images (right).

Figure 10 shows how our autoencoder was robust to reasonable amounts of noise. In fact, simply encoding and decoding the noisy input generates a

slice that was closer to the GT than the noisy slice, as depicted for different levels of noise for both b0 and b1000.

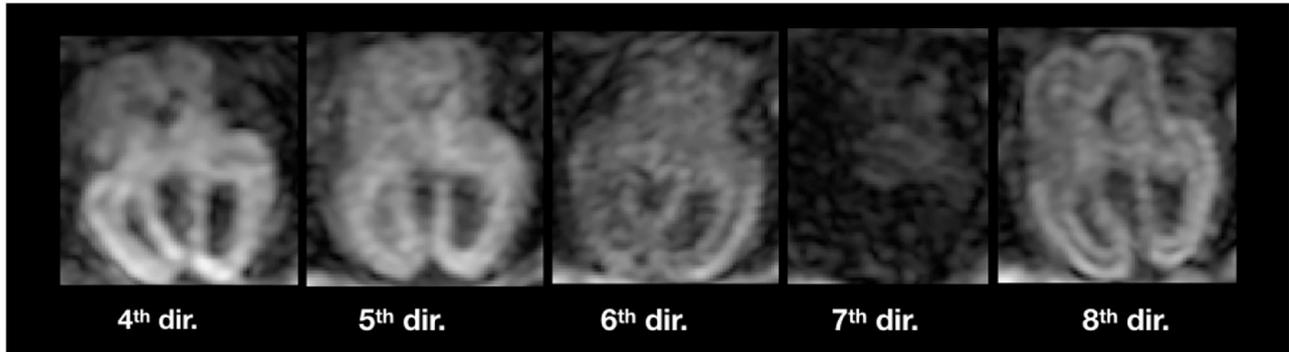


FIGURE 11 | Illustration of inter-volume motion in five different gradient directions of sub-4 (Table 2). Note the severe signal drop in the seventh direction because of motion.

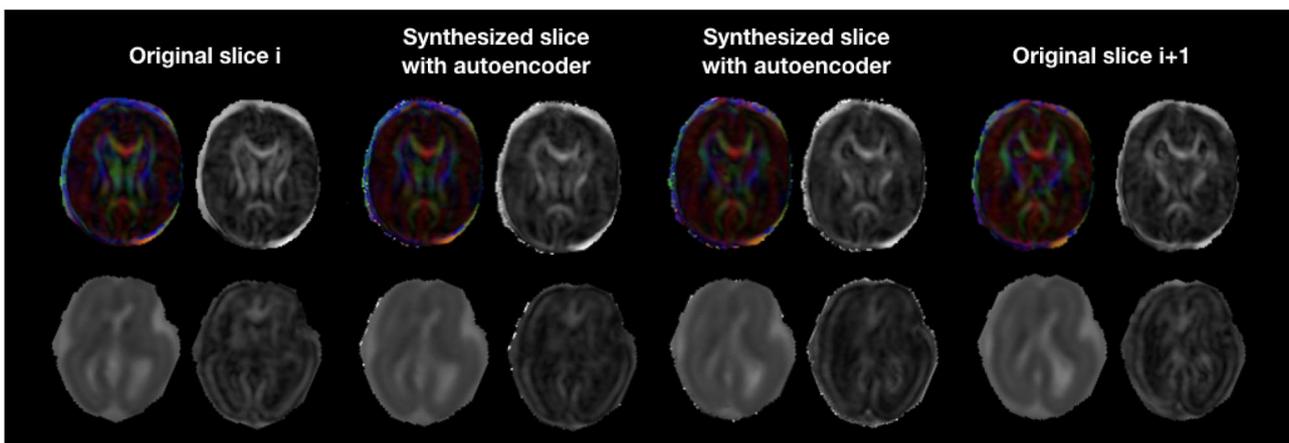


FIGURE 12 | Colored FA and FA (top row) illustration of autoencoder enhancement between two original adjacent fetal slices in a still subject (sub-1, 35 GW). The bottom row shows a similar illustration of MD and FA for a moving subject (sub-4, 23 GW).

3.2. Fetuses

Figure 11 illustrates inter-volume motion between five diffusion-weighted volumes where we also notice a severe signal drop in the seventh direction (sub-4, 23 GW).

The autoencoder trained on pre-term b0 images was able to coherently enhance fetal acquisitions both at b0 and DW-MRI volumes at b700. The network was able to learn low-level features that could generalize over anatomy, contrast, and b-values. Corresponding FA and colored FA for a still subject (sub-1, 35 GW) are illustrated in **Figure 12** (top) where we clearly see the coherence of the two synthesized images as we go from one original slice to the next one. In fact, both the corpus callosum and the internal/external capsules follow a smooth transition between the two slices. Similarly, **Figure 12** (bottom) exhibits MD and FA for a moving subject (sub-4, 23 GW) where we also notice, particularly for the MD, the smooth transition between the originally adjacent slices. FA and MD for the remaining subjects are shown in **Supplementary Figure S4**. Tractography on a fetal subject (sub-1, 35 GW) using both the original and autoencoder enhancement AE-1 DW-MRI is shown in **Supplementary Figure S5**.

The splenium and genu of the corpus callosum were only sufficiently visible in the three late GW subjects (sub-1, sub-2, and sub-6). **Figure 13** shows quantitative results for FA and MD in the two structures. Both maps fall into the range of reported values in the literature (59) for the respective gestational age, for original and autoencoder enhanced volumes.

4. DISCUSSION

In this work, we have shown that (1) autoencoders can be used for through-plane super-resolution in diffusion MRI, (2) training on b0 images can generalize to gradient diffusion volumes of different contrasts, and (3) as a proof of concept, training on pre-term anatomy can generalize to fetal images.

In fact, we have demonstrated how autoencoders can realistically enhance the resolution of DW-MRI perinatal brain images. We have compared it to conventionally used methods such as trilinear, tricubic, and B-spline interpolations both qualitatively and quantitatively for pre-term newborns of the dHCP database. Resolution enhancement was performed at the

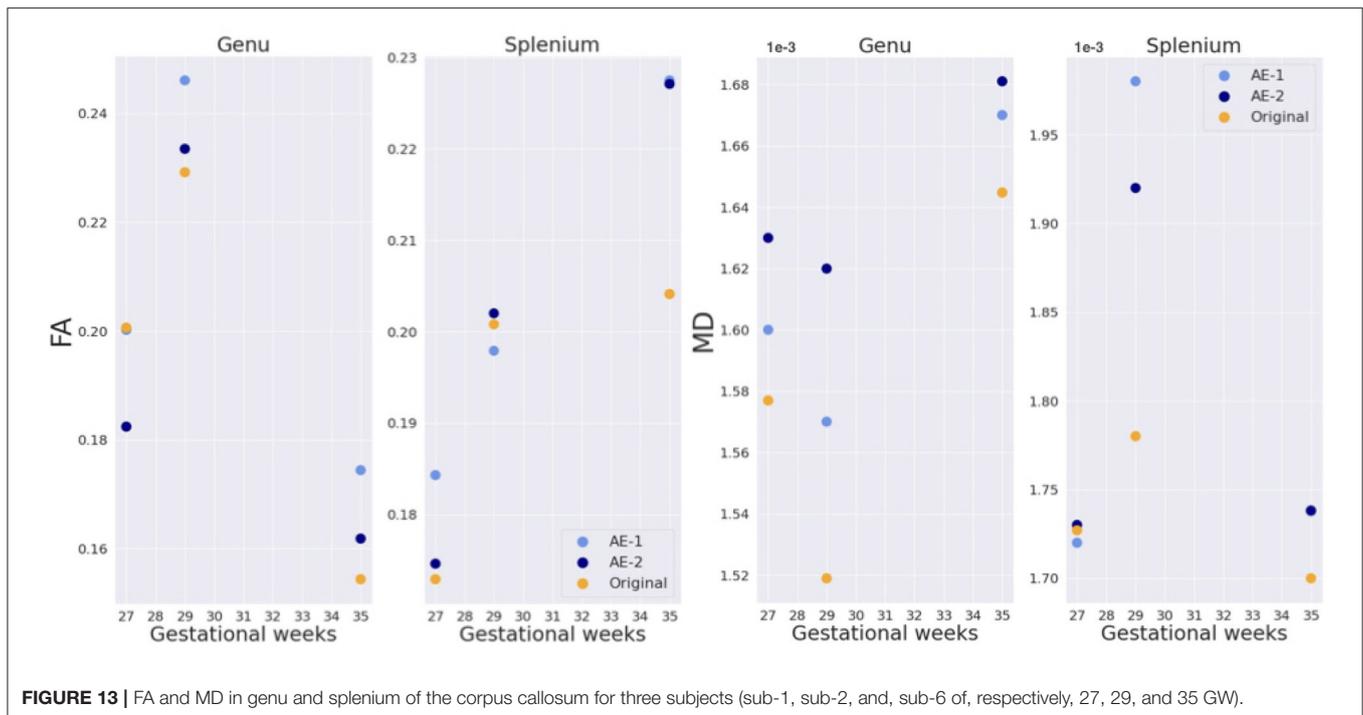


FIGURE 13 | FA and MD in genu and splenium of the corpus callosum for three subjects (sub-1, sub-2, and, sub-6 of, respectively, 27, 29, and 35 GW).

diffusion signal level and the downstream benefits propagated to the DTI maps.

Additionally, our network that was solely trained on non-diffusion weighted images (b_0) was able to generalize to a b_{1000} contrast. In fact, the most intuitive approach is to infer b_{1000} images using a network trained on b_{1000} . We have indeed tried but the network did not converge for the majority of the folds. This might be due to the high variability of b_{1000} images across directions and their inherently low SNR. However, in the 1-fold that the network converged, it slightly underperformed the network that was trained on b_0 only, on both b_{1000} pre-term and b_{700} fetal images. Moreover, being b -value independent is a desirable property since different b -values are used in different centers, in particular for clinical fetal imaging (400, 500, 600, 700 s/mm^2) (6, 10, 12, 29, 60). In fact, the same b_0 network trained on pre-term data was generalized to b_{700} fetal images where we qualitatively show its advantage, hence supporting the utility of pre-term data for fetal imaging, such as in Karimi et al. (61), where they have used pre-term colored FA and DW-MRI fetal scans to successfully predict fetal colored FA using a convolutional neural network. Furthermore, FA and MD of the corpus callosum, which were generated using the autoencoder enhanced volumes, are in the range of values provided by a recent study (59). This is a necessary but non sufficient condition for the validity of our framework in fetal data.

Notably, our trained network was able to reduce the noise from the data by learning the main features across images for different noise levels. This can be explained by two points. First, our autoencoder was exposed to different low levels of noise (as the dHCP data was already denoised) and hence the encoded features of the latent space are ought to be noise independent. Second, generative autoencoders intrinsically yield high SNR

outputs due to the desired smoothness property of the latent space (62).

The proposed framework could be applied to correct for anisotropic voxel sizes and can be used for slice outliers recovery in case of extreme motion artifacts for example. In fact, the artificially removed middle slices in our experiments can represent corrupted slices that may need to be discarded or replaced using interpolation (36–38). Our autoencoder can hence be used to recover these damaged slices using neighboring ones.

The power of our method compared to conventional interpolations resides in two points. First, the amount of data used to predict/interpolate the middle slice. While only two slices will be used in traditional interpolation approaches, our method will in addition take advantage of the thousands of slices to which the network has been exposed and from which the important features have been learned (without any supervision) in the training phase. Second, based on the manifold hypothesis, our method performs interpolations in the learned encoding space, which is closer to the intrinsic dimensionality of the data (63), and hence all samples from that space will be closer to the true distribution of the data compared to a naive interpolation in the pixel/voxel space.

Although our network performed quantitatively better than conventional interpolation methods in pre-term subjects, its output is usually smoother and hence exhibits lesser details. This is a well-known limitation of generative autoencoders, such as variational autoencoders, and the consequence of the desirable property of making the latent space smooth (62). Generative Adversarial Networks (64) can be an interesting alternative to overcome this issue. However, they have other drawbacks as being more unstable and less straightforward to train (65)

than autoencoders. But if trained properly, they can achieve competitive results.

In this work, qualitative results only were provided on fetal DW-MRI. We are limited by the lack of ground truth in this domain, hence our results are a proof of concept. The future release of the fetal dHCP dataset will be very valuable to further develop our framework and proceed to its quantitative assessment for fetal DW-MRI.

In future work, we want to add random Rician noise in the training phase to increase the network robustness and predictive power. We also want to extend the autoencoding to the angular domain by using spherical harmonics decomposition for each 4D voxel and hence enhancing both spatial and angular resolutions (66).

Although *unsupervised learning via* autoencoders has been recently used in DW-MRI to cluster individuals based on their microstructural properties (67), this is to the best of our knowledge, the first *unsupervised learning study* for super-resolution enhancement in DW-MRI using autoencoders.

As diffusion fetal imaging suffers from low through-plane resolution, super-resolution using autoencoders is an appealing method to artificially but realistically overcome this caveat. This can help depict more precise diffusion properties through different models, such as DTI or ODFs, and potentially increase the detectability of fiber tracts that are relevant for the assessment of certain neurodevelopmental disorders (29).

DATA AVAILABILITY STATEMENT

Part of the analyzed datasets were publicly available. This data can be found here: <http://www.developingconnectome.org/data-release/data-release-user-guide/>.

ETHICS STATEMENT

The studies involving human participants were reviewed and approved by the Cantonal Ethical Committee, Zürich. Written

informed consent to participate in this study was provided by the participants' legal guardian/next of kin.

AUTHOR CONTRIBUTIONS

HK performed the technical analysis, wrote the manuscript, provided original idea, and integrated all revisions. EC-R and HK contributed to the conceptualization of the research project. EC-R, PD, GG, AJ, and HL revised the manuscript. HL helped in the data generation. PD helped in the technical analysis. GG, MK, and YA-G helped in the processing of the fetal data. YA-G and MK acknowledged the manuscript. AJ provided the fetal data. MB conceptualized, designed and supervised the research project, contributed to the manuscript and to the final revision, and provided funding. All authors contributed to the article and approved the submitted version.

FUNDING

This work was supported by the Swiss National Science Foundation (regular project 205321-182602, Ambizione grant PZ00P2_185814, and NCCR SYNAPSY 185897). AJ was supported by the Prof. Max Cloetta Foundation, the EMDO Foundation, the Vontobel Foundation, the Novartis Foundation for Medical-Biological Research, and the URPP Adaptive Brain Circuits in Development and Learning (AdaBD) of the University of Zürich.

ACKNOWLEDGMENTS

We thank Athena Taymourtash for the help in the acceleration of the technical analysis and Samuel Lamon (MD) for the help in the segmentation of the corpus callosum.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fneur.2022.827816/full#supplementary-material>

REFERENCES

- Basser PJ, Pierpaoli C. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *J Magn Reson.* (2011) 213:560–70. doi: 10.1016/j.jmr.2011.09.022
- Johansen-Berg H, Behrens TE. *Diffusion MRI: From Quantitative Measurement to in vivo Neuroanatomy.* Cambridge, MA: Academic Press (2013).
- Bui T, Daire JL, Chalard F, Zaccaria I, Alberti C, Elmaleh M, et al. Microstructural development of human brain assessed in utero by diffusion tensor imaging. *Pediatr Radiol.* (2006) 36:1133–40. doi: 10.1007/s00247-006-0266-3
- Huang H, Xue R, Zhang J, Ren T, Richards LJ, Yarowsky P, et al. Anatomical characterization of human fetal brain development with diffusion tensor magnetic resonance imaging. *J Neurosci.* (2009) 29:4263–73. doi: 10.1523/JNEUROSCI.2769-08.2009
- Huang H, Vasung L. Gaining insight of fetal brain development with diffusion MRI and histology. *Int J Dev Neurosci.* (2014) 32:11–22. doi: 10.1016/j.ijdevneu.2013.06.005
- Jakab A, Kasprian G, Schwartz E, Gruber GM, Mitter C, Prayer D, et al. Disrupted developmental organization of the structural connectome in fetuses with corpus callosum agenesis. *Neuroimage.* (2015) 111:277–88. doi: 10.1016/j.neuroimage.2015.02.038
- Righini A, Bianchini E, Parazzini C, Gementi P, Ramenghi L, Baldoli C, et al. Apparent diffusion coefficient determination in normal fetal brain: a prenatal MR imaging study. *Am J Neuroradiol.* (2003) 24:799–804.
- Kasprian G, Brugger PC, Weber M, Krssák M, Krampl E, Herold C, et al. In utero tractography of fetal white matter development. *Neuroimage.* (2008) 43:213–24. doi: 10.1016/j.neuroimage.2008.07.026
- Khan S, Vasung L, Marami B, Rollins CK, Afacan O, Ortinau CM, et al. Fetal brain growth portrayed by a spatiotemporal diffusion tensor MRI

- atlas computed from in utero images. *Neuroimage*. (2019) 185:593–608. doi: 10.1016/j.neuroimage.2018.08.030
10. Deprez M, Price A, Christiaens D, Estrin GL, Cordero-Grande L, et al. Higher order spherical harmonics reconstruction of fetal diffusion MRI with intensity correction. *IEEE Trans Med Imaging*. (2019) 39:1104–13. doi: 10.1109/TMI.2019.2943565
 11. Kim K, Habas PA, Rousseau F, Glenn OA, Barkovich AJ, Koob M, et al. Reconstruction of a geometrically correct diffusion tensor image of a moving human fetal brain. In: *Medical Imaging 2010: Image Processing*. Vol. 7623. *International Society for Optics and Photonics*. San Diego, CA (2010). p. 76231I.
 12. Fogtmann M, Seshamani S, Kroenke C, Cheng X, Chapman T, Wilm J, et al. A unified approach to diffusion direction sensitive slice registration and 3-D DTI reconstruction from moving fetal brain anatomy. *IEEE Trans Med Imaging*. (2013) 33:272–89. doi: 10.1109/TMI.2013.2284014
 13. Marami B, Scherrer B, Afacan O, Erem B, Warfield SK, Gholipour A. Motion-robust diffusion-weighted brain MRI reconstruction through slice-level registration-based motion tracking. *IEEE Trans Med Imaging*. (2016) 35:2258–69. doi: 10.1109/TMI.2016.2555244
 14. Kuklisova-Murgasova M, Estrin GL, Nunes RG, Malik SJ, Rutherford MA, Rueckert D, et al. Distortion correction in fetal EPI using non-rigid registration with a Laplacian constraint. *IEEE Trans Med Imaging*. (2017) 37:12–9. doi: 10.1109/TMI.2017.2667227
 15. Rousseau F, Glenn OA, Iordanova B, Rodriguez-Carranza C, Vigneron DB, Barkovich JA, et al. Registration-based approach for reconstruction of high-resolution in utero fetal MR brain images. *Acad Radiol*. (2006) 13:1072–81. doi: 10.1016/j.acra.2006.05.003
 16. Gholipour A, Estroff JA, Warfield SK. Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI. *IEEE Trans Med Imaging*. (2010) 29:1739–58. doi: 10.1109/TMI.2010.2051680
 17. Kuklisova-Murgasova M, Quaghebeur G, Rutherford MA, Hajnal JV, Schnabel JA. Reconstruction of fetal brain MRI with intensity matching and complete outlier removal. *Med Image Anal*. (2012) 16:1550–64. doi: 10.1016/j.media.2012.07.004
 18. Tourbier S, Bresson X, Hagmann P, Thiran JP, Meuli R, Cuadra MB. An efficient total variation algorithm for super-resolution in fetal brain MRI with adaptive regularization. *Neuroimage*. (2015) 118:584–97. doi: 10.1016/j.neuroimage.2015.06.018
 19. Kainz B, Steinberger M, Wein W, Kuklisova-Murgasova M, Malamateniou C, Keraudren K, et al. Fast volume reconstruction from motion corrupted stacks of 2D slices. *IEEE Trans Med Imaging*. (2015) 34:1901–13. doi: 10.1109/TMI.2015.2415453
 20. Ebner M, Wang G, Li W, Aertsen M, Patel PA, Aughwane R, et al. An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain MRI. *Neuroimage*. (2020) 206:116324. doi: 10.1016/j.neuroimage.2019.116324
 21. Taymourtash A, Kebiri H, Tourbier S, Schwartz E, Nanning K-H, Licandro R, et al. 4D iterative reconstruction of brain fMRI in the moving fetus. *arXiv [Preprint]*. (2021). arXiv: 2111.11394. Available online at: <https://arxiv.org/pdf/2111.11394.pdf>
 22. Kebiri H, Lajous H, Alemán-Gómez Y, Girard G, Rodríguez EC, Tourbier S, et al. Quantitative evaluation of enhanced multi-plane clinical fetal diffusion MRI with a crossing-fiber phantom. In: *International Workshop on Computational Diffusion MRI*. Strasbourg: Springer (2021). p. 12–22.
 23. Kebiri H, Lajous H, Alemán-Gómez Y, Ledoux JB, Fornari E, Jakab A, et al. *Dataset: Quantitative Evaluation of Enhanced Multi-Plane Clinical Fetal Diffusion MRI With a Crossing-Fiber Phantom*. (2021). Available online at: https://zenodo.org/record/5153507#.YZ-uw_xKjRZ
 24. Ning L, Setsompop K, Michailovich O, Makris N, Shenton ME, Westin CF, et al. A joint compressed-sensing and super-resolution approach for very high-resolution diffusion imaging. *Neuroimage*. (2016) 125:386–400. doi: 10.1016/j.neuroimage.2015.10.061
 25. Liao C, Chen Y, Cao X, Chen S, He H, Mani M, et al. Efficient parallel reconstruction for high resolution multishot spiral diffusion data with low rank constraint. *Magn Reson Med*. (2017) 77:1359–66. doi: 10.1002/mrm.26199
 26. Mani M, Jacob M, Kelley D, Magnotta V. Multi-shot sensitivity-encoded diffusion data recovery using structured low-rank matrix completion (MUSSELS). *Magn Reson Med*. (2017) 78:494–507. doi: 10.1002/mrm.26382
 27. Ramos-Llordén G, Ning L, Liao C, Mukhometzianov R, Michailovich O, Setsompop K, et al. High-fidelity, accelerated whole-brain submillimeter in vivo diffusion MRI using gSlider-spherical ridgelets (gSlider-SR). *Magn Reson Med*. (2020) 84:1781–95. doi: 10.1002/mrm.28232
 28. Dyrby TB, Lundell H, Burke MW, Reislev NL, Paulson OB, Ptito M, et al. Interpolation of diffusion weighted imaging datasets. *Neuroimage*. (2014) 103:202–13. doi: 10.1016/j.neuroimage.2014.09.005
 29. Jakab A, Tuura R, Kellenberger C, Scheer I. In utero diffusion tensor imaging of the fetal brain: a reproducibility study. *Neuroimage Clin*. (2017) 15:601–12. doi: 10.1101/132704
 30. Coupé P, Manjón JV, Chamberland M, Descoteaux M, Hiba B. Collaborative patch-based super-resolution for diffusion-weighted images. *Neuroimage*. (2013) 83:245–61. doi: 10.1016/j.neuroimage.2013.06.030
 31. Alexander DC, Zikic D, Ghosh A, Tanno R, Wotschel V, Zhang J, et al. Image quality transfer and applications in diffusion MRI. *Neuroimage*. (2017) 152:283–98. doi: 10.1016/j.neuroimage.2017.02.089
 32. Elsaid NM, Wu YC. Super-resolution diffusion tensor imaging using SRCNN: a feasibility study. *Annu Int Conf IEEE Eng Med Biol Soc*. (2019) 2019:2830–4. doi: 10.1109/EMBC.2019.8857125
 33. Dong C, Loy CC, He K, Tang X. Image super-resolution using deep convolutional networks. *IEEE Trans Pattern Anal Mach Intell*. (2015) 38:295–307. doi: 10.1109/TPAMI.2015.2439281
 34. Chatterjee S, Sciarra A, Dinnwald M, Mushunuri RV, Podishetti R, Rao RN, et al. ShuffleUNet: super resolution of diffusion-weighted MRIs using deep learning. *arXiv [Preprint]*. arXiv:210212898. (2021) doi: 10.23919/EUSIPCO54536.2021.9615963
 35. Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Munich: Springer (2015). p. 234–41.
 36. Niethammer M, Bouix S, Aja-Fernández S, Westin CF, Shenton ME. Outlier rejection for diffusion weighted imaging. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Brisbane, QLD: Springer (2007). p. 161–8.
 37. Chang LC, Jones DK, Pierpaoli C. RESTORE: robust estimation of tensors by outlier rejection. *Magn Reson Med*. (2005) 53:1088–95. doi: 10.1002/mrm.20426
 38. Andersson JL, Graham MS, Zsoldos E, Sotiropoulos SN. Incorporating outlier detection and replacement into a non-parametric framework for movement and distortion correction of diffusion MR images. *Neuroimage*. (2016) 141:556–72. doi: 10.1016/j.neuroimage.2016.06.058
 39. Sander J, de Vos BD, Išgum I. Unsupervised super-resolution: creating high-resolution medical images from low-resolution anisotropic examples. In: *Medical Imaging 2021: Image Processing*. vol. 11596. *International Society for Optics Photonics*. San Diego, CA (2021). p. 115960E.
 40. Chung H, Kim J, Yoon JH, Lee JM, Ye JC. Simultaneous super-resolution and motion artifact removal in diffusion-weighted MRI using unsupervised deep learning. *arXiv [Preprint]*. arXiv:210500240. (2021). doi: 10.48550/arXiv.2105.00240
 41. Hutter J, Tournier JD, Price AN, Cordero-Grande L, Hughes EJ, Malik S, et al. Time-efficient and flexible design of optimized multishell HARDI diffusion. *Magn Reson Med*. (2018) 79:1276–92. doi: 10.1002/mrm.26765
 42. Bastiani M, Andersson JL, Cordero-Grande L, Murgasova M, Hutter J, Price AN, et al. Automated processing pipeline for neonatal diffusion MRI in the developing Human Connectome Project. *Neuroimage*. (2019) 185:750–63. doi: 10.1016/j.neuroimage.2018.05.064
 43. Makropoulos A, Robinson EC, Schuh A, Wright R, Fitzgibbon S, Bozek J, et al. The developing human connectome project: a minimal processing pipeline for neonatal cortical surface reconstruction. *Neuroimage*. (2018) 173:88–112. doi: 10.1016/j.neuroimage.2018.01.054
 44. Makropoulos A, Gousias IS, Ledig C, Aljabar P, Serag A, Hajnal JV, et al. Automatic whole brain MRI segmentation of the developing neonatal brain. *IEEE Trans Med Imaging*. (2014) 33:1818–31. doi: 10.1109/TMI.2014.2322280
 45. Salehi SSM, Hashemi SR, Velasco-Annis C, Ouaalam A, Estroff JA, Erdogmus D, et al. Real-time automatic fetal brain extraction in fetal MRI by deep

- learning. In: *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. Washington, DC: IEEE (2018). p. 720–4.
46. Gorgolewski K, Burns CD, Madison C, Clark D, Halchenko YO, Waskom ML, et al. Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in python. *Front Neuroinform.* (2011) 5:13. doi: 10.3389/fninf.2011.00013
 47. Veraart J, Fieremans E, Novikov DS. Diffusion MRI noise mapping using random matrix theory. *Magn Reson Med.* (2016) 76:1582–93. doi: 10.1002/mrm.26059
 48. Tustison NJ, Avants BB, Cook PA, Zheng Y, Egan A, Yushkevich PA, et al. N4ITK: improved N3 bias correction. *IEEE Trans Med Imaging.* (2010) 29:1310–20. doi: 10.1109/TMI.2010.2046908
 49. Avants BB, Tustison N, Song G. et al. Advanced normalization tools (ANTS). *Insight J.* (2009) 2:1–35. doi: 10.54294/uvnhin
 50. Modat M, Cash DM, Daga P, Winston GP, Duncan JS, Ourselin S. Global image registration using a symmetric block-matching approach. *J Med Imaging.* (2014) 1:024003. doi: 10.1117/1.JMI.1.2.024003
 51. Ioffe S, Szegedy C. Batch normalization: accelerating deep network training by reducing internal covariate shift. In: *International Conference on Machine Learning*. Lille: PMLR (2015). p. 448–56.
 52. Chollet F, et al. keras. (2015).
 53. Glorot X, Bengio Y. Understanding the difficulty of training deep feedforward neural networks. In: *Proceedings of the thirteenth international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings*. Sardinia (2010). p. 249–56.
 54. Kingma DP, Ba J. Adam: a method for stochastic optimization. *arXiv [Preprint]* arXiv:1412.6980. (2014). doi: 10.48550/arXiv.1412.6980
 55. Abadi M, Barham P, Chen J, Chen Z, Davis A, Dean J, et al. Tensorflow: a system for large-scale machine learning. In: *12th USENIX Symposium on Operating Systems Design and Implementation (SOSDI\$ 16)*. Savannah, GA (2016). p. 265–83.
 56. Tournier JD, Calamante F, Connelly A. MRtrix: diffusion tractography in crossing fiber regions. *Int J Imaging Syst Technol.* (2012) 22:53–66. doi: 10.1002/ima.22005
 57. Gudbjartsson H, Patz S. The Rician distribution of noisy MRI data. *Magn Reson Med.* (1995) 34:910–4. doi: 10.1002/mrm.1910340618
 58. Garyfallidis E, Brett M, Amirbekian B, Rokem A, Van Der Walt S, Descoteaux M, et al. Dipy, a library for the analysis of diffusion MRI data. *Front Neuroinform.* (2014) 8:8. doi: 10.3389/fninf.2014.00008
 59. Wilson S, Pietsch M, Cordero-Grande L, Price AN, Hutter J, Xiao J, et al. Development of human white matter pathways in utero over the second and third trimester. *Proc Natl Acad Sci USA.* (2021) 118:e2023598118. doi: 10.1073/pnas.2023598118
 60. Marami B, Salehi SSM, Afacan O, Scherrer B, Rollins CK, Yang E, et al. Temporal slice registration and robust diffusion-tensor reconstruction for improved fetal brain structural connectivity analysis. *Neuroimage.* (2017) 156:475–88. doi: 10.1016/j.neuroimage.2017.04.033
 61. Karimi D, Jaimes C, Machado-Rivas F, Vasung L, Khan S, Warfield SK, et al. Deep learning-based parameter estimation in fetal diffusion-weighted MRI. *Neuroimage.* (2021) 243:118482. doi: 10.1016/j.neuroimage.2021.118482
 62. Berthelot D, Raffel C, Roy A, Goodfellow I. Understanding and improving interpolation in autoencoders via an adversarial regularizer. *arXiv [preprint]*. arXiv:180707543. (2018). doi: 10.48550/arXiv.1807.07543
 63. Chollet F. *Deep learning With Python*. Simon and Schuster (2017).
 64. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, et al. Generative adversarial nets. *Adv Neural Inf Process Syst.* (2014) 27.
 65. Mescheder L, Geiger A, Nowozin S. Which training methods for GANs do actually converge? In: *International Conference on Machine Learning*. Stockholm: PMLR (2018). p. 3481–90.
 66. Ma J, Cui H. Hybrid graph convolutional neural networks for super resolution of DW images. In: *Computational Diffusion MRI*. Strasbourg: Springer (2021). p. 201–12.
 67. Chamberland M, Genc S, Tax CM, Shastin D, Koller K, Raven EP, et al. Detecting microstructural deviations in individuals with deep diffusion MRI tractometry. *Nat Computat Sci.* (2021) 1:598–606. doi: 10.1038/s43588-021-00126-8
- Conflict of Interest:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.
- Publisher's Note:** All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.
- Copyright © 2022 Kebiri, Canales-Rodríguez, Lajous, de Dumast, Girard, Alemán-Gómez, Koob, Jakob and Bach Cuadra. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.