



Retrospective motion estimation for fetal brain MRI

Chloé Mercier, Sylvain Faisan, Alexandre Pron, Nadine Girard, Guillaume Auzias, Thierry Chonavel, François Rousseau

► To cite this version:

Chloé Mercier, Sylvain Faisan, Alexandre Pron, Nadine Girard, Guillaume Auzias, et al.. Retrospective motion estimation for fetal brain MRI. IPTA 2023: International Conference on Image Processing Theory, Oct 2023, Paris, France. hal-04161325

HAL Id: hal-04161325

<https://hal.science/hal-04161325>

Submitted on 25 Aug 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Retrospective motion estimation for fetal brain MRI

Chloe Mericer¹, Sylvain Faisan², Alexandre Pron³, Nadine Girard⁴, Guillaume Auzias³,
Thierry Chonavel¹, François Rousseau⁵

¹ IMT Atlantique, Lab-STICC UMR 6285 CNRS, Brest

² ICube, UMR 7357, Université de Strasbourg, CNRS, Strasbourg,

³ Aix-Marseille Université, CNRS, Institut de Neurosciences de la Timone, Marseille

⁴ IMT Atlantique, LaTIM U1101 INSERM, Brest

July 2023

Abstract

Magnetic Resonance Imaging (MRI) is performed on fetus, to study cerebral development. However, the image is affected by child's and mother's motion. Ultra-fast MRI sequences allow to reduce artefact in slices, but motion still occurs between them. We propose a method based on orthogonal slices intersection to correct motion as well as a new solution for outliers removal.

Key Words : Fetal Brain, Registration, Magnetic Resonance Imaging

1 Introduction

Magnetic resonance imaging (MRI) has been used to study the developing fetal brain since the 1980s. However, the motion (of the mother and the fetus) has always been a real challenge, limiting the exploratory power of the acquired images. In the context of antenatal imaging, the full image of the brain is actually a stack of 2D slices. Those acquisitions are generally made along the three axes of space in order to provide the radiologist with a 3D "vision" of the brain. The acquisition time of a slice is generally sufficiently short (less than 1 second) to "freeze" the motion. As a result, the subject's motion induces mainly geometric distortion artifacts, i.e. the stack of 2D slices does not directly reflect the 3D geometry of the brain. Thus, it is necessary to estimate the motion retrospectively in order to reconstruct a 3D image of the fetal brain [1].

The main methods of reconstruction of fetal data, called SVR for "slice-to-volume registration", are based on two steps: estimation of the relative motion followed by fusion of the data [2–4]. In case of antenatal imaging, the registration problem is of the 2D-3D type, i.e. we have to estimate the motion between slices and a reference volume. This reference volume is also the image that we want to reconstruct and is therefore unknown. From a first estimation of the reference volume, the alignment of each slices is estimated by minimizing a criterion of alignment between the current slice and the reference volume. The latter is then recalculated from the set of transformations estimated for each slice. The quality of the reconstructed volume depends strongly on the quality of the registration of the slices. This process is repeated in an iterative way until convergence of the algorithm. In order to make these approaches more robust to the subject's motion, deep learning methods have been developed [5, 6]. However, methods based on iterative reconstruction remain insufficiently robust for the analysis of large databases of images acquired in clinical routine. Thus, it is necessary to detect misaligned slices in order not to include them in the reconstruction step [7, 8] or to reduce their influence in the reconstruction [9].

To solve this problem, one solution is to completely separate motion correction of the slices from the reconstruction step by using the intersection of the orthogonal slices and imposing their correspondence to the 3D intersections [10]. This approach allows for solving independently the problem of correcting the motion of the slices and reconstruction of a 3D volume. In this paper, we develop a method for estimating the motion of slices associated with a detection of misregistered slices using a machine learning approach. The proposed method is called ROSI, for "Registration based on Orthogonal Slices Intersection". The evaluation conducted on synthetic and real data shows the interest of the proposed approach compared to SVR methods.

2 Materials and Methods

2.1 Motion estimation

We consider three images, namely I_1 , I_2 , and I_3 (the algorithm can be extended to accommodate more images if available). Each image is composed of a stack of slices, with each slice having a unique identifier denoted as

k . Moreover, $f(k)$ represents the index (i.e., 1, 2, or 3) of the image from which slice k originates. Additionally, brain masks have been automatically segmented in the images [11], associating each k with a corresponding mask denoted as m_k . It is worth noting that the intensities of each image I_1 , I_2 , and I_3 have been normalized using z-score normalization. To this end, the mean and standard deviation of each image have been computed based solely on the brain mask.

Slices associated with I_1 , I_2 , and I_3 will be considered aligned when their intensities coincide along their intersections. This is achieved by estimating, for each slice k , a rigid transformation (in homogeneous coordinates) denoted as M_k , which aligns all slices into a common reference system. This transformation can be formally defined as follows:

$$M_k = T(c_k)M(\Theta_k, t_k)T(-c_k)R_{f(k)}R_{k,2t3d} \quad (1)$$

In (1), $R_{k,2t3d}$ represents a rigid transformation that converts the 2-D coordinates of slice k into the 3-D voxel coordinates of the image $I_{f(k)}$. Next, $R_{f(k)}$ converts the voxel coordinates into the world reference system (in millimeters). The matrix to be estimated is denoted as $M(\Theta_k, t_k)$. There are 6 parameters to estimate: 3 for rotation (Θ_k) and 3 for translation (t_k). To rotate around the point with coordinates c_k , the process involves first performing a translation of $-c_k$ (denoted as $T(-c_k)$), then applying $M(\Theta_k, t_k)$, and finally performing the inverse translation. Note that c_k is equal to $R_{f(k)}R_{k,2t3d}(x_k)$, where x_k represents the 2-D coordinates of the centroid of the mask associated with slice k .

The registration criterion used to align images is based on the principle, as described in [10], that well-registered slices should exhibit identical intensity profiles along their intersection (see Fig. 1).




Figure 1: Visualisation of the intersection between two orthogonal slices and their corresponding profiles. Two misregistered slices (a) have distinct profiles along their intersection (b) whereas two well-registered slices (c) display similar profiles (d).

The criterion is calculated as follows. Firstly, let us consider two slices from distinct images, denoted as k and k' . Assuming the intersection between these slices exists, it takes the form of a line. The equation of this line is determined in the reference system using the transformations M_k and $M_{k'}$. By applying the inverse transformations, we can derive the equations of the corresponding line in slice k and slice k' . The two lines are sampled regularly (every millimeter) so that each point sampled in the slice k has its "corresponding point" in the slice k' and vice versa. In the following, it is assumed that point v in slice k corresponds to point v' in slice k' , i.e., $M_k(v) = M_{k'}(v')$. Moreover, we ensure that the points v (resp. v') sample all structures of interest that are on the segment of the slice k (resp. k') (thanks to m_k and $m_{k'}$). Consequently, if the slices are misaligned, it is possible that v (or v') may not fall within the support of slice k (or slice k'). To address this issue, we assume that the slices have a value of zero outside their respective supports. Finally, for a pair (k, k') , we compute:

$$S^2(k, k') = \sum_v (s_k(v) - s_{k'}(v'))^2 \mathbb{1}_{m_k(v)=1 \text{ or } m_{k'}(v')=1} \quad (2)$$

$$N(k, k') = \sum_v \mathbb{1}_{m_k(v)=1 \text{ or } m_{k'}(v')=1}, \quad (3)$$

where $\mathbb{1}_E = \begin{cases} 1 & \text{if } E \text{ occurs} \\ 0 & \text{otherwise} \end{cases}$, $s_k(v)$ (resp. $s_{k'}(v')$) represents the intensity of slice k (resp. k') at pixel v (resp. v'), and where $m_k(v)$ (resp. $m_{k'}(v')$) represents the value of the mask k (resp. k') at pixel v (resp. v'). The mask value is 1 for structures of interest, 0 otherwise. Finally, the criterion to be optimized is the following:

$$\frac{\sum_{k,k'; k > k'} S^2(k, k')}{\sum_{k,k'; k > k'} N(k, k')}. \quad (4)$$

The motion between slices is not rigid in maternal tissues. Since the criterion is computed on the union of masks, the presence of maternal motion around the brain can impact the registration accuracy. Therefore, similar to [10], the intensity of voxels outside the mask is reduced to get closer to zero (this is not shown in Eq. 2 and 3).

2.2 Implementation

The criterion is optimized using an alternating block optimization method, where parameters associated with slice k (Θ_k and t_k) are updated while parameters associated with other slices remain unchanged. It is worth noting that the computation of the criterion in (4) can benefit from the fact for each update that most parameters remain constant for each update. Finally, the parameter optimization for a slice is performed using the Nelder-Mead method.

Algorithm 1 provides a detailed description of the proposed optimization scheme. In summary, this is an iterative optimization process where parameters for each slice are optimized until their changes fall below a certain threshold T (referred to as "local convergence"). The algorithm repeats this process until all slices achieve local convergence within a single application of the Nelder-Mead method.

Algorithm 1 Alternating block optimization procedure

Input : $A = \{1, \dots, n\}$, n the number of slices.

Output: Θ_k and t_k for each slice k .

```

while global convergence has not been achieved do
  Set  $B$  to  $A$ 
  while  $B$  is not empty do
    for each slice  $k$  in  $B$  do
      Update:  $(\Theta_k^{old}, t_k^{old}) \rightarrow (\Theta_k^{new}, t_k^{new})$  with Nelder Mead method
       $x = ||(\Theta_k^{old}, t_k^{old}) - (\Theta_k^{new}, t_k^{new})||^2$ 
      if  $x < T$  then ▷ local convergence is achieved for slice  $k$ 
        remove  $k$  from  $B$ 
      else
        global convergence is not achieved
      end if
    end for
  end while
end while

```

Three hyperparameters may influence the results : the initial simplex size (ds), the final simplex size (fs) and the threshold T . Initially, we choose, $ds = 4$, $fs = 0.25$ and $T = 25$. We conducted experiments using different values for ds , fs , and T , and observed that their choice does not significantly impact the final result across a wide range of values.

2.3 Rejection of misregistered slices

Inaccurate estimation of slice motion can result in errors in volume reconstruction. To address this issue, we propose a new strategy for detecting misregistered slices. Unlike existing methods in the literature [8, 10], which rely on a single feature threshold, we use a classification algorithm (random forest) based on five different features.

Using the notation of (4), we compute the mean squared deviation (MSD) associated with each slice k as follows:

$$\frac{\sum_{k'} S^2(k, k')}{\sum_{k'} N(k, k')}. \quad (5)$$

However, the presence of noise in the data acts as a nuisance variable. For instance, even if the slice is well-aligned, significant noise can result in a high MSD. Consequently, we normalize the *MSD* as follows:

$$\sum_{i=1, i \neq f(k)}^3 \frac{1}{\sigma_{f(k)}^2 + \sigma_i^2} \frac{\sum_{k'; f(k')=i} S^2(k, k')}{\sum_{k'; f(k')=i} N(k, k')}, \quad (6)$$

where $\sigma_{f(k)}$ (resp. σ_i) represents the standard deviation of the noise that corrupts image $I_{f(k)}$ (resp. I_i). The standard deviation $\sigma_{f(k)}$ (resp. σ_i) is estimated using the strategy proposed in [12] on the structures of interest defined by the related mask.

Additionally, we calculate a Dice index for each slice k to assess the alignment of the masks. The Dice index is defined as $2i/s$ where:

$$i = \sum_{k'} M(k, k'), \quad s = \sum_{k'} P(k, k') + Q(k, k'), \quad (7)$$

where

$$M(k, k') = \sum_v \mathbb{1}_{m_k(v)=1 \text{ and } m_{k'}(v)=1}, \quad (8)$$

and where

$$P(k, k') = \sum_v \mathbb{1}_{m_k(v)=1}, \quad Q(k, k') = \sum_v \mathbb{1}_{m_{k'}(v)=1}. \quad (9)$$

Note that $P(k, k')$ is a function of k' and $Q(k, k')$ a function of k because both k and k' are necessary to compute the intersection.

The volume of the mask can be considered as a confounding factor as it also influences the dice index. For example, if the registration is satisfactory, a larger mask volume for a slice k would result in a larger associated Dice index.

We employed two different strategies to address the problem. First, we incorporated the confounding variable (mask volume) into the features. Specifically, we divided the mask volume of slice k by the mask volume of the reference slice k_{max} , where k_{max} corresponds to the slice with the largest mask, subject to the constraint $f(k) = f(k_{max})$. Finally, we introduced an additional feature that serves the same purpose as the Dice index but is not influenced by the mask volume. This index, referred to as Diff, is calculated as follows:

$$\frac{(2i - s)}{n_k}, \quad (10)$$

where n_k represents the number of slices that intersect with k . Finally, we also consider the standard deviation of the slice intensity as a feature.

3 Results

3.1 Data

Due to the lack of ground truth in fetal MRI, our algorithm was evaluated using synthetic data generated from the DHCP [13] database. Synthetic data were created from a 3D image by simulating three orthogonal stacks of slices (axial, sagittal, and coronal) with a thickness of 3 mm and a resolution of 0.5x0.5 mm. Data are generated with three different levels of motion, expressed in degrees for rotations and in millimeters for translations: low ($\Theta_k \in [-1, 1]$, $t_k \in [-1, 1]$), medium ($\Theta_k \in [-3, 3]$, $t_k \in [-3, 3]$), and large ($\Theta_k \in [-5, 5]$, $t_k \in [-5, 5]$). As described in [14], the *point spread function* (PSF) is set to a Gaussian function, with width equal to the slice thickness. In total, we generate five data sets (using different 3D images) for each level of motion. An example of simulated data is given in Fig. 2.

In addition to synthetic data, MRI from 22 normal fetuses were included in this study (8 M, 5 F, 9 Unknown) covering the 26th until the 36th week of gestation (mean=32.85, sigma=2). The retrospective use of MRI data acquired in clinical routine at the hospital of La Timone was approved by the Institutional Review Board of Aix-Marseille Université (ref 2022-04-14-003). All acquisitions were ultra fast spin echo (HASTE) sequences acquired on two different Siemens scanners (Skyra (3T) n=10 (images) and SymphonyTim (1.5T) n=12 (images). The absence of abnormal neuroradiological findings was assessed for all subjects by an expert neuroradiologist (N.G.). Fig. 3 displays real data acquired on a fetus in the axial orientation.




Figure 2: Example of simulated data depicting the acquisition in the axial orientation. The data is generated with varying levels of motion and visualized in the sagittal view.




Figure 3: Example of real data acquired in the axial orientation, displayed in axial, coronal, and sagittal views.

3.2 Evaluation of Registration

The quality of the registration is assessed using the target registration error (TRE). We denote by M_k the transformation applied to slice k , and by \hat{M}_k its estimated counterpart. Given two intersecting slices k and k' , we can compute, as described in section 2.1, points v (from slice k) and v' (from slice k') such that $M_k(v) = M_{k'}(v')$. We consider only points v and v' that belong to the masks ($m_k(v) = m_{k'}(v') = 1$). The target registration error (TRE) is the distance in millimeters in the reference system between $\hat{M}_k(v)$ and $\hat{M}_{k'}(v')$.

For a slice k , we compute an average distance defined by:

$$TRE_k = \frac{1}{\#\{(k', v')\}} \sum_{k'} \sum_{v'} \|\hat{M}_k(v) - \hat{M}_{k'}(v')\| \quad (11)$$

When the slice is well aligned, TRE_k is small.

Results are compared with Nifty-MIC [8]. First, the quality and accuracy of the registration are evaluated without rejecting potentially misregistered slices. Fig. 4 shows for both methods the average TRE per slice (TRE_k) before and after registration for each level of motion simulated. After registration, the TRE_k values obtained with ROSI are substantially lower compared to those obtained with NiftyMIC. For ROSI, the majority of slices achieve a TRE_k below 1.5 mm after registration, regardless of the initial motion. However, it is worth noting that a few slices still exhibit poor registration, with TRE_k values exceeding 1.5 mm.

The accuracy of the motion estimation can also be validated by assessing the quality of the reconstruction. This is achieved by calculating the peak signal-to-noise ratio ($PSNR$) and the structural similarity ($SSIM$). Note that we use the same reconstruction method as NiftyMIC. For each level of motion, the $PSNR$ and $SSIM$ values are computed for the five reconstructed images, and then averaged. Tab. 1 presents the mean $PSNR$ obtained for both methods. The mean $PSNR$ obtained with ROSI is higher for all three levels of motion. As motion increases, the $PSNR$ decreases, but the decrease is less pronounced compared to what is observed with NiftyMIC. The same trend is observed for the mean $SSIM$ values (Tab. 2).

Table 1: mean PSNR (dB) obtained on 3-D reconstruction from simulated data. Standard deviation is given in parenthesis.

Motion	NiftyMIC	ROSI
Low	30.32 (0.79)	33.37 (0.78)
Medium	25.08 (0.63)	32.25 (0.60)
Large	23.57 (0.29)	31.86 (0.69)

Finally, Fig. 5 presents 3D images reconstructed from the simulated data, in axial view. Our method exhibits sharper and more precise contours compared to NiftyMIC, regardless of the level of motion.




Figure 4: TRE_k (in millimeters) estimated before (on ordinate axis) and after (on abscisse axis) registration, with ROSI (in black) and NiftyMIC (in grey)

Table 2: mean SSIM obtained on 3-D reconstruction from simulated data. Standard deviation is given in parenthesis.

Motion	NityMIC	ROSI
Low	0.93 (0.01)	0.96 (0.01)
Medium	0.81 (0.04)	0.96 (0.01)
Large	0.75 (0.05)	0.95 (0.01)




Figure 5: Example of 3D reconstruction obtained from simulated data with low motion (first row), medium motion (second row) and large motion (third row). Data were reconstructed with NiftyMIC (second column). For ROSI (third column), motion between slices was corrected using ROSI and a 3D image was reconstructed with the same approach as NiftyMIC.

3.3 Rejection of misregistered slices

To increase the number of misregistered slices, which is crucial for training, we expanded the range of simulated motion. Specifically, we set the motion range to $\Theta_k \in [-8.8]$ and $T_k \in [-8.8]$. Additionally, Gaussian

random noise was introduced into data, with distinct standard deviations for each image I_1 , I_2 , and I_3 . Slices with a TRE_k greater than 1.5 mm are classified as misregistered, and otherwise considered as well registered. Approximately 10% of the slices were identified as misregistered. The classification algorithm was trained using 20 registration results and tested on an independent set of 10 results.

The performance of the proposed method is compared to state-of-the-art approaches, including the SLIMMER method [10] and NiftyMIC [8]. SLIMMER uses the mean squared deviation (MSD) feature and sets the threshold to 1.25 times the median MSD value. NiftyMIC employs the normalized cross-correlation (NCC) as a feature with a threshold of 0.8. Fig. 6 depicts the confusion matrix obtained with the different methods. Additionally, we compare our results with a classifier that uses only the features MSD , $DICE$, or $DIFF$.




Figure 6: Confusion matrix for the detection of misregistered slices obtained with ROSI (a), SLIMMER (b), and NiftyMIC (c).




Figure 7: Confusion matrix for the detection of misregistered slices obtained with a classifier that uses only one feature: $DICE$ (a), MSE (b), and $DIFF$ (c).

Results demonstrate that the classifier trained on all five features outperforms other methods. As shown in Figure 7, using a single feature leads to poor detection of misregistered slices.

3.4 Results on real data

The accuracy of the motion estimation is validated through an evaluation of the reconstruction quality. It is important to note that as previously, we employ, the same reconstruction method as NiftyMIC for this purpose.

Tests conducted on real data have revealed that the results obtained with ROSI are generally comparable to those obtained with NiftyMIC. Furthermore, it is worth noting that ROSI is capable of reconstructing some data that NiftyMIC fails to reconstruct. Fig. 8 illustrates two examples of data that failed to be reconstructed by NiftyMIC but was successfully reconstructed using ROSI. In panels (c) and (d), we observe holes in the reconstruction obtained with NiftyMIC. However, these holes disappear when using the proposed method.

4 Conclusion

We have proposed a novel registration method specifically designed for fetal MRI, along with a new approach for rejecting poorly registered slices. Our method is based on the intersection of orthogonal slices, making independent the registration step and the reconstruction one. In synthetic data evaluations, our method outperformed state-of-the-art approaches. On real data, our method successfully registered data that state of the art NiftyMIC method was unable to handle.

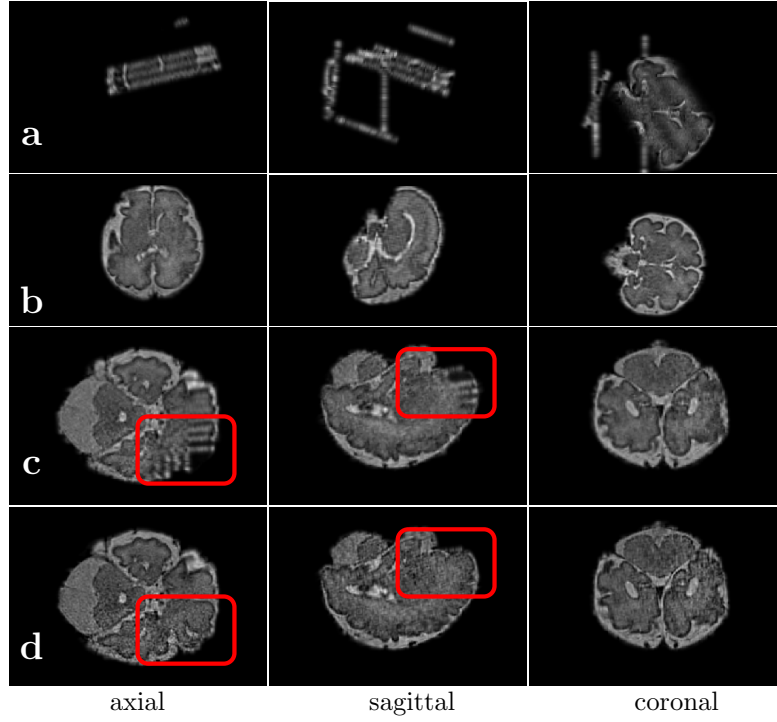


Figure 8: Example of 3D reconstruction obtained on real data. (a) and (b) represent reconstructed data, registered with NiftyMIC (a) and ROSI (b). The reconstruction failed with NiftyMIC but performs well with ROSI. (c) and (d) represent reconstructed data, registered with NiftyMIC (c) and ROSI (d). The red boxes highlight the parts of the images where the reconstruction is incorrect using NiftyMIC.

Acknowledgements: This work was supported by the French National Research Agency (projects ANR-19-CHIA-0015, ANR-19-CE45-0014, ANR-21-NEU2-0005).

References

- [1] C. Studholme and F. Rousseau, “Quantifying and modelling tissue maturation in the living human fetal brain,” *International Journal of Developmental Neuroscience*, vol. 32, pp. 3–10, 2014.
- [2] F. Rousseau, O. A. Glenn, B. Iordanova, C. Rodriguez-Carranza, D. B. Vigneron, J. A. Barkovich, and C. Studholme, “Registration-based approach for reconstruction of high-resolution in utero fetal MR brain images,” *Academic radiology*, vol. 13, no. 9, pp. 1072–1081, 2006.
- [3] A. Gholipour, J. A. Estroff, and S. K. Warfield, “Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI,” *IEEE transactions on medical imaging*, vol. 29, no. 10, pp. 1739–1758, 2010.
- [4] B. Kainz, M. Steinberger, W. Wein, M. Kuklisova-Murgasova, C. Malamateniou, K. Keraudren, T. Torsney-Weir, M. Rutherford, P. Aljabar, J. V. Hajnal *et al.*, “Fast volume reconstruction from motion corrupted stacks of 2d slices,” *IEEE transactions on medical imaging*, vol. 34, no. 9, pp. 1901–1913, 2015.
- [5] J. Xu, D. Moyer, B. Gagoski, J. E. Iglesias, P. E. Grant, P. Golland, and E. Adalsteinsson, “Nesvor: Implicit neural representation for slice-to-volume reconstruction in MRI,” *IEEE Transactions on Medical Imaging*, 2023.
- [6] W. Shi, H. Xu, C. Sun, J. Sun, Y. Li, X. Xu, T. Zheng, Y. Zhang, G. Wang, and D. Wu, “Affirm: Affinity fusion-based framework for iteratively random motion correction of multi-slice fetal brain MRI,” *IEEE Transactions on Medical Imaging*, 2022.
- [7] M. Kuklisova-Murgasova, G. Quaghebeur, M. A. Rutherford, J. V. Hajnal, and J. A. Schnabel, “Reconstruction of fetal brain MRI with intensity matching and complete outlier removal,” *Medical image analysis*, vol. 16, no. 8, pp. 1550–1564, 2012.
- [8] M. Ebner, G. Wang, W. Li, M. Aertsen, P. A. Patel, R. Aughwane, A. Melbourne, T. Doel, S. Dymarkowski, P. De Coppi *et al.*, “An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain MRI,” *NeuroImage*, vol. 206, p. 116324, 2020.

- [9] S. Tourbier, X. Bresson, P. Hagmann, J.-P. Thiran, R. Meuli, and M. B. Cuadra, “An efficient total variation algorithm for super-resolution in fetal brain MRI with adaptive regularization,” *NeuroImage*, vol. 118, pp. 584–597, 2015.
- [10] K. Kim, P. A. Habas, F. Rousseau, O. A. Glenn, A. J. Barkovich, and C. Studholme, “Intersection based motion correction of multislice MRI for 3-d in utero fetal brain image formation,” *IEEE transactions on medical imaging*, vol. 29, no. 1, pp. 146–158, 2009.
- [11] M. Ebner, G. Wang, W. Li, M. Aertsen, P. A. Patel, R. Aughwane, A. Melbourne, T. Doel, A. L. David, J. Deprest *et al.*, “An automated localization, segmentation and reconstruction framework for fetal brain MRI,” in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part I*. Springer, 2018, pp. 313–320.
- [12] S. Faisan, C. Heinrich, F. Rousseau, A. Lallement, and J. Zallat, “Joint filtering-estimation of stokes vector images based on a non-local means approach,” *Journal of the Optical Society of America A*, vol. 29, no. 9, 2012.
- [13] A. D. Edwards, D. Rueckert, S. M. Smith, S. Abo Seada, A. Alansary, J. Almalbis, J. Allsop, J. Andersson, T. Arichi, S. Arulkumaran *et al.*, “The developing human connectome project neonatal data release,” *Frontiers in neuroscience*, vol. 16, p. 886772, 2022.
- [14] S. Jiang, H. Xue, A. Glover, M. Rutherford, D. Rueckert, and J. V. Hajnal, “MRI of moving subjects using multislice snapshot images with volume reconstruction (SVR): application to fetal, neonatal, and adult brain studies,” *IEEE transactions on medical imaging*, vol. 26, no. 7, pp. 967–980, 2007.