SELF-SUPERVISED ISOTROPIC SUPERRESOLUTION FETAL BRAIN MRI

Kay Lächler†‡, Hélène Lajous‖†, Michael Unser‡, Meritxell Bach Cuadra†‖, and Pol del Aguila Pla†‡

† CIBM Center for Biomedical Imaging, Switzerland

‡ Biomedical Imaging Group, École polytechnique fédérale de Lausanne, Lausanne, Switzerland

‖ Department of Radiology, Lausanne University Hospital (CHUV), and University of Lausanne (UNIL), Lausanne, Switzerland

2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI) | 978-1-6654-7358-3/23/$31.00 ©2023 IEEE | DOI: 10.1109/ISBI53787.2023.10230522

ABSTRACT

Superresolution T2-weighted fetal-brain magnetic-resonance imaging (FBMRI) traditionally relies on the availability of several orthogonal low-resolution series of 2-dimensional thick slices (volumes). In practice, only a few low-resolution volumes are acquired. Thus, optimization-based image-reconstruction methods require strong regularization using hand-crafted regularizers (*e.g.*, TV). Yet, due to *in utero* fetal motion and the rapidly changing fetal brain anatomy, the acquisition of the high-resolution images that are required to train supervised learning methods is difficult. In this paper, we sidestep this difficulty by providing a proof of concept of a self-supervised single-volume superresolution framework for T2-weighted FBMRI (SAIR). We validate SAIR quantitatively in a motion-free simulated environment. Our results for different noise levels and resolution ratios suggest that SAIR is comparable to multiple-volume superresolution reconstruction methods. We also evaluate SAIR qualitatively on clinical FBMRI data. The results suggest SAIR could be incorporated into current reconstruction pipelines.

***Index Terms*—** Image reconstruction, Image enhancement, Neural networks.

1. INTRODUCTION

Magnetic resonance imaging (MRI) of the *in utero* developing brain is a key complementary imaging tool to ultrasound, due to its proven ability to visualize the fetal anatomy with excellent soft tissue contrast. T2-weighted (T2w) imaging is the tool of choice to depict normal and pathological maturation in fetal MRI [1, 2]. Unpredictable fetal motion in the womb is a major challenge and leads to the imperative requirement of fast clinical acquisitions. In practice, ultra fast multi-slice single-shot sequences are acquired, resulting in several orthogonal low-resolution (LR) series (also referred to as volumes) during one session. At 1.5 T, these volumes have a very good in-plane spatial resolution ($\Delta_{\mathrm{x}} = \Delta_{\mathrm{y}} \approx 1$ mm) while the slice thickness is typically chosen between $\Delta_{\mathrm{z}} = 3$ mm and 5 mm in order to ensure a good signal-to-noise ratio. This results in a resolution ratio $r = \Delta_{\mathrm{z}}/\Delta_{\mathrm{x}} \in [3, 5]$.

Several methods have been proposed for the superresolution (SR) reconstruction of fetal T2w imaging [3–12]. These support fetal-brain explorations and allow for the automated quantitative study of isotropic high resolution (HR) 3-dimensional (3D) volumes [13]. Most fetal SR algorithms use multiple orthogonal LR volumes to reconstruct a single isotropic HR volume. Traditional approaches are based on iterative optimization schemes, which combine motion-estimation and image-reconstruction [3–8] steps. In the motion-estimation step, each volume is aligned to a single reference volume using volume-to-volume registration, and each slice within each volume is aligned using slice-to-volume registration. The image enhancement is then formulated as an ill-posed inverse problem with hand crafted regularizers (*e.g.*, Tikhonov, total-variation, or non local means). Prior to the SR reconstruction, these methods need to perform a fetal-brain segmentation and a bias-field correction [14].

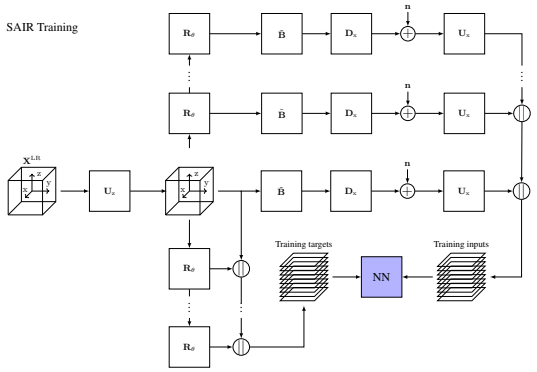

Fig. 1. Training procedure for the (SAIR) reconstruction pipeline.

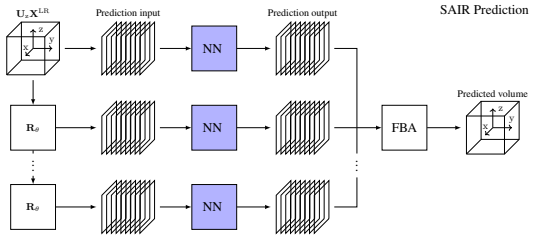

Fig. 2. Schematic description of the prediction procedure for the SAIR reconstruction pipeline.

Deep-learning SR methods are widely explored in MRI. Despite this, their application to fetal-brain MRI (FBMRI) is very limited. Few deep-learning strategies have been proposed within the context of multiple-volume SR reconstruction of the fetal brain [9, 11, 12]. Even then, they tackle the SR problem only partially, either by providing an initial step to optimization-based approaches [9, 12], or by improving only the in-plane (but not the trough-plane) spatial resolution [11]. In this paper, we propose a method that uses a single LR volume and that was designed to obtain a volume with isotropic resolution.

1.1. Related Work on Single-Volume Superresolution

To the best of our knowledge, only two methods have been presented for single-volume SR in diffusion and dynamic functional

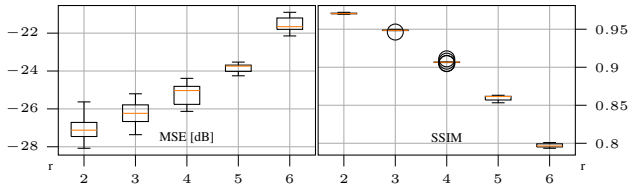

Fig. 3. Performance of the SAIR pipeline when applied to data generated by the simplified MRI model (1) applied to the HR volume described in Section 3.1. Results are reported for nine noise realizations in terms of the resolution ratio r.

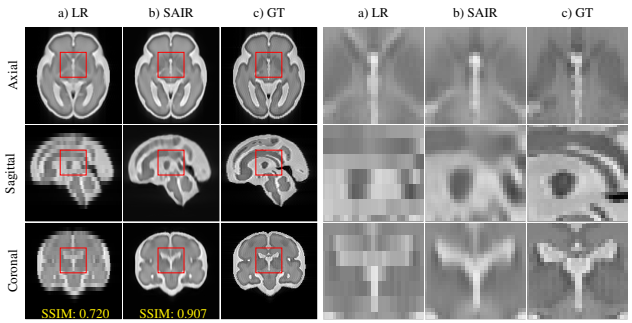

Fig. 4. Left panel: Reconstruction for $r = 4$ as described in Section 4.1.1. The results of b) SAIR reconstruction using a) one single LR series is compared to a synthetic 3D isotropic high resolution ground-truth image. The SSIM is reported in the bottom row. Right panel: Zoom over regions of interest.

fetal MRI [15, 16]—but none for T2w imaging. In [15], an autoencoder model is proposed to leverage HR diffusion-weighted MRI images from preterm babies to enhance the spatial resolution of a fetal diffusion MRI acquisition. This approach relies on training data from preexisting HR datasets, which are hard to obtain and come from a different population. To avoid the need for HR data, self-supervised SR approaches like the one we propose use internal information from the LR image to train the reconstruction procedure.

In fluorescence microscopy, Weigert *et al.* [17] proposed a SR reconstruction method where the information along the LR axis (by convention, z) is learned from the other two HR axes (x and y). The distortion in the LR axis was modeled mathematically and synthetically applied along one of the HR axes to construct pairs of samples to train a network. The resulting network was used to enhance the resolution along the LR axis. The work by Weigert *et al.* inspired our work, which transitions its key ideas to MRI. Our pipeline is similar to the one independently developed in [18]. Our network, however, has much fewer parameters and builds on a simpler MRI model. Furthermore, our quantitative empirical study avoids the inverse crime by using a realistic physics-based MR-acquisition simulator [19]. We also provide a thorough comparison with multiple-volume SR reconstruction methods, which is critical for FBMRI. In [16], a method that extends the same concepts to both the spatial and temporal domains is presented for dynamic functional fetal MRI.

1.2. Contribution

In this paper, we build on [17] to develop a single-volume self-supervised SR method for structural T2w FBMRI. Such methods can be of very high practical value in FBMRI. First, they avoid the need for the harmonization of intensities between the LR series, which is a major pitfall for quantitative applications with current multiple-volume approaches. Second, they can have a significant impact on T2 mapping strategies [20] by significantly reducing the acquisition time. Currently, these strategies can increase the scanning time by up to 12 min, in the case of the acquisition of three (or more) orthogonal series at six different echo times.

2. SINGLE-ACQUISITION ISOTROPIC-RESOLUTION RESTORATION

SAIR, our proposed pipeline, is based on a simplified model of the acquired LR volume $\mathbf{X}^{\mathrm{LR}}$ with respect to the true volume $\mathbf{X}$. In particular, we assume that

$$\mathbf{X}^{\mathrm{LR}} = \mathbf{D}_z \mathbf{B} \mathbf{X} + \mathbf{n}\,, \tag{1}$$

where $\mathbf{D}_z$ is the linear operator that performs downsampling along the z axis (by convention, the axial direction of the series) by a factor of r, $\mathbf{B}$ is a convolutional operator that approximates the frequency response of the MRI scanner (a.k.a. the slice-selection profile, here chosen as in [7]), and $\mathbf{n}$ is an additive white Gaussian noise.

Building on (1), we propose to create a training set based on a simple upsampling $\mathbf{X}^{\mathrm{up}} = \mathbf{U}_z \mathbf{X}^{\mathrm{LR}}$ (*e.g.*, bicubic) of the data volume along z as the training target, and an artificially degraded volume using model (1) along x as the training input $\mathbf{D}_x \tilde{\mathbf{B}} \mathbf{X}^{\mathrm{up}} + \mathbf{n}$, where $\tilde{\mathbf{B}}$ is an axis-permuted version of $\mathbf{B}$ (switching the roles of z and x). If the model represents the distortion caused by the MRI in z well enough, then a network that would perform well on this problem should be able to reconstruct a volume with isotropic resolution from the low-resolution volume $\mathbf{X}^{\mathrm{LR}}$, up to a permutation in the x and z axes. Because there is nothing special about the x axis, any rotation by an angle θ around the z axis of the upsampled volume $\mathbf{R}_\theta \mathbf{X}^{\mathrm{up}}$ is equally valid to augment the training dataset. Thus, we create the training dataset from n_{train} rotated volumes for evenly spaced angles between 0° and 180°. This pipeline is summarized in Figure 1. There, we see that the training inputs are upsampled in the x axis, so that the network is only trained to refine the upsampling operator of choice, $\mathbf{U}_x$. All the results in this paper were obtained with $n_{\mathrm{train}} = 10$ and $\mathbf{U}_z$ the bicubic upsampling operator.

At prediction time, we take an ensemble approach, combining the predictions of the network when applied to n_{pred} different rotations of the upsampled volume between 0° and 180° (see Figure 2). To combine these predictions, we use the Fourier-burst accumulation technique [21], summarized by

$$\mathbf{X}^{\mathrm{pred}} = \mathrm{FFT}^{-1}\left\{ \sum_{m=1}^{n_{\mathrm{pred}}} \mathbf{W}_m \odot \hat{\mathbf{X}}_m^{\mathrm{pred}} \right\}, \tag{2}$$

where $\hat{\mathbf{X}}_m^{\mathrm{pred}}$ is the FFT of $\mathbf{X}_m^{\mathrm{pred}} = \mathbf{R}_{-\theta_m} \mathrm{NN}(\mathbf{R}_{\theta_m} \mathbf{X}^{\mathrm{up}})$, where $\mathrm{NN}(\cdot)$ represents the trained neural network. The weighting mask $\mathbf{W}_m$ is computed as

$$\mathbf{W}_m = \frac{\left|\hat{\mathbf{X}}_m^{\mathrm{pred}}\right|^2}{\sum_{k=1}^{n_{\mathrm{pred}}} \left|\hat{\mathbf{X}}_k^{\mathrm{pred}}\right|^2}, \tag{3}$$

which corresponds to $p = 2$ in the expressions in [21]. All results in this paper were obtained with $n_{\mathrm{pred}} = 15$.

SAIR is a proof-of-concept pipeline in the context of FBMRI. Model (1) captures the most dominant effects of MR acquisition,

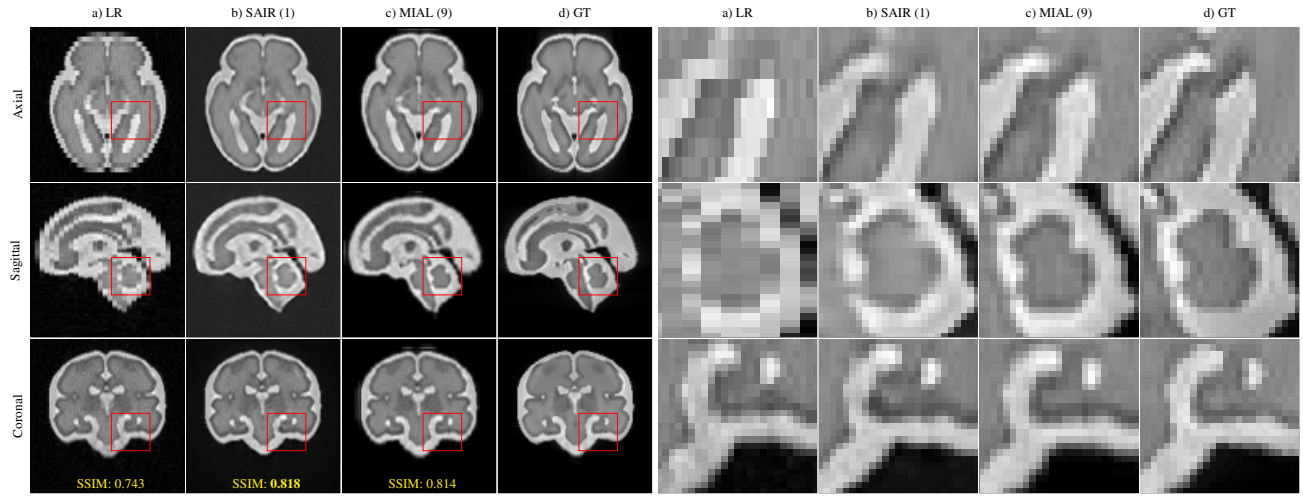

Fig. 5. Left panel: a) synthetic LR HASTE series simulated in the coronal orientation using FaBiAN for a subject of 30 weeks of gestational age; b) SAIR reconstruction from that single LR series; c) MIALSRTK reconstruction combining nine orthogonal LR series (MIAL). The SSIM was computed with respect to d) a synthetic 3D isotropic high resolution ground-truth image. Right panel: Zoom over regions of interest.

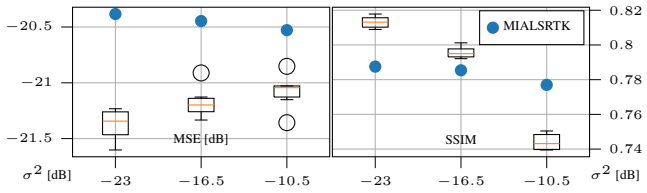

Fig. 6. Reconstruction performance as a function of the noise level. Reported for a single reconstruction with MIALSRTK using nine volumes, and in the form of a boxplot for SAIR for six different noise realizations.

and Fourier-burst accumulation [21] combines frequency information from several independent predictions. Our proposed architecture is a slice-wise (2D) U-Net with 32 convolutional initial channels, (7×7) convolutional kernels, a skip connection, and a single encoder/decoder step. Due to the use of a U-Net with skip connection in the fashion of [22], the network is trained to correct the artifacts in the input image. The architecture is very shallow, which is an advantage in the self-supervised context, in which the network is to be trained for each volume.

3. DATA SETS AND EVALUATION

3.1. Synthetic Low-Resolution Series

Synthetic, yet realistic T2w MR images of the fetal brain at different gestational ages (GA) were derived from a normative spatiotemporal MRI atlas of the fetal brain [23] through a simulated fetal-brain MR acquisition of a numerical phantom (FaBiAN) [19]. In this highly flexible and controlled environment, we have reproduced the clinical MR protocol for the half-Fourier acquisition single-shot turbo spin echo sequence (HASTE, Siemens Healthineers) performed at our local hospital (at 1.5 T: TR/TE, 1200 ms/90 ms; flip angle, 90°; echo train length, 224; echo spacing, 4.08 ms; field-of-view, (360×360) mm^2; voxel size, $(1.1 \times 1.1 \times 3.0)$ mm^3; inter-slice gap, 10%) [24, 25].

For a subject of 30 weeks of GA, nine orthogonal series (three in each orientation, with a shift of the field-of-view of 1.6 mm between each series) were simulated without motion. Random isotropic complex Gaussian noise (with mean 0 and standard deviation $\sigma_0 = 0.15$) was added to the k-space data to qualitatively match the noise characteristics of clinical acquisitions. Corresponding brain masks were automatically generated along with the 2D LR series. Additional data were simulated with two extra noise levels ($\sigma = 0.07$ and 0.30). Six independent realizations of the same LR series were generated for each noise level.

A 3D HR isotropic HASTE image of the fetal brain was simulated for each subject without noise or motion to serve as a reference for the quantitative evaluation of the corresponding SR reconstructions. For comparison, a multiple-volume SR reconstruction was generated with MIALSRTK [7, 26] using the nine simulated orthogonal LR series.

3.2. Clinical Data

Clinical MR images at 1.5 T (MAGNETOM Aera, Siemens Healthcare, Erlangen, Germany) were retrospectively selected for two normally-developing subjects of 28 and 33 weeks of GA. For comparison purposes, both subjects were SR-reconstructed with the MIALSRTK pipeline [7, 26] using all the LR series available for each fetus (three without relevant motion for the younger fetus and eight with little-to-moderate motion for the older fetus). The voxel size was $(1.125 \times 1.125 \times 3.3)$ mm^3.

4. EMPIRICAL RESULTS

4.1. Results on Simulated Data

We first analyze the performance of SAIR quantitatively on synthetic data. Because the ground truth (GT) HR images are available in this scenario, we evaluate reconstructions by computing the mean-squared error (MSE) and the mean structural similarity index (SSIM) over the brain area of every volume, as indicated by the corresponding GT mask.

4.1.1. *Simplified MRI Model*

We start by exploring the performance of SAIR with respect to the resolution ratio $r = \Delta_z/\Delta_x = \Delta_z/\Delta_y$. For this experiment, we use the simplified model (1) to simulate the MRI acquisitions of the subject described in Section 3.1. In Figure 4, we see an example of the resulting LR volume, its reconstruction by SAIR, and the GT HR volume for $r = 4$. Given the coarse resolution of the original data, SAIR does an impressive job at recovering many of the lost details and at improving the SSIM (MSE) from 0.720 (-19.05 dB, respectively) in the LR volume to 0.907 (-22.57 dB, respectively) in the reconstructed volume. In Figure 3, we report the performance of SAIR as a function of r. As expected, it decreases when r increases, but relatively high SSIMs (≈ 0.8) are still achieved for the extreme case $r = 6$.

4.1.2. *Realistic Synthetic Data*

We proceed by evaluating the performance of SAIR on the synthetic T2w MR images described in Section 3.1, compared to that of the MIALSRTK multiple-volume reconstruction with 9 different volumes. In Figure 5, a reconstruction obtained from SAIR (which uses a single volume) is compared to a reconstruction obtained from MIALSRTK (which uses 9 different volumes). There, we observe that 1) the reconstructions are of similar perceptual quality, with SAIR seeming advantageous in the overall view and MIALSRTK dominating in the zoomed view, and that 2) the obtained SSIMs are similar (0.818 for SAIR and 0.814 for MIALSRTK). Additionally, the MSEs are -22.35 dB for SAIR and -21.18 dB for MIALSRTK. However, we suggest here that MSE is not a reliable metric to evaluate the quality of these reconstructions, as it does not seem to match the perceived quality.

We further analyzed both how SAIR and MIALSRTK were affected by the level of noise. In Figure 6, we show the performance of both reconstruction methods for the three noise levels specified in Section 3.1. We observe that SAIR is more sensitive to noise. This reduced robustness was to be expected because SAIR uses only one volume, as opposed to the 9 volumes used by MIALSRTK. Indeed, the trend in the evolution of the SSIM between MIALSRTK and SAIR is similar, with the slopes being much more pronounced in the case of SAIR. We verify here that MSE is not a reliable metric to evaluate the quality of the reconstruction, as shown by the discrepancy between MSE and SSIM changes for MIALSRTK.

4.2. Results on Clinical Data

Finally, we present the SR reconstructions using SAIR and MIALSRTK on two real clinical LR volumes, one without motion (Figure 7) and the other one with little amplitude of fetal movements (Figure 8). Overall, both methods provide very similar results. Although the reconstruction from one single volume with SAIR looks a little blurrier than the one combining multiple volumes with MIALSRTK, the quality of the reconstruction is remarkable, especially when considering that, as opposed to MIALSRTK, SAIR does not include any mechanism to compensate for inter-slice motion.

5. DISCUSSION

In this work, we present the first study of a single-volume self-supervised superresolution method for the reconstruction of T2-weighted magnetic-resonance images (MRI) of the fetal brain compared to more conventional reconstruction techniques such as

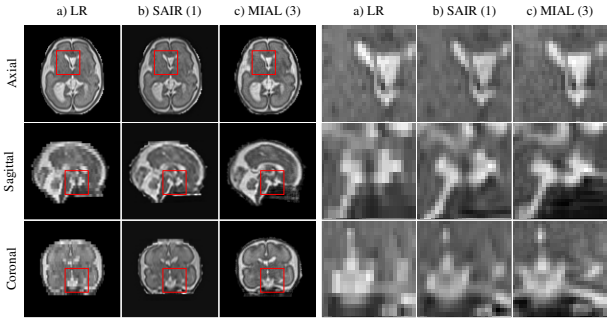

Fig. 7. Reconstruction of a clinical case without motion.

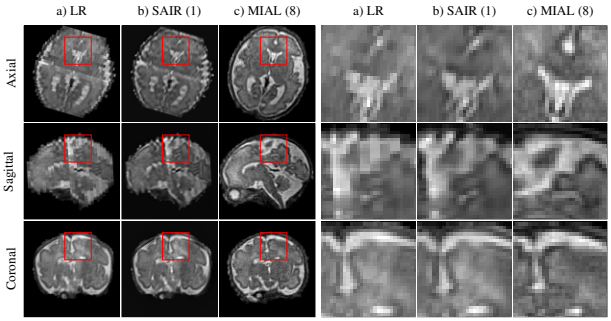

Fig. 8. Reconstruction of a clinical case with little motion artefacts.

MIALSRTK that combine multiple orthogonal low-resolution volumes. We exemplify the accuracy of our proposed single-acquisition isotropic resolution method on synthetic data without motion. Such simulations rely on a numerical phantom and are highly valuable to evaluate the proposed method. They provide not only more realistic LR images than the simplified MRI model commonly used, but also high-resolution isotropic ground-truth images. As proof of concept, we also show the applicability of SAIR on real clinical acquisitions with little amplitude of fetal movements. The resulting single-volume superresolution method will be very useful to replace the interpolation methods currently used to upsample the reference anatomy needed in the motion-estimation step. In addition, the reconstruction of a high-resolution volume of the fetal brain from one single series would be of paramount interest in the perspective of a clinical translation to minimize the acquisition time in this cohort of sensitive subjects. Further development will focus on integrating SAIR with motion-estimation methods, for instance, based on an age-matched HR simulated volume.

6. COMPLIANCE WITH ETHICAL STANDARDS

The use of the clinical MRI dataset was reviewed and approved by Commission cantonale (VD) d'éthique de la recherche sur l'être humain (CER-VD). Written informed consent for participation was not required for this study, in accordance with the national legislation and the institutional requirements.

7. ACKNOWLEDGEMENTS

This work was supported by the Swiss National Science Foundation (205321-182602). We acknowledge access to the facilities and

expertise of the CIBM Center for Biomedical Imaging, a Swiss research center of excellence founded and supported by Lausanne University Hospital (CHUV), University of Lausanne (UNIL), École polytechnique fédérale de Lausanne (EPFL), University of Geneva (UNIGE), and Geneva University Hospitals (HUG).

8. REFERENCES

[1] C. Garel, *MRI of the Fetal Brain: Normal Development and Cerebral Pathologies*, Berlin; New York: Springer, 2004.

[2] D. Prayer et al., "MRI of normal fetal brain development," *European Journal of Radiology*, vol. 57, no. 2, pp. 199–216, 2006.

[3] F. Rousseau et al., "Registration-based approach for reconstruction of high-resolution *in utero* fetal MR brain images," *Academic Radiology*, vol. 13, no. 9, pp. 1072–1081, 2006.

[4] A. Gholipour, J. A. Estroff, and S. 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.

[5] 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.

[6] M. Fogtmann, S. Seshamani, K. Kim, T. Chapman, and C. Studholme, "A unified approach for motion estimation and super resolution reconstruction from structural magnetic resonance imaging on moving objects," in *MICCAI Workshop on Perinatal and Paediatric Imaging: PaPI*, 2012, pp. 9–16.

[7] 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.

[8] B. Kainz et al., "Fast volume reconstruction from motion corrupted stacks of 2D slices," *IEEE Transactions on Medical Imaging*, vol. 34, no. 9, pp. 1901–1913, Sept. 2015.

[9] S. McDonagh et al., "Context-sensitive super-resolution for fast fetal magnetic resonance imaging," in *Molecular Imaging, Reconstruction and Analysis of Moving Body Organs, and Stroke Imaging and Treatment*, vol. 10555 of *Lecture Notes in Computer Science*, pp. 116–126. 2017.

[10] M. Ebner et al., "An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain MRI," *NeuroImage*, vol. 206, pp. 116324, 2020.

[11] L. Song et al., "Deep robust residual network for super-resolution of 2D fetal brain MRI," vol. 12, no. 1, pp. 406, 2022.

[12] J. Xu, D. Moyer, P. E. Grant, P. Golland, J. E. Iglesias, and E. Adalsteinsson, "SVoRT: Iterative transformer for slice-to-volume registration in fetal brain MRI," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2022)*, 2022, pp. 3–13.

[13] A. Makropoulos, S. Counsell, and D. Rueckert, "A review on automatic fetal and neonatal brain MRI segmentation," *NeuroImage*, vol. 170, pp. 231–248, 2018.

[14] A. U. Uus, A. Egloff Collado, T. A. Roberts, J. V. Hajnal, M. A. Rutherford, and M. Deprez, "Retrospective motion correction in foetal MRI for clinical applications: Existing methods, applications and integration into clinical practice," p. 20220071, 2022.

[15] H. Kebiri et al., "Through-plane super-resolution with autoencoders in diffusion magnetic resonance imaging of the developing human brain," *Frontiers in Neurology*, vol. 13, 2022.

[16] J. Xu, E. Abaci Turk, P. E. Grant, P. Golland, and E. Adalsteinsson, "STRESS: Super-resolution for dynamic fetal MRI using self-supervised learning," in *Medical Image Computing and Computer Assisted Intervention (MICCAI 2021)*, 2021, pp. 197–206.

[17] M. Weigert, L. Royer, F. Jug, and G. Myers, "Isotropic reconstruction of 3D fluorescence microscopy images using convolutional neural networks," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2017)*, 2017, pp. 126–134.

[18] C. Zhao, B. E. Dewey, D. L. Pham, P. A. Calabresi, D. S. Reich, and J. L. Prince, "SMORE: A self-supervised anti-aliasing and super-resolution algorithm for MRI using deep learning," *IEEE Transactions on Medical Imaging*, vol. 40, no. 3, pp. 805–817, 2021.

[19] H. Lajous et al., "A fetal brain magnetic resonance acquisition numerical phantom (FaBiAN)," *Scientific Reports*, vol. 12, no. 1, pp. 8682, 2022.

[20] H. Lajous et al., "T2 mapping from super-resolution-reconstructed clinical fast spin echo magnetic resonance acquisitions," in *Medical Image Computing and Computer Assisted Intervention (MICCAI 2020)*, 2020, pp. 114–124.

[21] M. Delbracio and G. Sapiro, "Removing camera shake via weighted Fourier burst accumulation," *IEEE Transactions on Image Processing*, vol. 24, no. 11, pp. 3293–3307, 2015.

[22] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Transactions on Image Processing*, vol. 26, no. 9, pp. 4509–4522, 2017.

[23] A. Gholipour et al., "A normative spatiotemporal MRI atlas of the fetal brain for automatic segmentation and analysis of early brain growth," *Scientific Reports*, vol. 7, no. 1, pp. 476, 2017.

[24] H. Lajous, J.-B. Ledoux, T. Hilbert, R. B. van Heeswijk, and M. Bach Cuadra, "Dataset T2 mapping from super-resolution-reconstructed clinical fast spin echo magnetic resonance acquisitions," 2020.

[25] M. Khawam et al., "Fetal brain biometric measurements on 3D super-resolution reconstructed T2-weighted MRI: An intra- and inter-observer agreement study," *Frontiers in Pediatrics*, vol. 9, no. 639746, pp. 651, 2021.

[26] S. Tourbier, P. De Dumast, H. Kebiri, P. Hagmann, and M. Bach Cuadra, "MIAL Super-Resolution Toolkit v2.0.1," Zenodo, 2020, DOI: 10.5281/zenodo.4392788.