

MIT Open Access Articles

STRESS: Super-Resolution for Dynamic Fetal MRI Using Self-supervised Learning

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Xu, Junshen, Abaci Turk, Esra, Grant, P Ellen, Golland, Polina and Adalsteinsson, Elfar. 2021. "STRESS: Super-Resolution for Dynamic Fetal MRI Using Self-supervised Learning." MEDICAL IMAGE COMPUTING AND COMPUTER ASSISTED INTERVENTION - MICCAI 2021, PT VII, 12907.

Published Version: 10.1007/978-3-030-87234-2_19

Publisher: Springer International Publishing

Permanent Link: <https://hdl.handle.net/1721.1/142674>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: <https://creativecommons.org/licenses/by-nc-sa/4.0/>



STRESS: Super-Resolution for Dynamic Fetal MRI using Self-Supervised Learning

Junshen Xu¹, Esra Abaci Turk², P. Ellen Grant^{2,3}, Polina Golland^{1,4}, and Elfar Adalsteinsson^{1,5}

¹ Department of Electrical Engineering and Computer Science, MIT,
Cambridge, MA, USA

junshen@mit.edu

² Fetal-Neonatal Neuroimaging and Developmental Science Center,
Boston Children’s Hospital, Boston, MA, USA

³ Harvard Medical School, Boston, MA, USA

⁴ Computer Science and Artificial Intelligence Laboratory, MIT,
Cambridge, MA, USA

⁵ Institute for Medical Engineering and Science, MIT, Cambridge, MA, USA

Abstract. Fetal motion is unpredictable and rapid on the scale of conventional MR scan times. Therefore, dynamic fetal MRI, which aims at capturing fetal motion and dynamics of fetal function, is limited to fast imaging techniques with compromises in image quality and resolution. Super-resolution for dynamic fetal MRI is still a challenge, especially when multi-oriented stacks of image slices for oversampling are not available and high temporal resolution for recording the dynamics of the fetus or placenta is desired. Further, fetal motion makes it difficult to acquire high-resolution images for supervised learning methods. To address this problem, in this work, we propose STRESS (Spatio-Temporal Resolution Enhancement with Simulated Scans), a self-supervised super-resolution framework for dynamic fetal MRI with interleaved slice acquisitions. Our proposed method simulates an interleaved slice acquisition along the high-resolution axis on the originally acquired data to generate pairs of low- and high-resolution images. Then, it trains a super-resolution network by exploiting both spatial and temporal correlations in the MR time series, which is used to enhance the resolution of the original data. Evaluations on both simulated and *in utero* data show that our proposed method outperforms other self-supervised super-resolution methods and improves image quality, which is beneficial to other downstream tasks and evaluations.

Keywords: Fetal MRI · Image super-resolution · Self-supervised learning · Deep learning.

1 Introduction

Fetal magnetic resonance imaging (MRI) is an important approach for studying the development of fetal brain *in utero* [18] and monitoring fetal function [15].

Due to unpredictable and rapid fetal motion, dynamic fetal MRI, which aims at capturing fetal motion and dynamics of fetal function, is limited to fast imaging techniques, such as single-shot Echo-planar imaging (EPI) [2], with severe compromises in signal-to-noise ratio (SNR) and image resolution.

Super-resolution (SR) methods is frequently applied to fetal MRI to improve image quality. One well-established category of super-resolution methods for fetal MRI is based on slice-to-volume registration (SVR) [12,21,5]. In these methods, multiple stacks of slices at different orientations are acquired, which are then registered to reconstruct a static and motion-free volume of the chosen region of interest (ROI). However, multi-oriented stacks for oversampling the ROI may not be available. Besides, in some applications, instead of a static ROI, a time series of MR volumes capturing the dynamics of fetal brain, body or placenta is of interest [11,23,15,20]. For example, in [23] and [15], interleaved multi-slice EPI time series are used for fetal body pose tracking and placental function analysis respectively. Thus, it is still a challenge to enhance the resolution in dynamic fetal MRI.

Although supervised super-resolution methods achieved state-of-the-art results in natural images [14,24], the acquisition of HR MRI data with adequate SNR is time consuming and prone to motion artifacts, especially in fetal MRI. To avoid the need for HR data in supervised learning, self-supervised super-resolution (SSR) methods have been developed, which utilize internal information from LR images for super-resolution. For instance, the ZSSR [19] method downsample the LR images to generate lower resolution (LR_2) images and train a network to learn a mapping from LR_2 to LR, which is then applied to the original LR images to estimate the HR images. Similar ideas are also explored in the field of MRI [9,25]. Zhao *et al.* extended [9] and proposed SMORE [25] for SSR of MR volume with anisotropic resolution where the information along the LR axis are learned from the other two HR axes. They blur the volume along the one of the HR axes, extract pairs of training samples to train a network and use it to enhance resolution along the LR axis. However, these methods only applied to a single slice or a stack of images and cannot utilize the temporal information in dynamic imaging.

In this work, we propose a SSR framework for dynamic fetal MRI with interleaved acquisition, named STRESS (**S**patio-**T**emporal **R**esolution **E**nhancement with **S**imulated **S**cans). Using the characteristic of interleaved slice acquisition, we perform simulated acquisitions on the originally acquired data to generate pairs of low- and high-resolution images. We then train a SR network on the extracted data, which exploits both internal spatial information within each frame and temporal correlation between adjacent frames. A optional self-denoising network is also introduced to this framework, when input images are of low SNR. We evaluate the STRESS framework on both simulated and *in utero* data to demonstrate that it can not only enhance resolution of dynamic fetal imaging but also improve performance of downstream tasks.

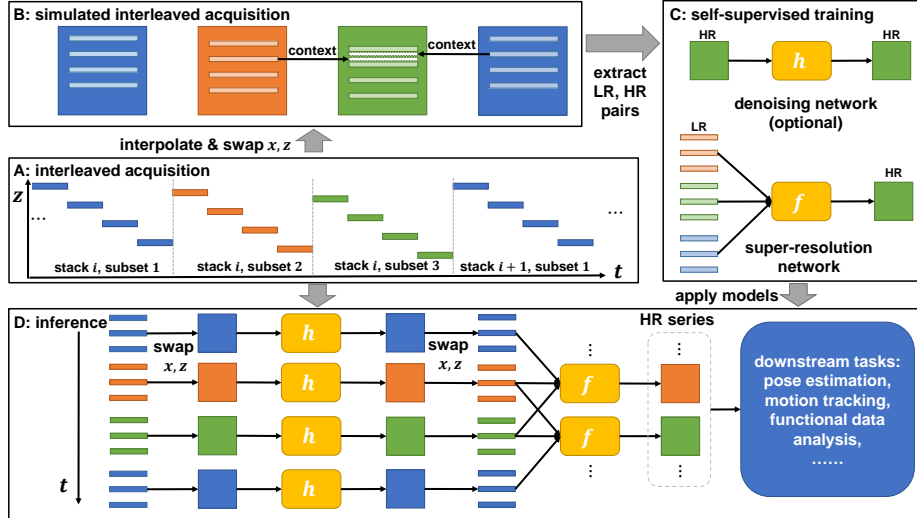


Fig. 1. The proposed STRESS workflow. A: Interleaved MRI acquisitions, e.g., $N_I = 3$. B: Acquired MR data are binned into different time frames. These frames are interpolated and transposed to produce a simulated object with motion. Then, we simulate a interleaved MR scan on this object and extract low- and high-resolution pairs from them C: We train the denoising network (optional) and super-resolution network in self-supervised manners. D: We apply the trained models to the originally or newly acquired data to generate a high-resolution MR volume series, which can be further used for other downstream tasks.

2 Methods

Fig. 1 shows the workflow of the proposed STRESS method, which can be divided into four parts: 1) interleaved slice acquisition, 2) simulated acquisition, 3) self-supervised training, and 4) inference. The details of each part are described in the following sections.

2.1 Interleaved acquisition

Interleaved slice acquisition is a widely used technique to avoid cross-excitation artifacts [3]. The number of slices skipped between two consecutive slice acquisitions is often referred to as the interleave parameter [16], N_I . For example, when $N_I = 2$, even slices are acquired after odd slices. Each image stack in interleaved acquisition are divided into N_I interleaved subsets. In dynamic imaging, multiple stacks are acquired. For simplicity, we refer to the i -th subset in the j -th stack as time frame F_k , where the index $k = N_I \times (i - 1) + j$. The acquisition time of each frame is only $1/N_I$ of the whole stack, making inter-slice motion artifacts within each frame milder. However, the spatial resolution of each frame

along the interleaved axis is also reduced by a factor of N_I . Therefore the interleave parameter can be considered as a trade-off between between spatial and temporal resolutions.

Our goal is to improve the spatial resolution of each frame to generate a HR MR series that has enough temporal resolution to capture fetal dynamics. Let $V_t(x, y, z)$ be the 3D dynamic object to be scanned, where t is time and (x, y, z) are the spatial variables. The acquisition of a slice at time t and location z is $V_t(\cdot, \cdot, z)$. Therefore, the k -th frame can be written as a set of slices, $F_k = \{V_t(\cdot, \cdot, z) | t = t(k, z), z \in \mathcal{Z}_k\}$, where $t(k, z)$ is the time when the slice at location z of the k -th frame is acquired, and \mathcal{Z}_k is the set of slice locations in the k -th frame.

2.2 Simulated interleaved acquisition

To generate HR and LR pairs for training a SSR network, we simulate the interleaved MR acquisition process with the acquired data. For each frame F_k , we interpolate it to make it an isotropic 3D volume denoted by $\tilde{F}_k(x, y, z)$. Then we swap the x - and z - axis¹ and result in a new 3D function \tilde{F}_k^T , i.e., $\tilde{F}_k^T(x, y, z) = \tilde{F}_k(z, y, x)$. $\tilde{F}_k^T(x, y, z)$ is an object of high resolution along the z -axis and having motion similar to V_t . Therefore, we can simulate interleaved acquisition along the z -axis to produce training pairs. The acquired frame in the simulated scan can be written as $S_k = \{\tilde{F}_k^T(\cdot, \cdot, z) | z \in \mathcal{Z}_k\}$. Let \tilde{S}_k be the volume generated by interpolating S_k along the z -axis. We can see that the y - z planes of \tilde{S}_k and \tilde{F}_k , i.e., $\tilde{S}_{k+l}(x, \cdot, \cdot)$ and $\tilde{F}_k^T(x, \cdot, \cdot)$ are pairs of LR and HR images. Besides, it is worth noting that the adjacent time frames provide contexts for estimating the missing slices in the target frame (Fig. 1 B). Therefore, it would be easier to learn a mapping from $\{\tilde{S}_{k+l}(x, \cdot, \cdot)\}_{l=-L}^L$ to $\tilde{F}_k^T(x, \cdot, \cdot)$, where L is the number of time frames used from each side.

2.3 Self-supervised training

Super-resolution: We extract image patches with size of $P \times P$ from the series of images, $\{\tilde{S}_{k+l}(x, \cdot, \cdot)\}_{l=-L}^L$, and concatenate them along the channel dimension to form input tensors $I_{LR} \in \mathbb{R}^{P \times P \times (2L+1)}$. Patches at the same spatial locations are also extracted from $\tilde{F}_k^T(x, \cdot, \cdot)$ as targets and denoted as $I_{HR} \in \mathbb{R}^{P \times P}$. A network f is trained to learn the mapping between I_{LR} and I_{HR} . L1 loss is used to improve the output sharpness, i.e., $\mathcal{L} = \|f(I_{LR}) - I_{HR}\|_1$. We adopt the EDSR [14] architecture for the SSR network f , with 16 residual blocks [8] and 64 feature channels.

Blind-spot denoising: Many fast imaging techniques for capturing fetal dynamics, e.g., EPI, suffer from low SNR [4]. Applying super-resolution algorithms to noisy images tends to emphasize image noise and results in images of

¹ We use x -axis here to keep the notation simple. In fact any axis within the x - y plane can be used.

low quality. To address this problem, we introduce an optional denoising network h to our framework, which can be applied when the original acquired images are of low SNR. The network h is a blind-spot denoising network (BDN) [13], i.e., the receptive field of h doesn't contain the central pixel. Therefore, when we train the network h to recover the input image I by minimizing the mean squared error, $\|h(I) - I\|_2^2$, the network will not become the identity function. Instead, $h(I)$ will approximate the mean of I , so that $h(I)$ can be considered as the denoised image. If BDN is enabled, we first train the denoising network h with images $I = F_k^T(x, \cdot, \cdot)$. Then, when training the SSR network f , we replace the target I_{HR} with $h(I_{HR})$ and the loss becomes $\mathcal{L} = \|f(I_{LR}) - h(I_{HR})\|_1$.

Training details: We set $L = N_I/2$ and $P = 64$, if not specifically indicated. All neural networks are trained on a Nvidia Tesla V100 GPU using an Adam optimizer [10] with a learning rate of 1×10^{-4} for 30000 iterations. We use batch sizes of 64 and 16 for network f and h respectively, which depend on GPU memory. Training images are randomly flipped along the two axes for data augmentation. Our models are implemented with PyTorch 1.5 [17].

2.4 Inference

After training the models, we can apply them to the original or newly acquired data. If BDN is enabled, we first perform image denoising on each frame by applying h to each slice, such that F_k becomes $\{h(V_t(\cdot, \cdot, z)) | t = t(k, z), z \in \mathcal{Z}_k\}$. Then, we interpolate it to generate a volume, $\tilde{F}_k(x, y, z)$. Finally, the trained super-resolution network f is applied to the y - z plane of $\tilde{F}_k(x, y, z)$ and its neighboring frames, which yields a super-resolved estimate \hat{V}_k , i.e., $\hat{V}_k(x, \cdot, \cdot) = f(\{\tilde{F}_{k+l}(x, \cdot, \cdot)\}_{l=-L}^L)$. This process is repeated for all k until we get a HR estimation of the whole series, which can be used for other downstream tasks.

3 Experiments and Results

In the experiments, we apply the following methods to fetal MR volume series: 1) cubic B-spline interpolation along the interleaved axis; 2) interpolation along the temporal direction (TI); 3) spatio-temporal interpolation (STI); 4) SMORE [25] and 5) STRESS. In SMORE, we adopt the same super-resolution network architecture and the same training hyperparameters as STRESS for fair comparison. The reference PyTorch implementation for STRESS is available on GitHub²

3.1 CRL fetal dataset

The CRL fetal atlas [6] consist of T2-weighted fetal brain MRI with gestational age (GA) ranging from 21 to 38 weeks. The images are reconstructed to volume with size of $135 \times 189 \times 155$ and isotropic resolution of 1 mm. To simulate fetal motion, we use the fetal landmark time series in [23]. Specifically, we use two

² <https://github.com/daviddmc/STRESS>

eyes and the midpoint of two shoulder to define the fetal pose and apply affine transformation to the MR volume to generate motion trajectories. There are 77 time series with length from 20 to 30 minutes in the landmark dataset. We randomly sample 10 1-min intervals from each series then apply the motion to the volumes, resulting in $18 \times 77 \times 10 = 13860$ data. We use 70% data for training and validation, 30% for test, data in the test set have different GAs from training and validation sets. We simulate MR scans with $N_I = 2, 4$ and 6, in-plane resolution of $1\text{mm} \times 1\text{mm}$ and slice thickness of 1mm. SR methods are applied to the noise-free data and also noisy data corrupted by Rician noise [7] with standard deviation $\sigma = 3\%$ of the maximum intensity. BDN is enabled when there is noise.

Table 1 shows the peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) [22] comparing to the ground truth. PSNR and SSIM are computed within a mask of non-background voxels. The proposed STRESS method outperforms the competing methods at different interleave parameters, with and without noise. Fig. 2 shows example slices of super-resolution results with $N_I = 4$ and Rician noise. Visual results also indicates that the outputs of STRESS have better image quality.

Models	$N_I = 2$				$N_I = 4$				$N_I = 6$			
	w/o noise		w/ noise		w/o noise		w/ noise		w/o noise		w/ noise	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
SI	32.69	.9883	28.42	.8849	23.90	.9049	22.98	.8114	19.71	.7422	19.39	.6686
TI	29.01	.9111	25.31	.8258	29.21	.9076	25.48	.8273	28.60	.9084	25.52	.8288
STI	31.29	.9682	27.94	.8846	26.87	.9390	25.75	.8711	23.89	.8769	23.37	.8182
SMORE	36.19	.9895	30.38	.9006	31.36	.9687	28.57	.8916	25.29	.8703	24.27	.8093
STRESS	<u>36.77</u>	<u>.9921</u>	<u>33.51</u>	<u>.9702</u>	<u>34.56</u>	<u>.9873</u>	<u>32.81</u>	<u>.9655</u>	<u>28.98</u>	<u>.9480</u>	<u>28.24</u>	<u>.9213</u>

Table 1. PSNR and SSIM of the super-resolution results on the CRL dataset, where ‘w/ noise’ means adding Rician noise with $\sigma = 3\%$ of the maximum intensity. The best results are underlined.

In addition, we also evaluate the performance of the STRESS method with and without BDN under different noise levels ($\sigma = 1\%, 3\%$, and 5% of the maximum intensity). The results are shown in Table 2. We can observe that the BDN makes a larger contribution to the performance of STRESS as the noise level increases.

3.2 Fetal EPI dataset

We also evaluate our method with an *in utero* fetal EPI dataset in [15], which consist of 111 volumetric MRI time series at a gestational age ranging from 25 to 35 weeks. MRIs were acquired on a 3T Skyra scanner (Siemens Healthcare, Erlangen, Germany). Interleaved, multislice, single-shot, gradient echo EPI sequence




Fig. 2. Visual results from CRL fetal dataset ($N_I = 4$, Rician noise $\sigma = 3\%$ of the maximum intensity), numbers in the parentheses are PSNR with ground truth data as reference.

Models	$\sigma = 1\%$		$\sigma = 3\%$		$\sigma = 5\%$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
STRESS w/o BDN	33.96	.9764	30.69	.9219	28.29	.8559
STRESS w/ BDN	33.99	.9826	32.81	.9655	31.09	.9425

Table 2. Evaluations of STRESS with and without BDN under different noise levels ($N_I = 4$).

was used for acquisitions with in-plane resolution of $3\text{mm} \times 3\text{mm}$, slice thickness of 3 mm, average matrix size of $120 \times 120 \times 80$; $\text{TR}=5 - 8\text{s}$, $\text{TE}=32 - 38\text{ms}$, $\text{FA}=90^\circ$, $N_I = 2$. Each subject was scanned for 10 to 30 min. We remove half of the slices at each frame to generate data with $N_I = 4$. We use 92 EPI series for training and 19 for testing. Due to the large voxel size in acquisition and the relatively high SNR, we disable BDN on this dataset. Besides, some volumes have matrix size less than 64, so we use $P = 32$ in this experiment.

Since ground truth is not available for the *in utero* dataset, we use the removed slices as reference to compute PSNR and SSIM. To further evaluate the quality of output images, we use fetal keypoint detection as a downstream task, where 15 fetal keypoints (ankles, knees, hips, bladder, shoulders, elbows, wrists and eyes) are detected from each time frame. Ground truth labels are manually annotated on the original data with $N_I = 2$. We apply a pretrained keypoint detection model [23] to the output volumes of each SR method. The percentage of correct keypoint (PCK) [1] are computed. $\text{PCK}(s) = N(s)/N \times 100\%$, where N is the total number of keypoints and $N(s)$ is the number of predicted keypoints with error less than threshold s .

Fig. 3 shows the evaluation of super-resolution results on the fetal EPI dataset. The proposed STRESS method achieves the highest PSNR and SSIM among all competing methods, which is also shown by the t-test. Besides, when using the super-resolution results for fetal keypoint detection, the results of STRESS also have the best performance in terms of PCK, indicating that the STRESS method is able to generate MR time series with high image quality which is beneficial to downstream tasks.

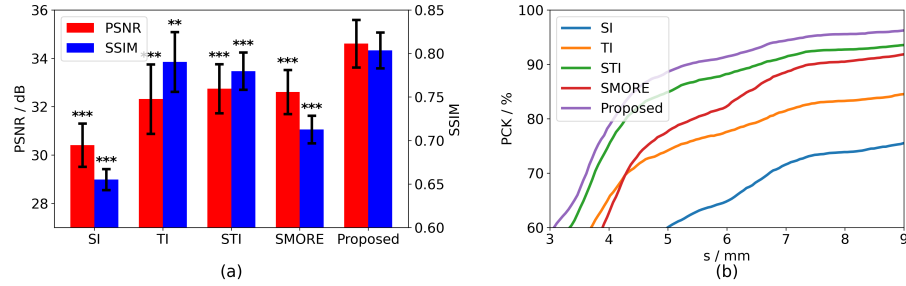


Fig. 3. Evaluation of super-resolution results from fetal EPI data with $N_I = 4$. Left: PSNR and SSIM comparing to the reference in the $N_I = 2$ data. Error bars show the corresponding standard deviations. **: p-value $< 10^{-2}$, ***: p-value $< 10^{-3}$. Right: PCK curves for fetal landmark detection using a pretrained model.




Fig. 4. Visual results from *in utero* fetal EPI dataset.

Fig. 4 shows example slices of super-resolution results in one frame of the fetal MR series. We can see that the results of the proposed STRESS method have the best perceptual quality. The output of SI is very blurred, since it only interpolates along the z -axis. The TI and STI methods utilize temporal information with simple interpolation and therefore introduce severe inter-slice misalignment to the images. Although SMORE achieves better image quality than interpolation methods, the boundary of fetal brain is unclear in the outputs of SMORE. The reason is that SMORE only take a single frame as input without the temporal context, so that it cannot restore the details in the body parts that are corrupted by fetal motion, such as the fetal brain. STRESS, however, utilizes both spatial and temporal information of the scan data during the self-supervised training process, and therefore recovers more image details.

4 Conclusions

This paper presents STRESS, a self-supervised super-resolution framework for dynamic fetal imaging with interleaved slice acquisition. STRESS trains a SR

network in a self-supervised manner, where low- and high-resolution training samples are extracted from simulated interleaved acquisitions. The SR network utilizes both internal spatial information within each frame and temporal correlation between adjacent frames to improve image quality and restore details corrupted by fetal motion. Evaluations on both simulated and *in utero* data shows that STRESS outperforms other competing methods. The experiments also demonstrate that STRESS is beneficial when serving as a data pre-processing step for further downstream analysis.

Acknowledgements

This research was supported by NIH U01HD087211, NIH R01EB01733 and NIH NIBIB NAC P41EB015902.

References

1. Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2d human pose estimation: New benchmark and state of the art analysis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3686–3693 (2014)
2. Diogo, M.C., Prayer, D., Gruber, G.M., Brugger, P.C., Stuhr, F., Weber, M., Bettelheim, D., Kasprian, G.: Echo-planar flair sequence improves subplate visualization in fetal mri of the brain. *Radiology* **292**(1), 159–169 (2019)
3. Dowling, J., Bourgeat, P., Raffelt, D., Fripp, J., Greer, P.B., Patterson, J., Denham, J., Gupta, S., Tang, C., Stanwell, P., et al.: Nonrigid correction of interleaving artefacts in pelvic mri. In: Medical Imaging 2009: Image Processing. vol. 7259, p. 72592P. International Society for Optics and Photonics (2009)
4. Gholipour, A., Estroff, J.A., Barnewolt, C.E., Robertson, R.L., Grant, P.E., Gagoski, B., Warfield, S.K., Afacan, O., Connolly, S.A., Neil, J.J., et al.: Fetal mri: a technical update with educational aspirations. *Concepts in Magnetic Resonance Part A* **43**(6), 237–266 (2014)
5. Gholipour, A., Estroff, J.A., Warfield, S.K.: Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain mri. *IEEE transactions on medical imaging* **29**(10), 1739–1758 (2010)
6. Gholipour, A., Rollins, C.K., Velasco-Annis, C., Ouaalam, A., Akhondi-Asl, A., Afacan, O., Ortinau, C.M., Clancy, S., Limperopoulos, C., Yang, E., et al.: A normative spatiotemporal mri atlas of the fetal brain for automatic segmentation and analysis of early brain growth. *Scientific reports* **7**(1), 1–13 (2017)
7. Gudbjartsson, H., Patz, S.: The rician distribution of noisy mri data. *Magnetic resonance in medicine* **34**(6), 910–914 (1995)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
9. Jog, A., Carass, A., Prince, J.L.: Self super-resolution for magnetic resonance images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 553–560. Springer (2016)
10. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2017)

11. Kochunov, P., Castro, C., Davis, D.M., Dudley, D., Wey, H.Y., Purdy, D., Fox, P.T., Simerly, C., Schatten, G.: Fetal brain during a binge drinking episode: a dynamic susceptibility contrast mri fetal brain perfusion study. *Neuroreport* **21**(10), 716 (2010)
12. Kuklisova-Murgasova, M., Estrin, G.L., Nunes, R.G., Malik, S.J., Rutherford, M.A., Rueckert, D., Hajnal, J.V.: Distortion correction in fetal epi using non-rigid registration with a laplacian constraint. *IEEE transactions on medical imaging* **37**(1), 12–19 (2017)
13. Laine, S., Karras, T., Lehtinen, J., Aila, T.: High-quality self-supervised deep image denoising. *arXiv preprint arXiv:1901.10277* (2019)
14. Lim, B., Son, S., Kim, H., Nah, S., Mu Lee, K.: Enhanced deep residual networks for single image super-resolution. In: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. pp. 136–144 (2017)
15. Luo, J., Turk, E.A., Bibbo, C., Gagoski, B., Roberts, D.J., Vangel, M., Tempany-Afdhal, C.M., Barnewolt, C., Estroff, J., Palanisamy, A., et al.: In vivo quantification of placental insufficiency by bold mri: a human study. *Scientific reports* **7**(1), 1–10 (2017)
16. Parker, D., Rotival, G., Laine, A., Razlighi, Q.R.: Retrospective detection of interleaved slice acquisition parameters from fmri data. In: *2014 IEEE 11th International Symposium on Biomedical Imaging (ISBI)*. pp. 37–40. IEEE (2014)
17. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: *Automatic differentiation in pytorch* (2017)
18. Saleem, N., S.: Fetal mri: An approach to practice: A review. *Journal of Advanced Research* **5**(5), 507–523 (2014)
19. Shocher, A., Cohen, N., Irani, M.: “zero-shot” super-resolution using deep internal learning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3118–3126 (2018)
20. Turk, E.A., Abulnaga, S.M., Luo, J., Stout, J.N., Feldman, H.A., Turk, A., Gagoski, B., Wald, L.L., Adalsteinsson, E., Roberts, D.J., et al.: Placental mri: effect of maternal position and uterine contractions on placental bold mri measurements. *Placenta* **95**, 69–77 (2020)
21. Uus, A., Zhang, T., Jackson, L.H., Roberts, T.A., Rutherford, M.A., Hajnal, J.V., Deprez, M.: Deformable slice-to-volume registration for motion correction of fetal body and placenta mri. *IEEE transactions on medical imaging* **39**(9), 2750–2759 (2020)
22. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
23. Xu, J., Zhang, M., Turk, E.A., Zhang, L., Grant, P.E., Ying, K., Golland, P., Adalsteinsson, E.: Fetal pose estimation in volumetric mri using a 3d convolution neural network. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 403–410. Springer (2019)
24. Zhang, Y., Tian, Y., Kong, Y., Zhong, B., Fu, Y.: Residual dense network for image super-resolution. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2472–2481 (2018)
25. Zhao, C., Dewey, B.E., Pham, D.L., Calabresi, P.A., Reich, D.S., Prince, J.L.: Smore: A self-supervised anti-aliasing and super-resolution algorithm for mri using deep learning. *IEEE transactions on medical imaging* (2020)