

# Context-Sensitive Super-Resolution for Fast Fetal Magnetic Resonance Imaging

Steven McDonagh<sup>1</sup>(✉), Benjamin Hou<sup>1</sup>, Amir Alansary<sup>1</sup>, Ozan Oktay<sup>1</sup>,  
Konstantinos Kamnitsas<sup>1</sup>, Mary Rutherford<sup>2</sup>, Jo V. Hajnal<sup>2</sup>,  
and Bernhard Kainz<sup>1,2</sup>

<sup>1</sup> Biomedical Image Analysis Group, Imperial College London, London, UK  
[steven.mcdonagh@imperial.ac.uk](mailto:steven.mcdonagh@imperial.ac.uk)

<sup>2</sup> Division of Imaging Sciences and Biomedical Engineering,  
Kings College London, London, UK

**Abstract.** 3D Magnetic Resonance Imaging (MRI) is often a trade-off between fast but low-resolution image acquisition and highly detailed but slow image acquisition. Fast imaging is required for targets that move to avoid motion artefacts. This is in particular difficult for fetal MRI. Spatially independent upsampling techniques, which are the state-of-the-art to address this problem, are error prone and disregard contextual information. In this paper we propose a context-sensitive upsampling method based on a residual convolutional neural network model that learns organ specific appearance and adopts semantically to input data allowing for the generation of high resolution images with sharp edges and fine scale detail. By making contextual decisions about appearance and shape, present in different parts of an image, we gain a maximum of structural detail at a similar contrast as provided by high-resolution data. We experiment on 145 fetal scans and show that our approach yields an increased PSNR of 1.25 dB when applied to under-sampled fetal data *cf.* baseline upsampling. Furthermore, our method yields an increased PSNR of 1.73 dB when utilizing under-sampled fetal data to perform brain volume reconstruction on motion corrupted captured data.

## 1 Introduction

Currently, 3D imaging of moving objects is limited by the time it takes to acquire a single image. The slower an imaging modality is, the more likely motion induced artefacts will occur within and between individual slices of a 3D volume. Very fast imaging modalities like Computed Tomography are not always applicable because of harmful ionising radiation, and ultrasound often suffers from poor image quality. Thus, Magnetic Resonance Imaging (MRI) is usually the modality of choice when; large fields of view, high anatomical detail, and non-invasive imaging is required. MRI is often applied to image involuntary moving objects such as the beating heart and examination of the fetus in-utero. Motion compensation for cardiac imaging can be achieved through ECG gating. However, fetal targets do not provide options for gated or tracked image acquisition to compensate for motion. Thus motion compensation is performed during

post-processing of oversampled input spaces, usually involving the acquisition of orthogonally oriented stacks of slices [8]. Oversampling with high resolution (HR) slices causes long scan times, which is uncomfortable and risky for patients like pregnant women. This limits the possible number of scan sequences during examination. However, improving image resolution is key to improving accuracy, understanding of anatomy and assessment of organ size and morphology. Imaging at lower resolution increases acquisition speed, thus partly mitigating the likelihood for motion between individual slices but at the cost of missing structural detail that could render the scan inappropriate for diagnostic purposes. Due to signal-to-noise ratio (SNR) limitations, the acquired slices are usually also *thick* compared to the in-plane resolution and thus negatively influence the visualization of anatomy in 3D.

Naïve up-sampling of fast but low resolution (LR) images is undesirable for the clinical practice, since results lack information. Information content cannot be increased by simply increasing the number of pixels with linear interpolation methods. Therefore, optimization-based super-resolution (SR) methods have been explored to generate rich volumetric information from oversampled input spaces. However, these methods are highly dependent on the quality and amount of input samples and depend on the choice of the objective function. Recent work, *e.g.* [4], on example-based SR has focused on incorporating additional prior image knowledge, and, in particular, deep neural networks have been employed to solve the single-image SR (SISR) problem. However, the majority of recent contributions typically place strong emphasis on natural images and therefore lack domain specific high-frequency detail prior knowledge [1].

**Contribution:** We present a novel approach to SISR in the context of motion compensation when using fast to acquire, low resolution volumes. Taking inspiration from recent investigation of network based SR for MRI modalities [15], we propose a network architecture with convolutional and transposed-convolutional layers and hypothesize that such a deep network architecture can be tailored to context sensitive applications, such as motion compensation of the fetal brain, and yield volume reconstruction improvements from low resolution input. Our network learns subject specific details from potentially motion corrupted input data and accurately reintroduces the expected fidelity allowing motion compensation and high quality reconstruction from fast low resolution input.

Our model is in particular data-adaptive since the upsampling is performed by learnable transposed-convolution layers instead of a fixed kernel. By performing the upsampling in the final layers of the network we avoid early redundant computation in a HR space, enabling a computational saving. Additionally by considering entire LR in-plane slice samples at training time, in comparison to image *patches*, we gain a large receptive field to enable the learning of spatial context, organ structure and anatomy.

We evaluate our method on 145 healthy fetal scans. The proposed approach shows improved qualitative results when compared visually to linear methods. Quantitative reconstruction performance, peak signal-to-noise-ratio (PSNR) and structural similarity index measure (SSIM) improve, accordingly. In particular,

we reach comparable reconstruction quality with half as many data samples, thus half of the currently required scan time, when compared to motion compensated reconstruction from high-resolution image acquisition.

**Related work:** The topic of SR has received much attention in the literature and a large body of work exists however, historically, algorithms exhibiting good performance in 2D domains such as satellite or facial imagery, are not necessarily ideal for 3D medical imaging. This is partly due to domain specific effects such as loss of spatial information caused by motion during slow target acquisition. Various algorithms have been shown to produce leading results [14] in differing domains.

SISR accounts for missing image information by using previously observed examples to optimise the LR-HR mapping between images or patches. In the medical imaging domain, data-adaptive patch-based approaches to SISR reconstruction [7, 13] have been shown to prevent the occurrence of well-known blurring effects, often found when using classical interpolation approaches. Interpolation techniques tend to increase the smoothness of images in an isotropic manner, however data-adaptive non-local methods allow for highly anisotropic reconstruction where required. In patch-based methods, the radius of 3D patch used to compute the similarity among voxels is often a free parameter and the choice of receptive field size typically affects computational cost when using iterative optimisation.

Learning based approaches also allow data-adaptive reconstruction and CNNs in particular have recently been successfully applied to context sensitive SISR for cardiac imaging. The work of [15] use a regression architecture based on [4] with a modified  $l_1$  objective function. The approach performed SR in the slice-select direction of lowest MRI resolution, *i.e.*, one-dimensionally and utilized transposed convolutional layers at the start of the network architecture to perform the upsampling, *prior* to convolutions, thus learning high level features in latter layers on (spatially) large feature maps.

Two-dimensional SR is a popular research area in natural image processing due to many applications requiring enhancement of a visual experience while limiting the amount of raw data that needs to be recorded, transferred or stored. Recent network-based approaches such as SRGAN [12] apply Generative Adversarial Networks (GAN) to achieve large up-sampling factors of up to four.

Motion compensation for MRI volume reconstruction typically incorporates a SR component. However to the best of our knowledge state-of-the-art network based SR techniques, capable of learning problem and sensor specifics from available data have not been harnessed for the upsampling step found in Slice-to-Volume frameworks for the reconstruction task. In this work we investigate the accuracy advantages that such an approach can contribute to the example of fetal MRI volume reconstruction.

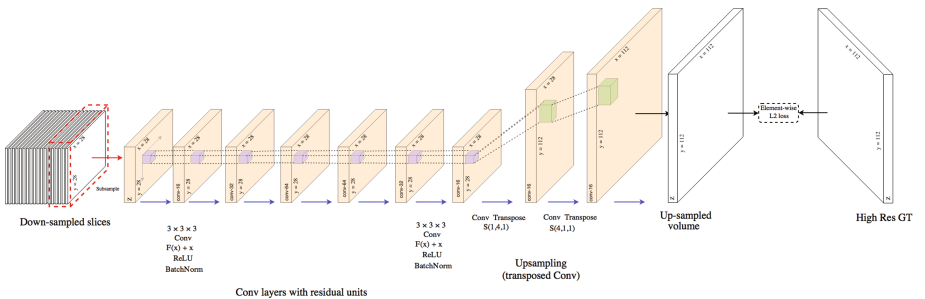
Contemporary SR components in MRI Slice-to-Volume reconstruction (SVR) tasks perform optimisation based incremental updates to the HR volume estimate. To achieve this, the SR problem for volume reconstruction has been modelled directly by considering minimisation of an error norm function and

use of Huber function statistics [5] or gradient-weighted averaging [10]. The ill-posed nature of modelling upsampling requires that the objective be regularised. Gholipour *et al.* [5] add a Tikhonov term to the cost for this purpose while Rousseau *et al.* [16–18] select a regularisation term that includes an approximation of Total Variation (TV) to better preserve edges. Tourbier *et al.* [21] apply fast convex optimization techniques for the SR problem also using an edge-preserving TV regularization. Murgasova *et al.* [11] used intensity matching and complete outlier removal for reconstruction. SR volume intensities are iteratively updated using the error gradients resulting from differences between simulated and observed slice samples. Transforming observed slice information to the upsampled volume space requires accurate yet potentially computationally expensive estimation of the sensor point spread function (PSF) and [8] developed a fast multi-GPU accelerated implementation for the task.

## 2 Method

The proposed approach implements a fully three-dimensional CNN architecture to infer upsampled MRI imagery, enabling HR input to be provided for subsequent SVR and motion compensation tasks. We define an architecture utilising 3D volumetric convolutions that have recently been shown to add value for medical imaging tasks considering 3D imagery [2,9]. Figure 1 provides a schematic of our upsampling network and architecture design details are provide in the **3D MRI CNN** subsection below. Figure 2 provides a schematic diagram indicating where the upsampling network component is implemented in a SVR reconstruction framework.

The architecture differs from recent network based MRI SR models [15] by generating feature maps in the LR image space *cf.* early redundant feature channel upsampling or fixed kernels [3], reducing memory and computation requirements while retaining the flexibility of learnable upsampling layers. As previously reported [19], early upsampling tends to introduce redundant computation in the



**Fig. 1.** Our proposed CNN network architecture for MRI super-resolution. See text for architecture details.



the common frugal strategy [20] of applying small ( $3 \times 3 \times 3$ ) convolution kernels and spending compute-budget alternatively on layer count to increase receptive field size.

By introducing two transposed convolution layers we perform the upscaling on in-plane sampling dimensions. In this manner, upscaling weights are learned specifically for the SR task where  $(x \uparrow U_x) * w_j = h_j^0$  and  $(h_1 \uparrow U_y) * w_j = h_j^1$  where  $\uparrow$  is a zero-padding upscaling operator and  $\{U_x, U_y\} = M/N$  are the in-plane upscaling factors. This allows for explicit optimization of the upsampling filters and facilitates training in an end-to-end manner for the SR task. By implementing trainable upsampling layers we improve upon the alternative strategy of initial independent linear upsampling, followed only by convolutional layers, as we gain an ability to learn upsampling weights specific to the SR task. In practice this often improves MRI image signal quality in image regions close to boundaries [15]. Residuals learned by the convolution layers and the upscaled transposed-convolutional output are used to reconstruct the final HR image. This allows the regression function to learn non-linearities such as the high frequency components of the signal.

Training involves evaluating the error function  $\Psi_{l_2}(\cdot)$  that calculates the difference between the reconstructed HR images and the ground truth volumes that were down-sampled to provide training data. Model weights are updated using standard back-propagation and adaptive moment estimation. In comparison to modified  $l_1$  losses [15] or recent *perceptual-quality* SR objective functions [12], we implement a standard voxel-wise  $l_2$  loss function to provide gradient information and emphasize voxel-wise difference to the ground-truth. An implementation of our model training strategy is made available online<sup>1</sup>.

**Fetal Brain Volume Reconstruction:** We combine our SR network with Slice-to-Volume registration (SVR) [8]. SVR requires multiple orthogonal stacks of 2D slices to provide improved reconstruction quality. By upsampling stacks prior to reconstruction we provide a means to acquire larger sets of low-resolution input. The motion-free 3D image is then reconstructed from the upsampled slices and motion-corrupted and misaligned areas are excluded during the reconstruction using an EM-based outliers rejection model.

### 3 Experiments

**Data:** We test our approach on clinical MR scans with varying gestational age. All scans have been ethically approved. The dataset contains 145 MR scans of healthy fetuses at gestational age between 20–25 weeks. The data has been acquired on a Philips Achieva 1.5T, the mother lying 20° tilt on the left side to avoid pressure on the inferior vena cava. Single-shot fast spin (ssFSE) echo T2-weighted sequences are used to acquire stacks of images that are aligned to the main axes of the fetus. Three to six stacks with a voxel size of

<sup>1</sup> <https://github.com/DLTK>.


1.25 mm  $\times$  1.25 mm  $\times$  2.5 mm per stack are acquired for the whole womb. Imagery is manually masked and cropped to isolate fetal brain regions.

**Experimental details:** We employ our 3D MRI network and separately two baseline SR strategies to upsample image stack inputs that serve as input to the SVR pipeline. SVR then performs motion compensation and volume reconstruction. We assess upsampled image quality directly and, additionally, investigate the effect of the proposed upsampling strategy on reconstruction quality, from the (initially) low resolution fetal data. We report three quantitative metrics: PSNR, structural similarity index (SSIM) and cross-correlation. In the first experiment, the data is randomly split into two subsets and used to train (100) and test (45) with our SR network. MRI stacks represent 46 individual patients and all image stacks, belonging to a particular patient, are found uniquely in either the train or test set. Images are cropped, intensity normalised and linearly downsampled by factors of 2 and 4 with respect to the in-plane stack axes. This resampling provides LR images to our network resulting in multiple training samples per volume with corresponding ground-truth label (HR source image). The network uses these training pairs to learn the LR to HR mapping. Note that image volume size choices introduce a trade-off between available contextual information and pragmatic memory constraints.


## 4 Evaluation and Results

**Image Quality Assessments:** We compare HR ground-truth 3D volumes with upsampled LR raw data by measuring PSNR, SSIM and cross-correlation. We report SSIM, in particular, due to the well-understood metric properties that afford assessment of local structure correlation and reduced noise sensitivity. LR test imagery is upsampled in-plane ( $X, Y$ ) by factors of 2, 4 to align with target ground-truth resolution. Quality metrics in Fig. 3 report improvements observed for an image upsampling factor of 2. This provides initial evidence in support of our hypothesis; *learning problem and sensor specific deconvolutional filters to perform MRI stack upsampling is of benefit for subsequent resolution-sensitive tasks such as motion compensation and HR volume reconstruction.*

By learning problem specific HR synthesis models, our 3D MRI CNN strategy outperforms the naïve baseline up-sampling, quantitatively improving the



**Fig. 3.** PSNR, SSIM and Cross Correlation metrics for 45 LR image stacks with voxel spacing (2.50  $\times$  2.50  $\times$  1.25) mm that are upsampled  $\times 2$  in-plane ( $X, Y$ ) and compared to ground-truth image stacks (1.25  $\times$  1.25  $\times$  1.25) mm using Linear, B-Spline, 3D MRI CNN methods.




**Fig. 4.** Orthogonal fetal MRI stacks showing in-plane stack axes per row. Low resolution input (left) is upsampled by two baselines (col *Linear*, *B-spline*) and our learning based approach (col *3D MRI CNN*) cf. ground-truth (GT) HR imagery. The learning based 3D MRI CNN, with modality specific priors, provides improved high frequency signal components cf. baselines.

quality of the inferred HR imagery. Figure 4 exhibits an example of qualitative improvement in orthogonal fetal MRI test-stack axes.


We additionally perform preliminary experiments towards integrating network-based SR components more tightly with an SVR pipeline by investigating the ability of the network to upsample LR voxel intensities that result from an initial volume reconstruction iteration. Successful integration of an iterative (learning-based) SR and volume reconstruction loop will facilitate the well understood mutual benefits of reduced-motion SR input and improved input fidelity for the motion correction task. Qualitative comparison of ( $\times 4$ ) LR volume-reconstructed input and resulting upsampled results are found in Fig. 5. The benefit of learning the upsampling with modality specific data can be observed to manifest as sharper edge gradients and improved high frequency signal components. The visual quality gap between the baselines and our method can be seen to widen as the prior information required to successfully upsample at larger factors make the task more challenging.

**Volume Reconstruction Improvement:** In our third experiment we evaluate SVR performance using LR input stacks, upsampled by the considered strategies, before initiating the volume reconstruction task. We additionally perform SVR reconstruction with original HR imagery to provide the “ground-truth” reference brain volumes. Employing the three quality metrics, introduced previously, we evaluate how well super-resolved LR stack reconstructions correspond to the reconstructions due to original high, in-plane, resolution imagery. Table 1 reports PSNR, SSIM and cross-correlation metrics for volume comparison (SR strategy with respect to “ground-truth” volume) for the 13 patients that define the MRI stack test set. Super-resolving the LR input data with the proposed learning based approach can be observed to facilitate reconstruction improvement,





**Fig. 5.** SR applied to LR ( $\times 4$ ) volume reconstructed input. Benefits of learning the specific non-linearities to recover sharp edge gradients and improved high frequency signal components of the modality become more evident *cf.* baselines as the amount of information required to upsample-successfully increases.



**Fig. 6.** (a) 2D slice through a fetal brain reconstruction, resulting from HR input-imagery. Attempting similar reconstruction from faster to acquire LR imagery, at half the in-plane resolution, results in highly degraded visual reconstruction quality (b) and gross DSSIM disparity (*ie.* red heatmap regions) (f) with respect to the HR reconstruction. Naïve up-sampling ( $\times 2$ ) of the LR in-plane input prior to reconstruction, with linear interpolation or B-splines, result in over-smoothed input. Loss of sharp gradient information and input-image fidelity can be seen to propagate to the respective reconstructions (c), (d) and disparity, with regard to the HR reconstruction, remains high (g), (h). Our 3D MRI CNN upsampling affords input closer to the original HR imagery and results in improved reconstructions (e) and reduced DSSIM (i) with visibly cooler heatmap regions (standard *jet* color scale). (Color figure online)

across the investigated metrics. Visual evidence supporting this claim is found in Fig. 6 (best viewed in color). Figure 6 displays 2D slices of patient fetal brain reconstructions resulting from the original HR input-imagery (far left) and identically spatially-located slices (a) resulting from (b) LR imagery (half the in-plane resolution), (c-d) input using naïve up-sampling strategies and (e) our 3D MRI CNN upsampling. Corresponding Structural Dissimilarity (DSSIM) error heatmaps (second row) provide improved visual spatial congruence between HR ground-truth and our method, supporting the claim that utilizing sensor specific priors is of marked benefit for the task of MRI fetal brain reconstruction from LR imagery.

**Table 1.** PSNR, SSIM and Cross-correlation evaluating disparity between reconstructed volumes using upsampled LR input (Linear, B-Spline, 3D MRI CNN) and ground-truth volumes.

Upsample	PSNR [dB]	SSIM	Cross-correlation
<i>No upsample</i>	18.466 ± 1.88	0.534 ± 0.15	0.699 ± 0.12
Linear	19.268 ± 1.14	0.665 ± 0.08	0.815 ± 0.06
B-Spline	19.985 ± 1.52	0.698 ± 0.12	0.836 ± 0.08
3D MRI CNN	<b>21.715 ± 1.84</b>	<b>0.779 ± 0.10</b>	<b>0.885 ± 0.07</b>

## 5 Discussion and Conclusion

We introduce a 3D MRI CNN to upsample low resolution MR data prior to performing volumetric motion compensation and SVR reconstruction. Our method produces upsampled images and uses them to reconstruct volumetric fetal brain representations that quantitatively outperform on reconstruction tasks that utilise conventional upscaling methods. This contribution helps to address the well-understood image resolution challenge in fetal brain MRI. Analysis of accuracy metrics, assessing upsampling quality, exhibit a mean PSNR increase of 1.25 dB. Furthermore, when utilizing the upsampled imagery as SVR input, reconstructed fetal brain volumes show improvements of up to 1.73 dB over the provided baseline. In addition to quality improvement, 3D MRI CNN upsampling provides a computationally efficient approach affording an ability to initially image at lower resolutions, with a shorter acquisition time, thus provides faster and safer scanning for high-risk patients like pregnant women.

The current work has implicitly provided evidence that the method learns the PSF of the investigated MRI data well. In future it would be valuable to investigate this further, explicitly. Real-world LR/HR samples, acquired from scanners at differing resolutions, would allow quantitative evaluation of the ability to reconstruct physical scanner PSF and would further allow investigation of a model’s ability to generalise to the reconstruction of PSFs not explicitly seen at training time. Further to this; the current work only investigates a single problem instance under one image modality. Future work will look to investigate the generalisability of the proposed framework to additional problem domains.

## References

1. Borman, S., et al.: Super-resolution from image sequences—a review. In: Midwest Symposium on Circuits and Systems, pp. 374–378. IEEE (1998)
2. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 424–432. Springer, Cham (2016). doi:[10.1007/978-3-319-46723-8\\_49](https://doi.org/10.1007/978-3-319-46723-8_49)
3. Dong, C., Deng, Y., Loy, C.C., Tang, X.: Compression artifacts reduction by a deep convolutional network, pp. 576–584, December 2015

4. Dong, C., et al.: Image super-resolution using deep convolutional networks. *IEEE Trans. PAMI* **38**(2), 295–307 (2016)
5. Gholipour, A., et al.: Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI. *TMI* **29**(10), 1739–1758 (2010)
6. Greenspan, H.: Super-resolution in medical imaging. *Comput. J.* **52**(1), 43–63 (2009)
7. Jia, Y., He, Z., Gholipour, A., Warfield, S.K.: Single anisotropic 3-D MR image upsampling via overcomplete dictionary trained from in-plane high resolution slices. *IEEE J. Biomed. Health Inf.* **20**(6), 1552–1561 (2016)
8. Kainz, B., et al.: Fast volume reconstruction from motion corrupted stacks of 2D slices. *Trans. Med. Imag.* **34**(9), 1901–1913 (2015)
9. Kamnitsas, K., Ferrante, E., Parisot, S., Ledig, C., Nori, A.V., Criminisi, A., Rueckert, D., Glocker, B.: DeepMedic for brain tumor segmentation. In: Crimi, A., Menze, B., Maier, O., Reyes, M., Winzeck, S., Handels, H. (eds.) *BrainLes 2016*. LNCS, vol. 10154, pp. 138–149. Springer, Cham (2016). doi:[10.1007/978-3-319-55524-9\\_14](https://doi.org/10.1007/978-3-319-55524-9_14)
10. Kim, K., et al.: Intersection based motion correction of multislice MRI for 3-D in utero fetal brain image formation. *Trans. Med. Imag.* **29**(1), 146–158 (2010)
11. Kuklisova-Murgasova, M., Quaghebeur, G., Rutherford, M.A., Hajnal, J.V., Schnabel, J.A.: Reconstruction of fetal brain MRI with intensity matching and complete outlier removal. *Med. Image Anal.* **16**(8), 1550–1560 (2012)
12. Ledig, C., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: *Computer Vision and Pattern Recognition (CVPR)* (2017)
13. Manjón, J.V., Coupé, P., Buades, A., Fonov, V., Collins, D.L., Robles, M.: Non-local MRI upsampling. *Med. Image Anal.* **14**(6), 784–792 (2010)
14. Nasrollahi, K., et al.: Super-resolution: a comprehensive survey. *Mach. Vis. Appl.* **25**(6), 1423–1468 (2014)
15. Oktay, O., et al.: Multi-input cardiac image super-resolution using convolutional neural networks. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) *MICCAI 2016*. LNCS, vol. 9902, pp. 246–254. Springer, Cham (2016). doi:[10.1007/978-3-319-46726-9\\_29](https://doi.org/10.1007/978-3-319-46726-9_29)
16. Rousseau, F., Kim, K., Studholme, C., Koob, M., Dietemann, J.-L.: On super-resolution for fetal brain MRI. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) *MICCAI 2010*. LNCS, vol. 6362, pp. 355–362. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-15745-5\\_44](https://doi.org/10.1007/978-3-642-15745-5_44)
17. Rousseau, F., et al.: Registration-based approach for reconstruction of high-resolution in utero fetal MR brain images. *Acad. Radiol.* **13**(9), 1072–1081 (2006)
18. Rousseau, F., et al.: BTK: an open-source toolkit for fetal brain MR image processing. *Comput. Methods Programs Biomed.* **109**(1), 65–73 (2013)
19. Shi, W., Caballero, J., Huszár, F., Totz, J., Aitken, A.P., Bishop, R., Rueckert, D., Wang, Z.: Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network, pp. 1874–1883 (2016)
20. Simonyan, K., et al.: very deep convolutional networks for large-scale image recognition. [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) abs/1409.1556 (2014)
21. Tourbier, S., Bresson, X., Hagmann, P., Thiran, J.P., Meuli, R., Cuadra, M.B.: An efficient total variation algorithm for super-resolution in fetal brain MRI with adaptive regularization. *NeuroImage* **118**, 584–597 (2015)