



HAL
open science

Joint Multi-Contrast Reconstruction of Fetal MRI based on Implicit Neural Representations

Steven Jia, Chloé Mercier, Alexandre Pron, Nadine Girard, Guillaume Auzias,
François Rousseau

► To cite this version:

Steven Jia, Chloé Mercier, Alexandre Pron, Nadine Girard, Guillaume Auzias, et al.. Joint Multi-Contrast Reconstruction of Fetal MRI based on Implicit Neural Representations. *Perinatal, Preterm and Paediatric Image Analysis. PIPPI 2024, Lecture Notes in Computer Science, 14747*, Springer Nature Switzerland, pp.14-23, 2025, *Lecture Notes in Computer Science*, 10.1007/978-3-031-73260-7_2 . hal-04650604v2

HAL Id: hal-04650604

<https://hal.science/hal-04650604v2>

Submitted on 24 Jul 2024

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.

Joint Multi-Contrast Reconstruction of Fetal MRI based on Implicit Neural Representations

Steven Jia¹, Chloé Mercier², Alexandre Pron¹, Nadine Girard^{1,3},
Guillaume Auzias^{1*}, and François Rousseau^{4*}

¹ Institut de Neurosciences de la Timone, UMR 7289, CNRS, Aix-Marseille
Université, 13005, Marseille, France

² IMT Atlantique, Lab-STICC UMR 6285, Brest, France

³ Aix-Marseille Univ, APHM, Service de Neuroradiologie, Hôpital de la Timone,
13005, Marseille, France

⁴ IMT Atlantique, LaTIM UMR 1101, Brest, France
`steven.jia@univ-amu.fr`

Abstract. Fetal cerebral brain magnetic resonance imaging (MRI) is critical for the detection of abnormal brain development before birth. A key image processing step is the reconstruction of a 3D high resolution volume from the acquired series of 2D slices. Several types of MR sequences are commonly acquired during a scanning session, but current reconstruction methods consider each sequence (or contrast) separately. Multi-contrast techniques have been proposed but they do not compensate for potential movement during the acquisition, which occurs almost systematically in the context of fetal MRI. In this work, we introduce a new method for the joint reconstruction of multiple 3D volumes from different contrasts. Our method combines the redundant and complementary information across several stacks of 2D slices from different acquisition sequences via an implicit neural representation, and includes a slice motion correction module. Our results on both simulations and real data acquired in clinical routine demonstrates the relevance and efficiency of the proposed method.

Keywords: MRI · Fetal Brain · Reconstruction · Multiscale latent representation

1 Introduction

Early detection of abnormal brain development in the fetus is crucial, as fetal cerebral brain magnetic resonance imaging (MRI) can identify these issues before birth. Brain malformations, affecting roughly 10% of children, can lead to significant lifelong neurological disabilities. To achieve this, developing non-invasive MRI-derived biomarkers and fetal-specific computational tools for predicting potential abnormal development is crucial.

* These authors contributed equally to this work

Fetal MRI exams differ from adult scans due to the need for minimal acquisition time to minimize motion artifacts from the fetus and mother. Unlike adult scans that acquire a single 3D volume, fetal MRI relies on acquiring multiple sets of thicker 2D slices –denoted as *stacks of 2D slices*– where motion can occur between slices. A crucial post-processing step for each acquisition sequence (e.g., T1-weighted, T2-weighted) is the reconstruction of a high-resolution 3D volume –denoted as *3DHR volume*– from the acquired stacks.

Most reconstruction methods [13, 7, 2, 1] rely on an iterative framework with two steps: slice motion estimation, and super-resolution of a discrete reconstructed 3DHR volume. The NeSVoR method [17], introduced the use of an Implicit Neural Representation (INR) to learn a continuous function. This function defines the association between any 3D coordinates in \mathbb{R}^3 and the corresponding intensity in the target 3DHR volume.

Note that in all these methods, only a single type of acquisition –denoted below as *contrast* is considered, while multiple sequences are commonly acquired during a fetal MRI exam. As a matter of illustration, the protocol used for the dHCP dataset [3] consisted of stacks of T1w slices but also T2w slices, as well as diffusion weighted sequences. Multi-contrast reconstruction techniques, which allow for the combination of inter-contrast information from the same anatomical object, are therefore particularly relevant in this context.

Several multi-contrast reconstruction techniques have been proposed for brain MRI.[14] and [10] introduced methods to improve super-resolution of an input stack of 2D slices by incorporating fine-grained information from a different 3DHR volume of the same subject. [18] proposed the first deep-learning approach using a two-stage convolutional neural network (CNN). The first CNN learns the reconstruction from the stacks, and the second injects information from the other high-resolution volume.[11] introduced an INR-based method for multi-contrast reconstruction, using the SIREN INR [15]. They reported an improvement of 7 dB of PSNR relative to traditional cubic-spline interpolation and of 2 db compared to single-contrast INR for joint reconstruction of T1w and T2w contrasts.

Other relevant works include [6, 9, 4, 8]. It is important to note that these methods typically address reconstruction directly from the raw k-space data acquired by the MRI scanner. In contrast, our present work focuses on reconstruction within the image domain (also known as super-resolution).

These techniques are however not appropriate for fetal brain MRI. The movement of the fetus during scans necessitates solutions that address this issue. In the present work, we introduce a novel method for joint reconstruction of multi-contrast 3DHR volumes. Our method exploits the redundant and complementary information across several stacks of 2D slices from different acquisition sequences and incorporates a slice motion correction module.

2 Methods

Let $I \in \mathbb{R}^{N_s \times N_p}$ denotes the acquired stacks of slices for a given contrast, where I_{ij} represents the intensity of the j -th pixel in the i -th slice, N_s the number of slices and N_p the number of pixels in each slice. Let $\Omega \subset \mathbb{R}^3$ the 3D physical space in which the 3DHR volume reconstruction is performed. Our approach is based on the reconstruction framework introduced in [17]. In that work, the reconstructed 3DHR volume V is computed from a stack of slices I for a given contrast by minimizing the following problem:

$$\operatorname{argmin}_{T, V, B, C, \sigma} \sum_{i=1}^{N_s} \mathcal{L}(I_i, \hat{I}_i) + \mathcal{R}_V \quad (1)$$

where the loss function \mathcal{L}_{ij} is defined as:

$$\mathcal{L}_{ij} = \frac{(I_{ij} - \hat{I}_{ij})^2}{2\sigma_{ij}^2} + \frac{1}{2} \log(\sigma_{ij}^2) \quad (2)$$

and T, V, B, C and σ are the parameters to be estimated. The estimated slice \hat{I}_i which is compared to the observed one I_i in Equation 2 is computed as:

$$\hat{I}_{ij} = C_i \int_{\Omega} B_i(x) M_{ij} V(x) dx \quad (3)$$

where M_{ij} is a matrix modeling both the estimated rigid motion T_i of the corresponding slice and a downsampling operator (related to the point spread function for a given MR acquisition), the estimated bias for this slice B_i , and C_i a slice-level scaling factor. ϵ_i is a Gaussian noise with null mean and covariance $\mathbb{E}[\epsilon_i(x)\epsilon_j(x)] = \sigma_i^2(x)\delta(x-y)$ with δ the Dirac function, such that the noise variance σ_{ij}^2 can be computed as:

$$\sigma_{ij}^2 = C_i^2 \int_{\Omega} B_i^2(x) M_{ij}^2 \sigma_i^2(x) dx \quad (4)$$

The second term of Equation 1 is a regularization factor for enforcing image regularity and reducing noise, defined as:

$$\mathcal{R}_V = \int_{\Omega} r(\|\nabla V(x)\|_2) dx \quad (5)$$

with r is the Huber (edge-preserving) loss function.

Multi-contrast reconstruction In this work, we expand the approach to reconstruct two 3DHR volumes simultaneously, each from a different contrast. Let consider two sets of stacks of 2D slices from two different contrasts I and $J \in \mathbb{R}^{N_{s_k} \times N_{p_k}}$ for $k \in \{I, J\}$, with N_{s_k} the number of slices and N_{p_k} the number of pixels in each slice, with $k \in \{I, J\}$. Our goal is to jointly reconstruct the corresponding 3DHR volumes V_I and V_J .

The multi-contrast extension of variational energy for slice motion estimation and joint reconstruction of 3DHR volumes can be written as follows:

$$\operatorname{argmin}_{T_{I,J}, V_{I,J}, B_{I,J}, C_{I,J}, \sigma_{I,J}} \sum_{i=1}^{N_{s_I}} \mathcal{L}(I_i, \hat{I}_i) + \mathcal{R}_{V_I} + \sum_{i=1}^{N_{s_J}} \mathcal{L}(J_i, \hat{J}_i) + \mathcal{R}_{V_J} \quad (6)$$

While both sets of 2D slices come from the same anatomical object, they represent different "perspectives" due to the use of two distinct MRI sequences. This can lead to variations in resolution (number of slices and pixels) and potential movement between slices.

While Equation 6 might imply separate calculations for each parameter, our goal is a joint reconstruction approach. To achieve this, we introduce a common latent representation. This essentially means a single underlying description of the anatomy (the fetal brain) to be reconstructed. This shared INR-based representation serves as the basis for estimating all the different parameters within a generative modeling framework, leading to a joint reconstruction of $V_{I,J}$.

Model Architecture The architecture of the proposed joint implicit network is shown in Figure 1. The network is divided into two parts: a shared multiscale latent representation and a generative modeling network. The latent representation is computed using a hash grid encoding [12] at multiple scales for feature extraction. To achieve this, for each position x within the 3D space, the network first identifies eight vertices surrounding that position. It then performs a linear combination of the feature vectors associated with these eight vertices to obtain a representative feature vector for position x . Notably, this feature vector simultaneously incorporates information from both contrasts and can be interpreted as a disentanglement module. The generative modeling network relies on a multi-layer perceptron (MLP) that decodes the shared multiscale features (represented by $\phi(x)$ in Equation 7) to predict the intensities for both contrasts at position x :

$$z(x) = MLP_V(\phi(x)) \quad (7)$$

where $V_1(x)$ and $V_2(x)$ are two distinct components of $z(x)$. The noise for each slice of each contrast (denoted as $\sigma_{1_{ij}}$ and $\sigma_{2_{ij}}$ for the j -th voxel at the i -th slice) is estimated using a dedicated network taking as inputs the MLP output $z(x)$, and a dedicated embedding e_{ij} for the slices.

This multi-contrast approach takes advantage of data from all available acquisitions to estimate the motion of each slice. This generative framework can be seen as a unified model of a multi-contrast registration of stacks based on a common representation of the fetal brain to be reconstructed. This ensures that both reconstructed volumes occupy the exact same physical space.

3 Experiments and Results

Implementation details

All the hyper parameters were set as in the implementation of NeSVoR available at <https://github.com/daviddmc/NeSVoR>. All models were implemented

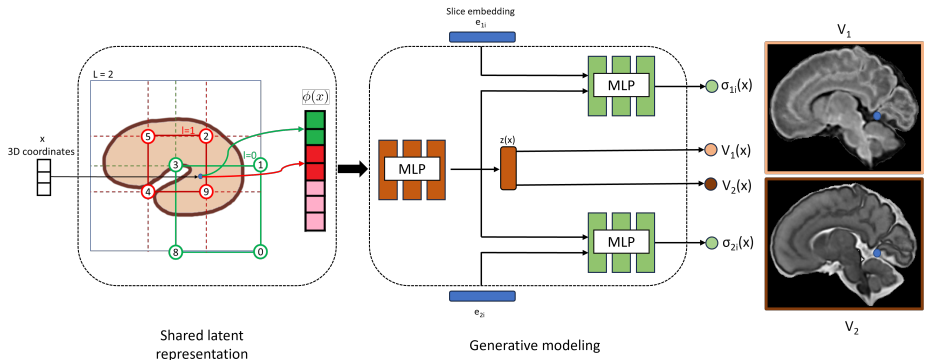


Fig. 1. Overview of the joint multi-contrast reconstruction INR from a coordinate of a slice during the training. The corresponding input x 3D coordinate can be found on the 3DHR volumes V_1 and V_2 .

using PyTorch version 1.13.1 with CUDA 11.7 and tested on a system equipped with a single NVIDIA RTX A2000 12GB GPU and a 12th Gen Intel(R) Core(TM) i7-12700 CPU.

3.1 Simulated Data

As mentioned in Section 1, obtaining ground truth 3D volumes for fetal MRI reconstruction is challenging. To address this, we designed multi-contrast simulations based on the fetal dHCP atlas [16].

We downsampled the T1w and T2w 3DHR volumes from the dHCP atlas to simulate sets of S stacks of 2D slices with a resolution of $0.5\text{mm} \times 0.5\text{mm} \times F$ mm, where F can be used to control for the thickness of the simulated stack, and S the number of stacks. Each stack was sampled following a principal orientation (axial, coronal or sagittal) and we simulated the inter-slice rigid motion as a combination of translation in $[-3\text{mm}; 3\text{mm}]$ and rotation in $[-3^\circ; 3^\circ]$. Figure 2 shows a set of simulated stacks used in our experiments. The number and orientation of stacks as input varies. 1 axial stack or 3 (ax, cor, sag) stacks are inputs for the T1w, and HR, 1 coronal stack or 3 (ax, cor, sag) stacks are input for the T2w contrast. The simulated data can then be reconstructed using either the mono-contrast version of NesVOR, i.e. aiming at reconstructing each contrast 3DHR volume independently, or considering jointly the two acquisitions using our proposed method. Each estimated 3DHR volume is then quantitatively compared to the original 3D volume from the atlas using the PSNR score.

Results from simulations The quantitative comparison between the mono-contrast and the multi-contrast reconstructions is reported in Table 1, which summarises six different experimental setups depending on stack resolution, slice motion and number of input stacks. In the experiments reported in the first 8 lines of the table, the joint reconstruction of the T1w is guided by the 3DHR




Fig. 2. Example of simulated data. Top: stack of axial 2D T1w slices. Middle: stack of sagittal 2D T1w slices. Bottom: stack of coronal 2D T2w slices.

T2w volume without downsampling. This set of experiments is a proof of concept of the improvement resulting from the injection of additional information from T2w volume into the reconstruction of the T1w 3DHR volume. The last 4 lines of Table 1 show experiments where both T1w and T2w input data are downsampled and contaminated by motion, to be consistent with real clinical data.

Using joint multi-contrast reconstruction leads to higher PSNR in almost all experiments. The gain in PSNR varies from 1 to 5 dB, and is higher in presence of motion. Figure 3 illustrates the qualitative improvement obtained in the experiment with 1 axial T1w stack with a resolution of $0.5 \times 0.5 \times 1.5\text{mm}$ with a motion range between $[-3\text{mm}; 3\text{mm}]$ and a 3DHR T2w volume.

3.2 Real Clinical Data

We also evaluated our method on real data acquired in clinical routine. The imaging protocol consisted of two different T2w sequences (contrasts): a half-Fourier single-shot turbo spin-echo (HASTE) sequence providing good signal-to-noise ratio with an excellent T2-weighted contrast and a TRUFISP sequence that is highly efficient at demonstrating cortical malformation and mid-line abnormalities [5]. It has to be noticed that, in this experiment, the input data consisted only of stacks of slices with motion, as no complete 3DHR volume was available for the method (see Figure 4 for an illustration, 3 stacks for TRUFISP and 3

Table 1. Quantitative comparison between the mono-contrast NeSVoR and our multi-contrast method on simulated data. 1 stack of T1w is axial, 1 stack of T2w is coronal, 3 stacks in T1w and T2w are axial, coronal and sagittal.

Method	T1w Res. F(mm)	T1w Num. of stacks	T1w Motion	T1w PSNR	T1w SSIM	T2w Res. F(mm)	T2w Num. of stacks	T2w Motion	T2w PSNR	T2w SSIM
Mono	6	1 (ax.)	0	21.8	0.4	0.5	HR	0	-	-
Joint	6	1 (ax.)	0	25	0.64	0.5	HR	0	-	-
Mono	6	3 (ax., cor., sag.)	0	27.2	0.74	0.5	HR	0	-	-
Joint	6	3 (ax., cor., sag.)	0	29.6	0.81	0.5	HR	0	-	-
Mono	1.5	1 (ax.)	3	18.6	0.29	0.5	HR	0	-	-
Joint	1.5	1 (ax.)	3	23.4	0.65	0.5	HR	0	-	-
Mono	1.5	3 (ax., cor., sag.)	3	34.4	0.94	0.5	HR	0	-	-
Joint	1.5	3 (ax., cor., sag.)	3	34.8	0.94	0.5	HR	0	-	-
Mono	6	3 (ax., cor., sag.)	3	26.2	0.74	6	3 (ax., cor., sag.)	3	22.6	0.53
Joint	6	3 (ax., cor., sag.)	3	26	0.72	6	3 (ax., cor., sag.)	3	25.3	0.66
Mono	6	1 (ax.)	3	18.8	0.36	6	1 (cor.)	3	22.5	0.58
Joint	6	1 (ax.)	3	19.3	0.41	6	1 (cor.)	3	24.6	0.69




Fig. 3. Left: T1w 3DHR volume reconstructed using mono-contrast NeSVoR [17]. Middle: T1w 3DHR volume reconstructed using our multi-contrast method. Right: T2w 3DHR volume given as input to our multi-contrast method in this experiment.

stacks for HASTE). To compensate for large motions before applying NeSVoR and our proposed method, all stacks were roughly registered using SVoRT [17].




Fig. 4. Illustration of the stacks from two different contrasts from the same fetus used in our experiment on real data. Left TRUFISP contrast. Right: HASTE contrast. Note that since the stacks were acquired sequentially, the movements across slices can substantially vary between the two contrasts.

Figure 5 shows the reconstructed 3DHR volumes for one subject (33 weeks post-conceptual age) with 3 stacks of TRUFISP contrast with a resolution per stack of $0.81\text{mm} \times 0.81\text{mm} \times 3.5\text{mm}$ and 3 stacks of HASTE contrast with a resolution per stack of $0.68\text{mm} \times 0.68\text{mm} \times 3\text{mm}$. The proposed multi-contrast approach provides a real visual gain for the reconstruction of the HASTE volume compared with the mono-contrast approach (NeSVoR) (top row), in particular a reduction in artefacts and a better reconstructed cortical plate. The arrows indicate regions where the difference is noticeable. It’s worth mentioning that this joint approach benefits both reconstructions. The contrast in the cortex (outer layer of the brain) in the TRUFISP reconstruction becomes clearer. Additionally, the joint approach avoids artifacts (indicated by the red circle) that might appear in mono-contrast reconstructions.

4 Discussion

The improvements observed in both simulated and clinical data demonstrate the effectiveness of combining information across several contrast data acquired on the same individual. The proposed generative modelling is based on the definition of a common multi-scale latent representation, enabling information from the stacks of different contrasts to be merged. This approach enables more accurate reconstruction of fine structures such as the cortical plate, and reduces




Fig. 5. Experiment on real data acquired in clinical routine. Top: HASTE contrast. Bottom: TRUFISP contrast. Left: mono-contrast, independent reconstruction using NeSVoR [17]. Right: our multi-contrast method. Red arrows point to regions where the improvements are particularly visible. Note that the joint reconstruction also improves the TRUFISP (red circle).

reconstruction artefacts. The shared multiscale latent representation is implemented using a hashgrid-based INR as proposed in NeSVoR [17]. The comparison with other approaches such as SIREN [15] used in [11] will be investigated in future work.

5 Conclusion

In this work, we proposed a joint multi-contrast reconstruction method for fetal brain MRI. A shared latent representation is used for slice motion correction and 3DHR volume reconstruction of both contrasts. Our results on both simulations and real data acquired in clinical routine demonstrate the relevance and efficiency of the proposed method for fetal MRI reconstruction.

Acknowledgments. The research leading to these results has also been supported by the ANR AI4CHILD Project, Grant ANR-19-CHIA-0015, the ANR SulcalGRIDS Project, Grant ANR-19CE45-0014, the ERA-NET NEURON MULTI-FACT Project, Grant ANR-21-NEU2-0005 and the ANR HINT Project, Grant ANR-22-CE45-0034 funded by the French National Research Agency. Data were provided by the developing Human Connectome Project, KCL-Imperial-Oxford Consortium funded by the European Research Council under the European Union Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement no. [319456]. We are grateful to the families who generously supported this trial.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Cordero-Grande, L., et al.: Fetal MRI by Robust Deep Generative Prior Reconstruction and Diffeomorphic Registration. *IEEE Transactions on Medical Imaging* **42**(3), 810–822 (Mar 2023). <https://doi.org/10.1109/TMI.2022.3217725>
2. Ebner, M., et al.: An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain MRI. *NeuroImage* **206**, 116324 (Feb 2020). <https://doi.org/10.1016/j.neuroimage.2019.116324>
3. Edwards, A.D., et al.: The Developing Human Connectome Project Neonatal Data Release. *Frontiers in Neuroscience* **16**, 886772 (May 2022). <https://doi.org/10.3389/fnins.2022.886772>
4. Feng, C.M., et al.: Exploring Separable Attention for Multi-Contrast MR Image Super-Resolution (Aug 2022). <https://doi.org/10.48550/arXiv.2109.01664>
5. Girard, N., et al.: MR imaging of acquired fetal brain disorders. *Child’s nervous system : ChNS : official journal of the International Society for Pediatric Neurosurgery* **19**(7-8), 490–500 (Aug 2003). <https://doi.org/10.1007/s00381-003-0761-x>
6. Huang, J., et al.: Fast multi-contrast MRI reconstruction. *Magnetic Resonance Imaging* **32**(10), 1344–1352 (Dec 2014). <https://doi.org/10.1016/j.mri.2014.08.025>
7. Kuklisova-Murgasova, M., et al.: Reconstruction of fetal brain MRI with intensity matching and complete outlier removal. *Medical Image Analysis* **16**(8), 1550–1564 (Dec 2012). <https://doi.org/10.1016/j.media.2012.07.004>
8. Lei, P., et al.: Decomposition-Based Variational Network for Multi-Contrast MRI Super-Resolution and Reconstruction. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 21239–21249. IEEE, Paris, France (Oct 2023). <https://doi.org/10.1109/ICCV51070.2023.01947>
9. Lyu, Q., et al.: Multi-Contrast Super-Resolution MRI Through a Progressive Network. *IEEE transactions on medical imaging* **39**(9), 2738–2749 (Sep 2020). <https://doi.org/10.1109/TMI.2020.2974858>
10. Manjón, J.V., et al.: MRI superresolution using self-similarity and image priors. *International Journal of Biomedical Imaging* **2010**, 425891 (2010). <https://doi.org/10.1155/2010/425891>
11. McGinnis, J., et al.: Single-subject Multi-contrast MRI Super-resolution via Implicit Neural Representations (Sep 2023), <http://arxiv.org/abs/2303.15065>
12. Müller, T., et al.: Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM Transactions on Graphics* **41**(4), 1–15 (Jul 2022). <https://doi.org/10.1145/3528223.3530127>
13. Rousseau, F., et al.: Registration-Based Approach for Reconstruction of High-Resolution In Utero Fetal MR Brain Images. *Academic Radiology* **13**(9), 1072–1081 (Sep 2006). <https://doi.org/10.1016/j.acra.2006.05.003>
14. Rousseau, F., et al.: A non-local approach for image super-resolution using intermodality priors. *Medical Image Analysis* **14**(4), 594–605 (Aug 2010). <https://doi.org/10.1016/j.media.2010.04.005>
15. Sitzmann, V., et al.: Implicit Neural Representations with Periodic Activation Functions (Jun 2020). <https://doi.org/10.48550/arXiv.2006.09661>

16. Uus, A., et al.: Multi-channel spatio-temporal mri atlas of the normal fetal brain development from the developing human connectome project. *IEEE Transactions on Medical Imaging* (6) (2023). <https://doi.org/https://doi.org/10.12751/g-node.ysgsy1>
17. Xu, J., et al.: NeSVoR: Implicit Neural Representation for Slice-to-Volume Reconstruction in MRI. *IEEE Transactions on Medical Imaging* **42**(6), 1707–1719 (Jun 2023). <https://doi.org/https://doi.org/10.1109/tmi.2023.3236216>
18. Zeng, K., et al.: Simultaneous single- and multi-contrast super-resolution for brain MRI images based on a convolutional neural network. *Computers in Biology and Medicine* **99**, 133–141 (Aug 2018). <https://doi.org/10.1016/j.combiomed.2018.06.010>