

IM-MoCo: Self-supervised MRI Motion Correction using Motion-Guided Implicit Neural Representations

Ziad Al-Haj Hemidi^[0000–0002–9823–802X], Christian Weihsbach^[0000–0002–3839–096X], and Mattias P. Heinrich^[0000–0002–7489–1972]

Institute of Medical Informatics, Universität zu Lübeck, Lübeck, Germany
z.alhajhemidi@uni-luebeck.de
[Code repository](#)

Abstract. Motion artifacts in Magnetic Resonance Imaging (MRI) arise due to relatively long acquisition times and can compromise the clinical utility of acquired images. Traditional motion correction methods often fail to address severe motion, leading to distorted and unreliable results. Deep Learning (DL) alleviated such pitfalls through generalization with the cost of vanishing structures and hallucinations, making it challenging to apply in the medical field where hallucinated structures can tremendously impact the diagnostic outcome. In this work, we present an instance-wise motion correction pipeline that leverages motion-guided Implicit Neural Representations (INRs) to mitigate the impact of motion artifacts while retaining anatomical structure. Our method is evaluated using the NYU *fastMRI* dataset with different degrees of simulated motion severity. For the correction alone, we can improve over state-of-the-art image reconstruction methods by +5% SSIM, +5 *db* PSNR, and +14% HaarPSI. Clinical relevance is demonstrated by a subsequent experiment, where our method improves classification outcomes by at least +1.5 accuracy percentage points compared to motion-corrupted images.

Keywords: Motion Correction · Reconstruction · Implicit Neural Representations · Magnetic Resonance Imaging

1 Introduction

Magnetic Resonance Imaging (MRI) is a non-invasive imaging technique providing detailed images of soft-tissue structures. However, common artifacts, particularly motion-related ones, degrade image quality, impacting clinical diagnosis and leading to high clinical costs [23]. Various retrospective strategies for motion correction (MoCo) have been developed [4, 24]. Retrospective correction handles motion artifacts post-acquisition, offering flexibility without requiring scanner modifications. This approach benefits from prior knowledge such as from navigators [7], and from the imaged object itself in so-called autofocus [3] or in data-consistency (DC) [9] methods. Unfortunately, these methods face challenges in

solving poorly conditioned, non-convex optimization problems with high computational complexity. Image-based deep learning (DL) models, like convolutional neural networks (CNNs) [1, 14] and generative adversarial networks [2], have been employed for direct motion-refined image estimation to overcome the challenges mentioned earlier. While promising, image-based DL-based methods can be unstable, struggle with pathology preservation, and introduce undesired alterations or hallucinations [4], raising concerns about their suitability for clinical applications. Specific approaches break down the problem into intermediate steps to address instability, maintain DC, and alleviate network hallucinations. For instance, CNNs are employed to identify subsets of k-space lines affected by motion, which can either be excluded [21] or down-weighted [5] in the context of a DC-based reconstruction. However, these methods do not explicitly model motion (transformations) during reconstruction, potentially compromising quality with many motion-corrupted lines. Efforts to improve efficiency integrate image-based MoCo CNNs into DC-based [8] or autofocus [15] iterative algorithms. Alternatively, they replace the optimization-based motion estimation with a DL network and incorporate this information into the reconstruction process [10]. However, they rely on supervised training and require large paired datasets, which are difficult to obtain in the medical domain. Recently, Implicit neural representations (INRs) were used to model images as continuous functions, spurring significant interest in instance-optimization-based reconstruction [19]. INRs do not need to be population-trained and have proven successful in self-supervised MRI reconstruction [6, 11], addressing issues present in prior methods, such as hallucinations. Despite this, the potential of INRs for MoCo remains untapped.

In this work, we introduce a novel MoCo pipeline utilizing INRs to address motion artifacts effectively while preserving anatomical structures in MRI scans. Our contributions are 1) Developing a motion-guided INR-based method called IM-MoCo, combining a k-space line motion detection network (*k*LD-Net), an Image INR, and a motion-INR. 2) Demonstrating its efficacy on the NYU *fastMRI* dataset [13, 25] through significant reductions in simulated motion-induced distortions compared to competing methods. 3) Highlighting its potential by improving pathology classification of *fastMRI+* annotations [26]. These findings position IM-MoCo as a promising advancement for enhancing MRI quality for scans with light and heavy motion corruption.

2 Materials & Methods

An overview of the proposed IM-MoCo pipeline is depicted in Fig. 1. Our pipeline comprises three steps: 1) Motion Detection with *k*LD-Net, 2) Generation of motion-free intensity priors with an Image INR, and 3) Forward application of predicted motion from a Motion INR.

2.1 Physical Model & Motion Simulation

The physical forward model, $K = \mathbf{F}I + \epsilon$, of a 2D rigid motion-induced MRI acquisition $K \in \mathbb{C}^{N_x \times N_y \times N_c}$ with N_x and N_y being the number of points in the




Fig. 1. IM-MoCo pipeline. The pre-trained *k*LD-Net takes a motion-corrupted *k*-space and outputs a motion mask, which is post-processed to yield a list of movement groups indicated by the different colors in *S* for better visualization. The number of movements and corresponding lines guide the Motion INR and the Image INR. The Image INR predicts the motion-free image and the Motion INR is used for guidance by optimizing the forward motion yielding a motion-corrupted *k*-space. The discrepancy between the motion-corrupted and the measured *k*-space is minimized using the data-consistency (DC) loss. The gradient entropy is a denoiser on the Image-INR output to impose crisp image priors. The final motion-corrected image is the output of the Image INR.

frequency and phase encoding directions, respectively, and N_c being the number of coils, can be described as follows:

$$K = \sum_{t=1}^T S_t \odot \mathcal{F} \mathbf{C} \mathbf{M}_t I, \quad (1)$$

where $I \in \mathbb{C}^{N_x \times N_y \times N_c}$ is the desired motion free image, $\mathbf{M}_t$ is the motion transform at time point t , $\mathbf{C} \in \mathbb{C}^{N_x \times N_y \times N_c}$ are the coil sensitivity maps, $\mathcal{F}$ is the Fourier transform, $\odot$ a sampling operator, and S_t is the sampling mask at time point t . The motion transform $\mathbf{M}_t$ combines rotation and planar translations in x and y directions. The sampling mask S_t is a binary mask indicating the *k*-space lines acquired at time t . Equation (1) describes motion in the image domain. However, the forward model can also be formulated in *k*-space as it is mathematically equivalent [18]. Irrespective of the space in which motion is

described, for the simulation of motion, it is crucial to synchronize the sampling of motion S_t with the MR k -space filling scheme to simulate realistic artifacts accurately. This work follows previously developed methods [16] and relies on the simulation of rigid-body motion in 2D single-coil cartesian sampled k -spaces filled sequentially left-to-right.

2.2 IM-MoCo

Motion Detection. In the first step of our pipeline, we present a k -space Line Detection Network [5, 21] (k LD-Net) to detect motion-corrupted lines in raw k -space data following a left-to-right sequential cartesian sampling scheme. The k LD-Net, built on a U-Net [22], takes a complex-valued motion-corrupted k -space as input, where the real and imaginary parts are concatenated in the channel dimension, and outputs a binary mask predicting the motion-corrupted lines denoted as S in Eq. (1). The four layers of the U-Net consist of convolutions with a kernel size of 3×3 , batch normalization, ReLU activation, and an average pooling layer with kernel size 2×2 . The number of channels (starting from 16) is doubled after each down-sampling operation and halved after each up-sampling operation. A final 1×1 convolution outputs the predicted mask $\hat{S}$. The network is trained with the binary cross entropy loss with logits between the predicted and ground truth masks. During inference, the raw prediction of the network is activated with *sigmoid* and thresholded (> 0.5) to yield a binary prediction mask. We consider a line corrupted if 20% of its frequencies are classified as so, and we generate an index list, e.g., $[1, 1, 0, 0, 0, 1, 1, 1]$, where 1 indicates a motion-corrupted line and 0 a motion-free one. We then post-process it to yield a list of movement groups, where each group is a list of line indices for one movement, e.g., $[[1, 1, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 1, 1, 1]]$ for the previously mentioned example. The following steps use the number of movements n and corresponding lines to guide the motion correction process.

Motion Correction. We use $\hat{S}$ from the k LD-Net (step 1) in a MoCo model, which is a combination of an Image INR (step 2) that acts as an implicit image prior and a Motion INR (step 3) to search for motion that satisfies the forward model (see Sec. 2.1). We can rewrite the forward model $\mathbf{F}$ from Eq. (1), for single-coil scans, as $\mathbf{F}_{\theta, \psi}$ as follows:

$$\hat{K} = \sum_{t=1}^T \hat{S}_t \odot \mathcal{F} \text{INR}_{\theta}(h_{\mathbf{M}}(n_{\mathbf{M}}, \mathbf{x}))_t \text{INR}_{\psi}(h_I(\mathbf{x})), \quad (2)$$

where INR_{ψ} is the image INR with optimizable parameters ψ , $h_I(\mathbf{x})$ is an encoding of the 2D coordinate $\mathbf{x}$, INR_{θ} is the motion INR with optimizable parameters θ , $h_{\mathbf{M}}(n_{\mathbf{M}}, \mathbf{x})$ is an encoding of the 2D coordinate $\mathbf{x}$ and $n_{\mathbf{M}}$, which is a sequence of linearly spaced numbers from -1 to 1 representing the number of movement groups. $\hat{S}_t$ is the predicted motion mask at time t . Each INR is built from MPLs taking encodings as input. To encode our data, we use hash grid encoding [20] for $h_I(\mathbf{x})$ and $h_{\mathbf{M}}(n_{\mathbf{M}}, \mathbf{x})$. This encoding introduces learnable feature grids that

enable a speed-up by replacing large MLPs with a look-up table of feature vectors and a much smaller MLP. The Image INR consists of three layers with 256 channels and ReLU activations, while the Motion INR has three layers with 64 channels and tanh activations. The INR_ψ predicts a 2-channel complex-valued intensity image, while INR_θ predicts n transformation grids for n movements.

In each iteration of the pipeline, we optimize the INR_θ and INR_ψ in an end-to-end fashion by feeding the encoding and the mask $\hat{S}$ through the forward model $\mathbf{F}_{\theta,\psi}$ which in term predicts a motion-free intensity image $\hat{I}$, applies the motion to the intensity image resulting in a motion corrupted image that is transferred to the k -space domain by the Fourier transform yielding $\hat{K}$. Using the DC-loss in k -space, we measure the discrepancy to the acquired k -space K_{acq} . We use the gradient entropy on $\hat{I}$ for regularization in image-space. The motion correction loss is then defined as follows:

$$\mathcal{L} = \underbrace{\frac{1}{N} \sum_{i=1}^N \|K_{acq,i} - \hat{K}_i\|_2^2}_{\text{DC-loss in } k\text{-space}} + \underbrace{\lambda \left(- \sum_{i=1}^N \nabla \hat{I}_i \cdot \log(\nabla \hat{I}_i) \right)}_{\text{Regularization in image-space}}, \quad (3)$$

where N is the number of data points, ∇ is the gradient operator, and λ is the regularization weight. The regularization is used to introduce image priors to the Image INR as a denoiser, forcing the INR to learn an artifact-free image. To ensure better motion estimates, high-frequencies are suppressed at the beginning of the optimization by setting a large weight λ and halving it every s step after i iterations [17]. The model is optimized until convergence, and the final motion-corrected image is the output of INR_ψ .

3 Experimental Results

3.1 Datasets & Motion Simulation

To evaluate our method, we use the MRI T_2 -weighted brain k -space data from the open-source NYU *fastMRI* database [13, 25] in our first experiment. We employed 300 2D slices from different patients and split them into 200/50/50 for training/validation/testing. We cropped the images to 320×320 and combined the coils. The second dataset contains T_1 -weighted and FLAIR brain *fastMRI* scans, which are part of the *fastMRI+* [26] annotations and were used for the downstream classification task. We classify two pathologies, *Nonspecific White Matter Lesion* and *Craniotomy*, resulting in 1116 slices of 60 subjects with a total of 2851 annotations. We split the slices into 889 (1460 annotations) for training, 174 (467 annotations) for validation, and 50 (141) for testing. For all experiments, two motion scenarios were simulated [16] with random rotations and translations between -10 and $+10$ in $^\circ$ and mm , respectively. In the Light Motion scenario between 6 and 10 movements were introduced during simulation, whereas the Heavy Motion scenario contained between 16 and 20 movements.

3.2 Experimental Setup

Experiment I: Motion Correction in Simulated Data. Here the image correction quality is assessed. We implemented the following methods: 1) AF [3], 2) U-Net [1], 3) AF+ [15], 4) IM-MoCo (ours). AF is an instance-wise autofocus-ing approach and optimizes motion parameters using Adam for 100 iterations. The U-Net is a population-trained model, mapping motion-corrupted images to motion-free ones, for 200 epochs using Adam with $3e^{-4}$ learning rate and an L1-SSIM loss. AF+ combines AF and the U-Net for deep image priors, trained with default settings using the available codebase [15]. IM-MoCo is optimized for 200 iterations using Adam with learning rates of $1e^{-2}$ for the INRs and an initial $\lambda = 1e^{-2}$, halved every 10th step after 100 iterations. All methods were evaluated on the T_2 -weighted 50 corrupted test scans for the two motion scenarios. Results were evaluated using the structural similarity index (SSIM), peak-signal-to-noise-ratio (PSNR) and the Haar Perceptual Similarity Index (HaarPSI) [12].

Experiment II: Downstream Classification Task. In this experiment, we demonstrate clinical relevance with a classification of pathologies in motion and motion-corrected images. We extract patches of size 124×124 around bounding boxes of the *fastMRI+* dataset and normalized between $[0, 1]$. We used a pre-trained *ResNet18* (PyTorch Hub¹) backbone for feature extraction and added one trainable linear layer as the classification head. Using the SGD optimizer, we trained the model for 100 epochs with a batch size of 20 and a learning rate of $1e^{-3}$. The trained model is then applied to the test set. We report the SSIM, PSNR, and HaarPSI for extracted patches around the bounding boxes and accuracy for the classification.

Implementation Details. The *kLD*-Net was population-trained with Adam (learning rate of $1e^{-4}$) for 4200 epochs and a batch size of 4. We used PyTorch for implementations and an NVIDIA GeForce RTX 2080 TI with 12 GB of VRAM. The code is publicly available on Github².

3.3 Results

Experiment I: Motion Correction in Simulated Data. Table 1 shows quantitative results from the motion correction experiment. For light and heavy motion, corrupted images yield SSIM values of 87% and 74%, PSNR values of 28db and 24db, and HaarPSI values of 70% and 56%, respectively. AF and U-Net improve SSIM by 7% and 10%, PSNR by 5db and 3db, and HaarPSI by 18% and 16%, respectively. However, AF+ worsens results by ca. 2% across all metrics. Compared to the best baseline, AF, our method improves SSIM/HaarPSI by $> 4\%$ for SSIM, $> 9\%$ for HaarPSI, and $> 10\text{ db}$ for PSNR. A qualitative comparison is shown in Fig. 2.

Experiment II: Downstream Classification Task. The results are summarized in Table 2. Motion-free images achieve an accuracy of 97%. For corrupted

¹ https://pytorch.org/hub/pytorch_vision_resnet/

² https://github.com/multimodallearning/MICCAI24_IMMoCo.git




Fig. 2. The visualization shows the median results of motion-corrected images of our IM-MoCo pipeline besides motion-corrupted, ground truth, and comparison methods. The first and third rows show the light and heavy correction results, respectively. The second and fourth rows show the residual error images.

Table 1. Quantitative results of experiment I: motion correction for all methods and motion scenarios. We report the results as mean $\pm$ standard deviation over the test set. The arrows indicate the direction of improvement.

Scenario	Method	SSIM $\uparrow$	PSNR (db) $\uparrow$	HaarPSI $\uparrow$
Light	Motion corrupted	87.26 ± 4.42	28.34 ± 2.97	70.48 ± 8.69
	AF [3]	94.47 ± 2.06	33.91 ± 2.37	88.49 ± 4.11
	U-Net [1]	91.39 ± 2.14	30.58 ± 2.33	81.58 ± 4.49
	AF+ [15]	85.18 ± 4.75	27.93 ± 2.79	70.82 ± 8.83
	IM-MoCo (ours)	98.25 ± 1.25	40.06 ± 3.33	97.20 ± 4.05
Heavy	Motion corrupted	74.06 ± 6.36	24.28 ± 2.50	56.56 ± 8.68
	AF [3]	87.19 ± 3.51	29.84 ± 2.32	78.91 ± 6.21
	U-Net [1]	84.55 ± 3.63	27.40 ± 2.34	72.72 ± 5.69
	AF+ [15]	70.61 ± 8.30	24.12 ± 2.67	56.91 ± 10.35
	IM-MoCo (ours)	92.77 ± 3.59	33.06 ± 3.59	87.29 ± 9.38

images, the accuracy is 96% for light and 94% for heavy motion, which is due to image quality loss as reflected in SSIMs of 89% and 77%, respectively. The U-Net’s accuracy is 90% for light and 88% for heavy while its SSIMs are 87% and 79%. IM-MoCo achieves an accuracy of 97% for light and 96% for heavy motion, with the highest SSIMs of 99% and 95%, respectively.

Table 2. Quantitative results of experiment II: image quality and classification accuracy improvements in patches of the *fastmri+* annotations. The arrows indicate the direction of improvement.

Scenario	Method	SSIM $\uparrow$	PSNR (db) $\uparrow$	HaarPSI $\uparrow$	Accuracy $\uparrow$
Motion-free	n.A.	n.A.	n.A.	n.A.	97.16
Light	Motion corrupted	89.93 $\pm$ 4.67	28.29 $\pm$ 4.07	76.12 $\pm$ 9.95	96.32
	U-Net [1]	87.37 $\pm$ 4.31	25.80 $\pm$ 2.61	70.87 $\pm$ 7.74	90.44
	IM-MoCo (ours)	99.00 $\pm$ 1.82	44.82 $\pm$ 6.44	97.33 $\pm$ 5.56	97.06
Heavy	Motion corrupted	77.03 $\pm$ 5.74	23.56 $\pm$ 2.18	58.87 $\pm$ 6.20	94.12
	U-Net [1]	79.45 $\pm$ 4.50	23.70 $\pm$ 2.18	59.82 $\pm$ 5.88	88.24
	IM-MoCo (ours)	95.26 $\pm$ 3.31	34.56 $\pm$ 5.61	88.48 $\pm$ 7.80	96.32

4 Discussion & Outlook

Experiment I demonstrated that our method effectively enhanced the quality of motion-corrupted images, surpassing comparison methods, which is mainly contributed by the implicit image priors and DC with the acquired k -space in the forward motion model. Notably, our approach exhibited robustness in addressing heavier motion scenarios as evidenced in Fig. 2 and Table 1. However, the efficacy of our method hinges on the accurate detection of the k lD-Net, which is contingent upon the acquisition sequence and motion simulation pattern. In Experiment II, we showcased the performance enhancement of a classification task compared to both corrupted and U-Net-corrected images as indicated in Table 2. This improvement in classification accuracy can be attributed to the prevention of overfitting on "healthy" features in instance-optimization methods in contrast to population training ones like the U-Net. Although our work primarily revolved around 2D rigid motion, we are confident in the method’s potential for extension to more complex motion, e.g., 3D rigid and non-rigid movements. This is feasible by adapting the INR’s input coordinates and applying 3D transformations. Experiments were conducted on single-coil data, but the extension to multi-coil can be achieved by incorporating coil sensitivity maps following the third step of the pipeline. Future investigations will explore the utilization of different sampling patterns, such as radial sampling, and the adaptation of the motion detection network to these patterns or pattern-independent methods. While real motion cases were not tested in this study, we are optimistic regarding the adaptability of our method through adjustments to the motion detection network.

5 Conclusion

We have introduced IM-MoCo, a pipeline that utilizes motion-guided INRs to mitigate the impact of motion artifacts in MRI scans. We tested our method on the *fastMRI* dataset and observed improved image quality and pathology classification performance. Our results suggest that IM-MoCo is a promising solution for enhancing the quality of MRI scans in the presence of challenging motion artifacts. In the future, we plan to investigate the applicability of our method to real 3D motion cases to validate its potential for clinical applications.

Acknowledgments. This work was supported by a grant from the French-German joint call for proposals on Artificial Intelligence. The authors thank the French Ministry of Higher Education, Research and Innovation (MESRI) and the German Ministry of Education and Research (BMBF) for funding this research. The work was carried out under the MEDICARE project (MESRI funding code: ANR-21-FAI1-0007, BMBF funding code: 01IS21094).

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

References

1. Al-Masni, M.A., Lee, S., Yi, J., Kim, S., Gho, S.M., Choi, Y.H., Kim, D.H.: Stacked U-Nets with self-assisted priors towards robust correction of rigid motion artifact in brain MRI. *NeuroImage* **259**, 119411 (2022)
2. Armanious, K., Jiang, C., Fischer, M., Küstner, T., Hepp, T., Nikolaou, K., Gatidis, S., Yang, B.: MedGAN: Medical image translation using GANs. *Computerized medical imaging and graphics* **79**, 101684 (2020)
3. Atkinson, D., Hill, D.L., Stoyale, P.N., Summers, P.E., Clare, S., Bowtell, R., Keevil, S.F.: Automatic compensation of motion artifacts in MRI. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* **41**(1), 163–170 (1999)
4. Chen, Z., Pawar, K., Ekanayake, M., Pain, C., Zhong, S., Egan, G.F.: Deep learning for image enhancement and correction in magnetic resonance imaging—state-of-the-art and challenges. *Journal of Digital Imaging* **36**(1), 204–230 (2023)
5. Eichhorn, H., Hammernik, K., Speiker, V., Epp, S.M., Rueckert, D., Preibisch, C., Schnabel, J.A.: Physics-Aware Motion Simulation For T2*-Weighted Brain MRI. In: *International Workshop on Simulation and Synthesis in Medical Imaging*. pp. 42–52. Springer (2023)
6. Feng, R., Wu, Q., Feng, J., She, H., Liu, C., Zhang, Y., Wei, H.: IMJENSE: Scan-specific Implicit Representation for Joint Coil Sensitivity and Image Estimation in Parallel MRI. *IEEE Transactions on Medical Imaging* (2023). <https://doi.org/10.1109/TMI.2023.3342156>
7. Gallichan, D., Marques, J.P., Gruetter, R.: Retrospective correction of involuntary microscopic head movement using highly accelerated fat image navigators (3D FatNavs) at 7T. *Magnetic resonance in medicine* **75**(3), 1030–1039 (2016)

8. Haskell, M.W., Cauley, S.F., Bilgic, B., Hossbach, J., Splitthoff, D.N., Pfeuffer, J., Setsompop, K., Wald, L.L.: Network accelerated motion estimation and reduction (NAMER): convolutional neural network guided retrospective motion correction using a separable motion model. *Magnetic resonance in medicine* **82**(4), 1452–1461 (2019)
9. Haskell, M.W., Cauley, S.F., Wald, L.L.: TArgeted Motion Estimation and Reduction (TAMER): data consistency based motion mitigation for MRI using a reduced model joint optimization. *IEEE transactions on medical imaging* **37**(5), 1253–1265 (2018)
10. Hossbach, J., Splitthoff, D.N., Cauley, S., Clifford, B., Polak, D., Lo, W.C., Meyer, H., Maier, A.: Deep learning-based motion quantification from k-space for fast model-based magnetic resonance imaging motion correction. *Medical Physics* **50**(4), 2148–2161 (2023). <https://doi.org/10.1002/mp.16119>
11. Huang, W., Li, H.B., Pan, J., Cruz, G., Rueckert, D., Hammernik, K.: Neural Implicit k-Space for Binning-Free Non-Cartesian Cardiac MR Imaging. In: *International Conference on Information Processing in Medical Imaging*. pp. 548–560. Springer (2023)
12. Kastrýulin, S., Zakirov, J., Pezzotti, N., Dylov, D.V.: Image Quality Assessment for Magnetic Resonance Imaging. *IEEE Access* **11**, 14154–14168 (2023). <https://doi.org/10.1109/access.2023.3243466>
13. Knoll, F., Zbontar, J., Sriram, A., Muckley, M.J., Bruno, M., Defazio, A., Parente, M., Geras, K.J., Katsnelson, J., Chandarana, H., et al.: fastmri: A publicly available raw k-space and dicom dataset of knee images for accelerated mr image reconstruction using machine learning. *Radiology: Artificial Intelligence* **2**(1), e190007 (2020)
14. Küstner, T., Armanious, K., Yang, J., Yang, B., Schick, F., Gatidis, S.: Retrospective correction of motion-affected MR images using deep learning frameworks. *Magnetic resonance in medicine* **82**(4), 1527–1540 (2019)
15. Kuzmina, E., Razumov, A., Rogov, O.Y., Adalsteinsson, E., White, J., Dylov, D.V.: Autofocusing+: Noise-resilient motion correction in magnetic resonance imaging. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 365–375. Springer (2022)
16. Lee, S.B., Jung, S., Jung, K.J., Kim, D.H.: Deep Learning in MR Motion Correction: a Brief Review and a New Motion Simulation Tool (view2Dmotion). *Investigative Magnetic Resonance Imaging* **24**, 196–206 (2020)
17. Lin, C.H., Ma, W.C., Torralba, A., Lucey, S.: Barf: Bundle-adjusting neural radiance fields. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5741–5751 (2021)
18. Loktyushin, A., Nickisch, H., Pohmann, R., Schölkopf, B.: Blind retrospective motion correction of MR images. *Magnetic resonance in medicine* **70**(6), 1608–1618 (2013)
19. Molaei, A., Aminimehr, A., Tavakoli, A., Kazerouni, A., Azad, B., Azad, R., Merhof, D.: Implicit Neural Representation in Medical Imaging: A Comparative Survey. In: *2023 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*. pp. 2373–2383 (2023). <https://doi.org/10.1109/ICCVW60793.2023.00252>
20. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)* **41**(4), 1–15 (2022)
21. Oksuz, I., Clough, J., Ruijsink, B., Puyol-Antón, E., Bustin, A., Cruz, G., Prieto, C., Rueckert, D., King, A.P., Schnabel, J.A.: Detection and correction of cardiac

- MRI motion artifacts during reconstruction from k-space. In: International conference on medical image computing and computer-assisted intervention. pp. 695–703. Springer (2019)
22. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18. pp. 234–241. Springer (2015)
 23. Slipsager, J.M., Glimberg, S.L., Sogaard, J., Paulsen, R.R., Johannesen, H.H., Martens, P.C., Seth, A., Marner, L., Henriksen, O.M., Olesen, O.V., et al.: Quantifying the financial savings of motion correction in brain mri: a model-based estimate of the costs arising from patient head motion and potential savings from implementation of motion correction. *Journal of Magnetic Resonance Imaging* **52**(3), 731–738 (2020)
 24. Zaitsev, M., Maclaren, J., Herbst, M.: Motion artifacts in MRI: A complex problem with many partial solutions. *Journal of Magnetic Resonance Imaging* **42**(4), 887–901 (2015). <https://doi.org/10.1002/jmri.24850>
 25. Zbontar, J., Knoll, F., Sriram, A., Murrell, T., Huang, Z., Muckley, M.J., Defazio, A., Stern, R., Johnson, P., Bruno, M., et al.: fastMRI: An open dataset and benchmarks for accelerated MRI. arXiv preprint arXiv:1811.08839 (2018)
 26. Zhao, R., Yaman, B., Zhang, Y., Stewart, R., Dixon, A., Knoll, F., Huang, Z., Lui, Y.W., Hansen, M.S., Lungren, M.P.: fastMRI+, Clinical pathology annotations for knee and brain fully sampled magnetic resonance imaging data. *Scientific Data* **9**(1), 152 (2022)