



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

# Physics-Informed Deep Learning for Motion-Corrected Reconstruction of Quantitative Brain MRI

Hannah Eichhorn<sup>1,2,\*</sup>, Veronika Spieker<sup>1,2</sup>, Kerstin Hammernik<sup>2</sup>, Elisa Saks<sup>3</sup>, Kilian Weiss<sup>4</sup>, Christine Preibisch<sup>3</sup>, and Julia A. Schnabel<sup>1,2,3,5</sup>

<sup>1</sup> Institute of Machine Learning in Biomedical Imaging, Helmholtz Munich, Germany

<sup>2</sup> School of Computation, Information & Technology, Technical University of Munich, Germany

<sup>3</sup> School of Medicine & Health, Technical University of Munich, Germany

<sup>4</sup> Philips GmbH Market DACH, Hamburg, Germany

<sup>5</sup> School of Biomedical Engineering & Imaging Sciences, King's College London, UK

\*[hannah.eichhorn@tum.de](mailto:hannah.eichhorn@tum.de)

**Abstract.** We propose PHIMO, a physics-informed learning-based motion correction method tailored to quantitative MRI. PHIMO leverages information from the signal evolution to exclude motion-corrupted k-space lines from a data-consistent reconstruction. We demonstrate the potential of PHIMO for the application of T2\* quantification from gradient echo MRI, which is particularly sensitive to motion due to its sensitivity to magnetic field inhomogeneities. A state-of-the-art technique for motion correction requires redundant acquisition of the k-space center, prolonging the acquisition. We show that PHIMO can detect and exclude intra-scan motion events and, thus, correct for severe motion artifacts. PHIMO approaches the performance of the state-of-the-art motion correction method, while substantially reducing the acquisition time by over 40%, facilitating clinical applicability. Our code is available at <https://github.com/compai-lab/2024-miccai-eichhorn>.

**Keywords:** Self-Supervised Learning · Motion Detection · Data-Consistent Reconstruction · T2\* Quantification · Gradient Echo MRI

## 1 Introduction

Quantitative magnetic resonance imaging (MRI) estimates physical tissue properties from a series of qualitative images with varying imaging parameters. In contrast to commonly employed qualitative structural imaging, this facilitates a consistent extraction of potential biomarkers across scanners and hospitals. With typical voxel sizes of 2-3 mm, the resolution of quantitative MRI is commonly lower than for qualitative MRI due to longer acquisition times.

Patient motion is a challenge for brain MRI in general, potentially impeding successful diagnoses. Quantitative MRI is particularly sensitive to motion due to its intrinsically long imaging times. In recent years, deep learning-based

approaches have demonstrated promising results for motion correction (MoCo) of brain MRI [17]. MoCo has previously been addressed as an image denoising problem, with convolutional [2,19] or generative adversarial networks [9,12,15]. Yet, exclusively relying on image data, these methods cannot ensure consistency with the acquired raw k-space data, posing a potential obstacle to their clinical translation. Data consistency (DC) is only achievable by integrating MoCo in the image reconstruction process [5,8,16]. Nevertheless, previous methods have mostly been developed for higher resolution qualitative MRI or - in the context of quantitative MRI - they do not enforce DC [19].

In this work, we propose physics-informed motion correction (PHIMO), which utilises information from the quantitative parameter estimation process to detect and exclude motion-corrupted k-space lines from a data-consistent reconstruction. We demonstrate the potential of PHIMO for the application of T2\* quantification from gradient echo (GRE) MRI, which enables oxygenation-sensitive imaging as part of the multi-parametric quantitative BOLD (mqBOLD) protocol [7]. GRE MRI is highly sensitive to motion due to the influence of magnetic field inhomogeneities, especially for larger echo times [13]. The current mqBOLD MoCo method [14] relies on redundant k-space acquisition, which significantly increases the total acquisition time from 3 min 39 s to 6 min 25 s.

To the best of our knowledge, this is the first MoCo approach that incorporates information from the quantitative parameter estimation process as a physics-informed loss. Our contributions are three-fold:

1. We train a single unrolled reconstruction network to recover high-quality images from undersampled k-space data for varying rates of excluded lines.
2. We train a multi-layer perceptron (MLP) to predict exclusion masks for motion-corrupted lines. Utilising the T2\* decay information allows us to perform MoCo in a self-supervised fashion for each subject individually.
3. We evaluate PHIMO on multicoil raw data acquired with and without head motion and compare the results to the current mqBOLD MoCo which requires redundant data acquisition.

## 2 Background

### 2.1 Motion during MR Image Acquisition

The MRI multicoil forward model in the presence of motion includes the sampling mask  $\mathbf{S}_t$ , incorporating the line-wise k-space acquisition pattern, the Fourier transform  $\mathcal{F}$ , the coil sensitivity maps  $\mathbf{C}$ , and the motion transform  $\mathbf{U}_t$ , which are applied to the motion-free image  $x$  for each time point  $t$ , yielding the motion-affected k-space data  $\hat{y}$  [1]:

$$\hat{y} = \sum_{t=1}^T \mathbf{S}_t \mathcal{F} \mathbf{C} \mathbf{U}_t x \quad (1)$$

Head motion is typically treated as rigid-body motion with random timings due to the patient’s carelessness or discomfort. Consequently,  $\mathbf{U}_t$  consists of rotation

and translation transforms, but can also incorporate second-order motion effects, such as motion-induced field inhomogeneity changes [3]. Under the assumption that the patient approximately stays in one position except for individual motion events, an undersampling mask  $\Omega$  exists, which excludes the individual motion events, so that the masked motion-corrupted data is close to the masked motion-free data  $y$ :

$$\Omega \hat{y} \simeq \Omega y \quad (2)$$

## 2.2 T2\* Quantification from Multi-Echo GRE MRI

The effective transverse relaxation time T2\* can be quantified from a series of GRE images  $x = [x_1, \dots, x_E]$  acquired at various echo times  $t_e$  for  $e = 1, \dots, E$ . The time evolution of the signal magnitude  $s_e = |x_e|$  for a single voxel is commonly modelled as a mono-exponential decay [13]:

$$s_e = s_0 \cdot \exp\left(-\frac{t_e}{T2^*}\right) \quad (3)$$

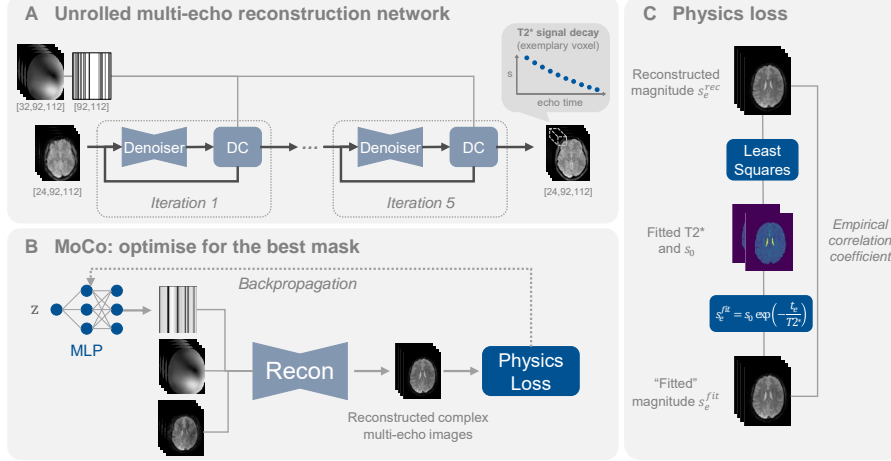
T2\* and the signal magnitude  $s_0$  at  $t_e = 0$  follow from least-squares fitting.

## 3 Methods

### 3.1 PHIMO

Inspired by MoCo through outlier-rejecting bootstrap aggregation [15], we use Eq. 2 to split MoCo into two subproblems: undersampled reconstruction and, in contrast to [15], a physics-informed detection of motion-corrupted k-space lines. Specifically, Eq. 2 allows us to train a reconstruction network on randomly undersampled *motion-free* data  $\Omega y$  and, during inference, recover a high-quality image from the undersampled *motion-corrupted* data  $\Omega \hat{y}$ , given that  $\Omega$  excludes the motion-corrupted k-space lines. An overview of PHIMO is provided in Fig. 1.

**Supervised Training of Unrolled Reconstruction** First, we train an unrolled multi-echo reconstruction network on randomly undersampled *motion-free* data. The reconstruction alternates five times between a CNN-based denoiser and a gradient descent DC step (see Fig. 1A), without weight sharing between iterations and the step sizes of the gradient descent layers implemented as learnable parameters [4]. The denoiser consists of five 2D convolutional layers (kernel size:  $3 \times 3$ , 64 filters) and ReLU activations. Real and imaginary parts of the input image, as well as multiple echoes are stacked in the channel dimension. The network is trained for 4000 epochs with Adam [11] and a batch size of eight. We use a learning rate of  $1 \times 10^{-4}$  and structural similarity (SSIM) [18] as loss, calculated on real and imaginary parts separately. During training, the rate of excluded lines in the undersampling masks is varied between 0.05 and 0.5, with a fixed fully-sampled center of 10 lines.



**Fig. 1.** Overview of PHIMO. (A) Training an unrolled reconstruction network on randomly undersampled *motion-free* multi-echo images. (B) MoCo of *motion-corrupted* data by optimising an MLP to predict slice-wise exclusion masks in a self-supervised fashion for each subject individually. (C) Calculation of empirical correlation coefficient as physics-informed loss for optimising the MLP in (B).

**Self-Supervised k-Space Line Detection** In the presence of *motion-corrupted* data, we perform MoCo by detecting and excluding motion-corrupted k-space lines from the reconstruction. Therefore, we train an MLP for each subject to predict an undersampling mask that excludes motion-corrupted k-space lines (Fig. 1B). We introduce a physics loss, which takes into account that for multi-echo GRE MRI, motion increasingly impacts the image with increasing echo time. Thus, motion disturbs the mono-exponential signal evolution (Eq. 3) [14]. As illustrated in Fig. 1C, the physics loss compares the reconstructed signal intensities  $s_e^{rec}$  and the “fitted” intensities  $s_e^{fit}$  resulting from inserting the fitted parameters  $T2^*$  and  $s_0$  into Eq. 3. We calculate the physics loss  $L_{phys}$  as the empirical correlation coefficient, with the mean over all echoes  $\bar{s} = \frac{1}{E} \sum_e s_e$ :

$$L_{phys} = \frac{\sum_e (s_e^{rec} - \bar{s}^{rec})(s_e^{fit} - \bar{s}^{fit})}{\sqrt{\sum_e (s_e^{rec} - \bar{s}^{rec})^2 (s_e^{fit} - \bar{s}^{fit})^2}} \quad (4)$$

$L_{phys}$  is calculated voxel-wise across echoes and averaged within a brain mask excluding cerebrospinal fluid (CSF). We combine  $L_{phys}$  with a regularisation  $L_{reg}$  on the variation of the predicted masks for adjacent slices. Following the interleaved, multi-slice acquisition scheme,  $L_{reg}$  is calculated between all  $Z$  pairs of predicted masks,  $\Omega_z$  and  $\Omega_{z+2}$ , within the batch:

$$L = L_{phys} + \lambda L_{reg} = L_{phys} + \lambda \sum_{z=1}^Z |\Omega_z - \Omega_{z+2}| \quad (5)$$

The MLP embeds the slice index  $z$  using a mapping function  $f: \mathbb{R}^1 \rightarrow \mathbb{R}^3$ , followed by two fully connected layers of size 23 and 46 and a sigmoid activation. It results in an output mask of size 92, with values ranging between 0 and 1. We do not binarise the output mask to allow for a smooth loss minimisation. Note that we also do not restrict the output masks to a fully-sampled center to allow for an exclusion of motion-corrupted lines in the k-space center. The MLP is optimised with Adam, a batch of 4 slices and a regularisation strength of  $\lambda = 0.1$ . Differentiable least-squares fitting is performed with `torch.linalg.lstsq`. We perform early stopping when  $L_{reg}$  does not improve for 50 epochs. Optimising the MLP takes less than 14 minutes per subject.

### 3.2 Data

We have acquired multi-coil k-space data from 15 volunteers ( $26.6 \pm 2.8$  years, 5 females) on a 3T Philips Elition X MR scanner (Philips Healthcare, Best, The Netherlands), using a multi-slice 2D GRE sequence (12 echoes,  $TE_1 = \Delta TE = 5$  ms,  $TR = 2300$  ms, voxel size:  $2 \times 2 \times 3$  mm, 32-channel head coil). The study has been approved by the local ethics committee (approval numbers 440/18 S-AS, 2023-386-S-SB). We have conducted repeated scans under two conditions: without voluntary head motion and with the subject instructed to move randomly, imitating sneezing or coughing. In one case, referred to as *motion timing experiment*, we have provided precise timings for the volunteer to move, aiming to assess the motion detection performance of PHIMO. Additionally, we have acquired half- and quarter-resolution data in both conditions to compare PHIMO to the current mQbOLD MoCo (“HR/QR-MoCo”) [14]. The datasets are divided subject-wise into train, validation, and test sets (6/2/7 subjects), including only slices with more than 10% brain voxels (193/61 train/validation slices). For testing, we manually exclude inferior slices to disentangle motion from severe susceptibility artifacts, which need to be addressed separately [6], resulting in 128 test slices. The input of the reconstruction network is normalised by the maximum magnitude of the zero-filled image.

### 3.3 Evaluation

We compare PHIMO to the state-of-the-art HR/QR-MoCo [14] and outlier-rejecting bootstrap aggregation (OR-BA) [15]. The latter averages reconstructions with 15 random masks, for which we train a separate unrolled reconstruction network with variable density masks at a fixed exclusion rate of 0.5. To differentiate between the methods’ performances for more severe and minor motion, we categorise the seven test subjects based on the visual quality of the motion-corrupted images, resulting in four subjects with severe and three with minor motion. T2\* maps are evaluated based on mean absolute error (MAE), SSIM and feature similarity (FSIM) [20]. Therefore, all acquisitions are aligned to the motion-free acquisition via 3D registration of the stacked slices. Wilcoxon signed rank tests and False-Discovery Rate correction are employed for statistical testing. HR/QR-MoCo and segmentation of anatomical scans are performed

in MATLAB (R2022b) and SPM12 with custom programs [10]. Other computations are performed in Python 3.8.12, using PyTorch 2.0.1 and MERLIN [4].

## 4 Experiments and Results


In Fig. 2 we compare exclusion masks derived from OR-BA and PHIMO with the reference mask inferred from the verbal instructions of the motion timing experiment. This example demonstrates that PHIMO can detect individual motion events. In contrast, OR-BA on average assigns nearly uniform weights to all k-space lines and cannot detect specific lines.

Example T2\* maps in Fig. 3 demonstrate that - in contrast to OR-BA - PHIMO effectively suppresses motion artefacts for the subject with severe motion, close to the performance of HR/QR-MoCo. In the case of minor motion, OR-BA leads to blurring, while PHIMO preserves the image quality and is comparable to the motion-free reference. Note that, despite the volunteer’s cooperation, residual motion may still be present in the so-called motion-free reference acquisition.

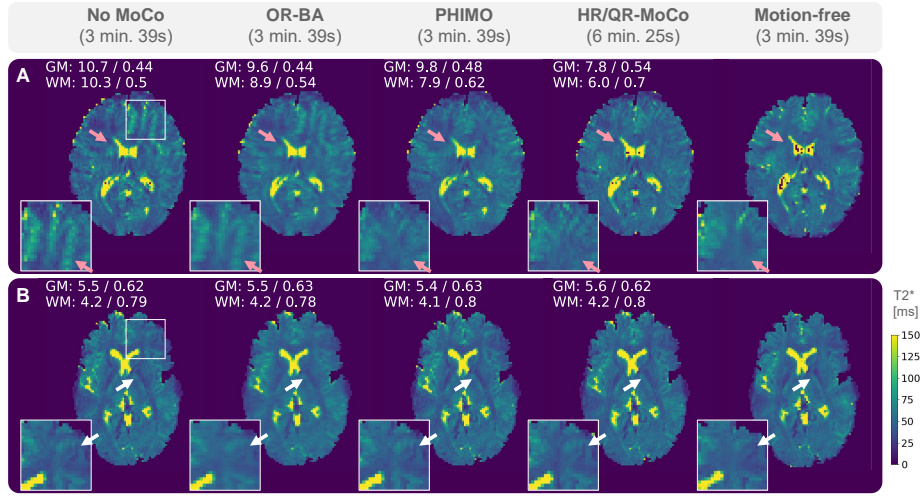
The qualitative observations are supported by the quantitative analysis in Fig. 4. In case of more severe motion, all metrics are notably improved for PHIMO, approaching the values of HR/QR-MoCo. OR-BA yields inconsistent results: for MAE it results in significantly but not substantially worse values than PHIMO, while PHIMO more clearly outperforms OR-BA for the structure-sensitive metrics SSIM and FSIM. For minor motion, the values for the motion-corrupted maps are already satisfactory, but PHIMO and HR/QR-MoCo achieve small improvements, particularly in white matter regions. The performance of OR-BA remains inconsistent across the evaluated metrics.

## 5 Discussion and Conclusion

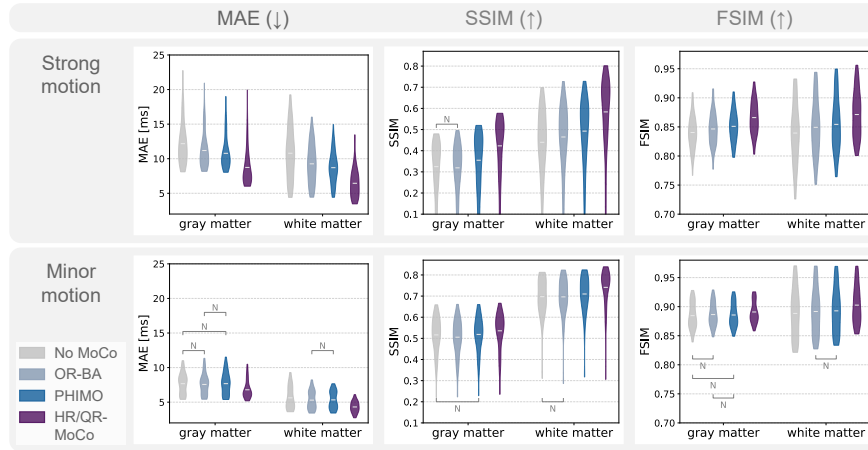
We propose PHIMO, a physics-informed MoCo method tailored to quantitative MRI. To the best of our knowledge, PHIMO is the first approach that incorporates information from the signal evolution in order to down-weight motion-corrupted k-space lines in a DC-based reconstruction. For the application of T2\*



**Fig. 2.** Exclusion masks of two example slices for the motion timing experiment (single subject), derived from OR-BA, PHIMO and the verbal instructions (reference). The corresponding reconstructions with reduced motion artefacts for PHIMO are provided in the supplementary material (Fig. S1).



**Fig. 3.** Example T2\* maps for subjects with severe (A) and minor motion (B). From left to right: motion-corrupted (No MoCo), OR-BA, proposed PHIMO, HR/QR-MoCo and motion-free reference. Pink arrows indicate (suppressed) motion artifacts, white arrows increased blurring for OR-BA. Gray/white matter (GM/WM) MAE values in ms as well as SSIM values relative to the motion-free maps are indicated in the top left corner. The corresponding exclusion masks of OR-BA and PHIMO are presented in the supplementary material (Fig. S2).



**Fig. 4.** Image quality metrics MAE, SSIM and FSIM, calculated relative to the motion-free T2\* maps, for the 4 test subjects with severe (top) and the 3 test subjects with minor motion (bottom). All metrics are calculated on T2\* maps registered to the motion-free reference and evaluated in gray and white matter, separately. Gray brackets indicate comparisons with no statistical significance (N) at a level of  $p = 0.001$ .

quantification from multi-echo GRE MRI, we have demonstrated that PHIMO can detect motion events, suppress strong motion artefacts and thus, significantly improve the quality of motion-corrupted T2\* maps.

To investigate the impact of our proposed physics-informed loss for excluding motion events and satisfying equation 2, we have compared PHIMO to OR-BA, which we have implemented as a bootstrap aggregation averaging reconstructions with random masks, in line with the idea of Oh et al. [15]. The visual examples in Fig. 3 and the structure- and feature-based quantitative metrics SSIM and FSIM in Fig. 4 demonstrate that PHIMO notably outperforms OR-BA. A reason for that is that OR-BA samples random masks and, thus, down-weights all k-space lines uniformly on average (Fig. 2), thereby lacking the ability to adequately detect and exclude corrupted lines.

MAE does not seem to reflect PHIMO’s - visually apparent - superiority to OR-BA in suppressing strong motion artefacts. In general, reliable image quality evaluation is an ongoing challenge in the MoCo and image reconstruction community [17]. In our case, MAE’s inconsistency could be due to PHIMO excluding central k-space lines in the case of detected motion events, which allows for a reduction of strong motion artefacts, as highlighted in Fig. 3. At the same time, excluding central k-space lines might lead to some T2\* estimation errors, since the k-space center contains overall signal intensity and contrast information. Nevertheless, it appears preferable to exclude severe motion outliers while slightly misestimating T2\* (PHIMO) rather than including motion measurements and not correcting severe artefacts (OR-BA).

We have further compared PHIMO to HR/QR-MoCo [14], which is the state-of-the-art within the mqBOLD technique [10]. HR/QR-MoCo combines the motion-corrupted scan with additionally acquired half- and quarter-resolution scans, effectively sampling the k-space center three times. The performance of PHIMO as well as HR/QR-MoCo depends on the exact motion pattern. Both approaches build on the assumption that the subject generally maintains a consistent position, interrupted by random motion events. Additionally, PHIMO is more severely impacted by subject motion that occurs during the acquisition of the k-space center, since excluding low-frequency contrast information affects the reconstruction quality more severely than missing high-frequency information on edges and details. In our qualitative and quantitative comparison, HR/QR-MoCo outperforms PHIMO for more severe and minor motion scenarios, as it essentially results in a weighted average of three acquisitions, which even for a motion-free scan leads to noise reduction. However, PHIMO achieves a competitive image quality, particularly when taking into account that it does not rely on additionally acquired half- and quarter-resolution scans. Consequently, PHIMO reduces the total acquisition time by over 40 %, which is a key for facilitating clinical applicability of the entire mqBOLD technique.

**Conclusion** We have introduced PHIMO, a MoCo method tailored to quantitative MRI, which utilises information from the MR signal evolution to detect motion events with a subject-specific, self-supervised MLP. We have demonstrated



PHIMO’s capabilities in detecting and adequately down-weighting motion-corrupted k-space lines in a DC-based reconstruction for T2\* quantification from GRE MRI. Compared to the state-of-the-art motion correction method, which requires redundant data acquisition, PHIMO achieves competitive image quality and reduces the overall acquisition time from 6 min 25 s to 3 min 39 s, positively impacting clinical workflow. Importantly, PHIMO is applicable to various other quantitative MRI tasks, as long as the parameter estimation is differentiable.

**Acknowledgments.** H.E. and V.S. are partially supported by the Helmholtz Association under the joint research school ”Munich School for Data Science - MUDS”.

**Disclosure of Interests.** K.W. is an employee of Philips GmbH Market DACH.

## References

1. Atkinson, D.: Retrospective correction of motion in MR images. In: van der Kouwe, A.J., Andre, J.B. (eds.) *Motion Correction in MR*, vol. 6, pp. 259–267. Academic Press (2022). <https://doi.org/10.1016/B978-0-12-824460-9.00022-4>
2. Chatterjee, S., Sciarra, A., Dünnwald, M., Oeltze-Jafra, S., Nürnberger, A., Speck, O.: Retrospective motion correction of MR images using prior-assisted deep learning. In: *Proceedings of the 34th Conference on NeurIPS (2020)*. <https://doi.org/10.48550/arXiv.2011.14134>
3. Eichhorn, H., Hammernik, K., Spieker, V., Epp, S.M., Rueckert, D., Preibisch, C., Schnabel, J.A.: Physics-aware motion simulation for T2\*-weighted brain MRI. In: *Simulation and Synthesis in Medical Imaging (2023)*. [https://doi.org/10.1007/978-3-031-44689-4\\_5](https://doi.org/10.1007/978-3-031-44689-4_5)
4. Hammernik, K., Küstner, T.: Machine enhanced reconstruction learning and interpretation networks (MERLIN). In: *Proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM) (2022)*
5. 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). <https://doi.org/10.1002/mrm.27771>
6. Hirsch, N.M., Preibisch, C.: T2\* mapping with background gradient correction using different excitation pulse shapes. *American Journal of Neuroradiology* **34**(6), E65–E68 (2013). <https://doi.org/10.3174/ajnr.A3021>
7. Hirsch, N.M., Toth, V., Förschler, A., Kooijman, H., Zimmer, C., Preibisch, C.: Technical considerations on the validity of blood oxygenation level-dependent-based MR assessment of vascular deoxygenation. *NMR in Biomedicine* **27**(7), 853–862 (2014). <https://doi.org/10.1002/nbm.3131>
8. 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 MRI motion correction. *Medical physics* (2022). <https://doi.org/10.1002/mp.16119>
9. Johnson, P.M., Drangova, M.: Conditional generative adversarial network for 3D rigid-body motion correction in MRI. *Magnetic Resonance in Medicine* **82**(3), 901–910 (2019). <https://doi.org/10.1002/mrm.27772>

10. Kaczmarz, S., Hyder, F., Preibisch, C.: Oxygen extraction fraction mapping with multi-parametric quantitative BOLD MRI: Reduced transverse relaxation bias using 3D-GraSE imaging. *NeuroImage* **220**, 117095 (2020). <https://doi.org/10.1016/j.neuroimage.2020.117095>
11. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Proceedings of the 3rd ICLR (2015). <https://doi.org/10.48550/arXiv.1412.6980>
12. 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). <https://doi.org/10.1002/mrm.27783>
13. Magerkurth, J., Volz, S., Wagner, M., Jurcoane, A., Anti, S., Seiler, A., Hattingen, E., Deichmann, R.: Quantitative T2\* -mapping based on multi-slice multiple gradient echo flash imaging: Retrospective correction for subject motion effects: Movement correction in T2\* mapping. *Magnetic Resonance in Medicine* **66**(4), 989–997 (2011). <https://doi.org/10.1002/mrm.22878>
14. Nöth, U., Volz, S., Hattingen, E., Deichmann, R.: An improved method for retrospective motion correction in quantitative T2\* mapping. *NeuroImage* **92**, 106–119 (2014). <https://doi.org/10.1016/j.neuroimage.2014.01.050>
15. Oh, G., Lee, J.E., Ye, J.C.: Unpaired MR Motion Artifact Deep Learning Using Outlier-Rejecting Bootstrap Aggregation. *IEEE Transactions on Medical Imaging* **40**(11), 3125–3139 (2021). <https://doi.org/10.1109/TMI.2021.3089708>
16. Singh, N.M., Dey, N., Hoffmann, M., Fischl, B., Adalsteinsson, E., Frost, R., Dalca, A.V., Golland, P.: Data consistent deep rigid MRI motion correction. In: *Medical Imaging with Deep Learning*. Proceedings of Machine Learning Research, vol. 227, pp. 368–381. PMLR (10–12 Jul 2024), <https://proceedings.mlr.press/v227/singh24a.html>
17. Spieker, V., Eichhorn, H., Hammernik, K., Rueckert, D., Preibisch, C., Karampinos, D.C., Schnabel, J.A.: Deep learning for retrospective motion correction in MRI: A comprehensive review. *IEEE Transactions on Medical Imaging* **43**(2), 846–859 (2024). <https://doi.org/10.1109/TMI.2023.3323215>
18. Wang, Z., Bovik, A., Sheikh, H., Simoncelli, E.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004). <https://doi.org/10.1109/TIP.2003.819861>
19. Xu, X., Kothapalli, Satya V. V. N., Liu, J., Kahali, S., Gan, W., Yablonskiy, D.A., Kamilov, U.S.: Learning-based motion artifact removal networks for quantitative R2\* mapping. *Magnetic Resonance in Medicine* **88**(1), 106–119 (2022). <https://doi.org/10.1002/mrm.29188>
20. Zhang, L., Zhang, L., Mou, X., Zhang, D.: FSIM: a feature similarity index for image quality assessment. *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society* **20**(8), 2378–2386 (2011). <https://doi.org/10.1109/TIP.2011.2109730>