MotionDPS: Motion-Compensated 3D MRI Reconstruction

Antonio Ortiz-Gonzalez*
Life and Medical Sciences Institute
University of Bonn
Germany, Bonn 53115
aortizgo@uni-bonn.de

Erich Kobler*
Institute for Machine Learning
Johannes Kepler University Linz
Austria, Linz 4040
erich.kobler@jku.at

Lukas Schletter
German Center for Neurodegenerative Diseases
Germany, Bonn 53127
Lukas.Schletter@dzne.de

Alexander Effland
Institute for Applied Mathematics
University of Bonn
Germany, Bonn 53115
effland@iam.uni-bonn.de

Abstract

We introduce MotionDPS, a unified Bayesian framework for motion-compensated 3D MRI that jointly estimates the anatomical image, rigid-body motion, and coil sensitivity maps directly from motion-corrupted k-space data. MotionDPS integrates pretrained 3D complex-valued score-based diffusion models as expressive anatomical priors within a physics-based forward model, tightly coupling learned generative modeling with data-consistent reconstruction. Inference alternates diffusion posterior image updates with proximal optimization steps for motion and coil estimation, enabling fully unsupervised reconstruction without requiring paired motion-free training data or external motion tracking. Experiments on simulated and in vivo brain MRI datasets demonstrate consistently improved image quality and motion robustness compared to state-of-the-art methods.

1 Introduction

Magnetic Resonance Imaging (MRI) is highly sensitive to patient motion due to long acquisition times and sequential k-space sampling. Even small movements introduce phase inconsistencies that lead to blurring, ghosting, and geometric distortions. These motion artifacts reduce diagnostic quality and require repeated scans in up to 20% of clinical examinations, incurring substantial costs [1].

Despite decades of research, retrospective motion correction remains a major unsolved challenge in clinical MRI. Classical approaches rely on explicit forward models with handcrafted regularization [2, 3], while learning-based methods can improve image quality but are often trained on synthetic corruption, operate slice-wise, and may not fully enforce 3D k-space consistency [4, 5]. Diffusion models provide powerful learned priors for inverse problems via posterior sampling [6], yet existing motion-aware diffusion methods are either limited to 2D settings or do not jointly infer all latent variables [7]. Consequently, a unified framework that jointly models anatomy, motion, and coil sensitivities within a fully 3D, physics-consistent formulation remains unexplored. To address this gap, we propose *MotionDPS*, a unified Bayesian framework for motion-compensated 3D MRI that jointly estimates the motion-free image, rigid motion trajectory, and coil sensitivity maps directly from motion-corrupted k-space. MotionDPS combines a pretrained 3D complex-valued diffusion prior

*These authors contributed equally.

The Third Austrian Symposium on AI, Robotics, and Vision (AIROV26).

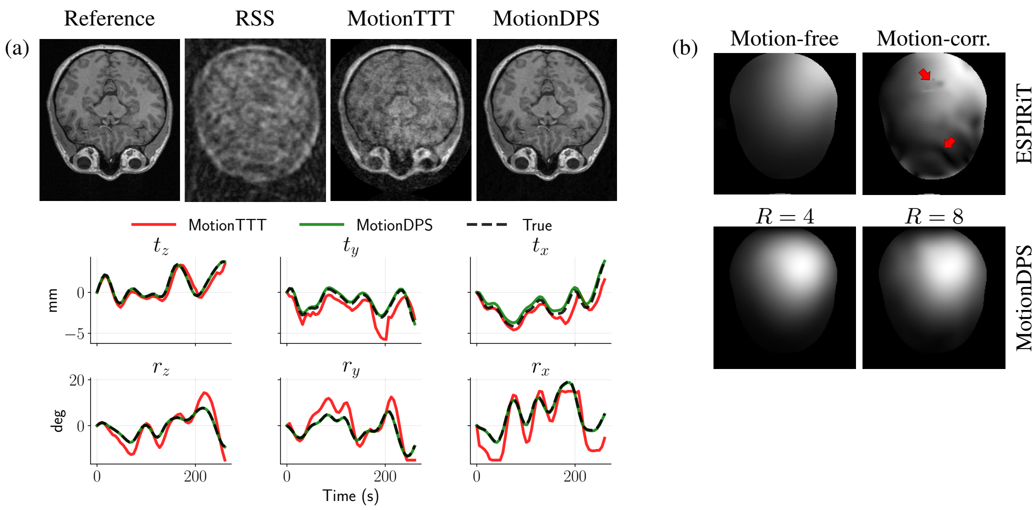

Figure 1: a) Reconstruction results for a volume from the CC359 dataset under severe motion. b) Coil sensitivity maps estimation for a FastMRI sample. In the presence of motion, the ESPIRiT estimate ($R = 4$) exhibits visible motion-induced artifacts (red arrows).

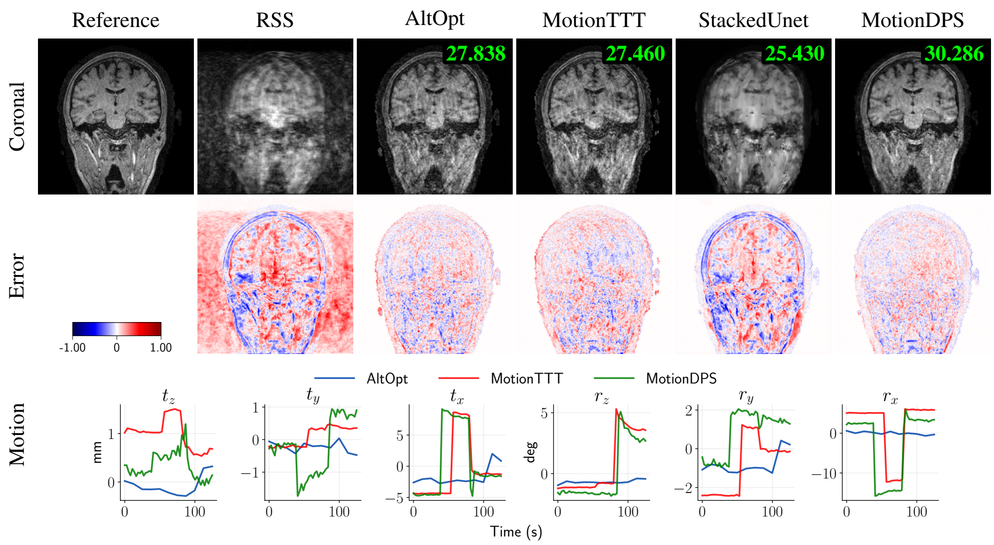

Figure 2: Coronal reconstructions for a severely motion-corrupted scan from the PMoC3D dataset. For each method, we show the reconstructed image, the error map, and the estimated motion parameters. PSNR values (green, top right) are computed relative to the motion-free reference image.

with physics-based data consistency and explicit smoothness priors for motion and coils. Inference alternates diffusion posterior image updates with efficient proximal optimization for motion and coil estimation. To our knowledge, this is the first fully 3D framework that jointly estimate all variables within a unified diffusion-based formulation.

2 Methods

Physics-based imaging and motion model: To model patient motion in accelerated multi-coil MRI, the acquisition is described by a sampling trajectory $\mathcal{T}$ specifying the k-space coefficients acquired at each discrete time point t. We assume that K coefficients are acquired simultaneously at any time point $t \in \{1, \ldots, T\}$, i.e. $\mathcal{T} \colon \{1, \ldots, T\} \to \{1, \ldots, D\}^K$. The k-space data acquired at a single time point $\mathbf{z}_t \in \mathbb{C}^{C\times K}$ read as

$$\mathbf{z}_t = \mathbf{M}_t\mathbf{F}\left(\mathbf{c}\odot\mathcal{W}(\mathbf{x},\mathbf{v}_t)\right) + \boldsymbol{\epsilon}_t = \mathcal{A}_t(\mathbf{x},\mathbf{c},\mathbf{v}_t) + \boldsymbol{\epsilon}_t, \tag{1}$$

where $\mathbf{x} \in \mathbb{C}^D$ is the unknown image comprising of D voxels, $\mathbf{c} \in \mathbb{C}^{C\times D}$, the unknown spatially varying sensitivity maps of the C receiver coils, and $\boldsymbol{\epsilon} \in \mathbb{C}^{C\times D}$ is additive measurement noise. The element-wise multiplication operator $\odot\colon \mathbb{C}^{C\times D}\times\mathbb{C}^D \to \mathbb{C}^{C\times D}$ is given by $\mathbf{c}\odot\mathbf{x} = (\mathrm{diag}(\mathbf{c}_1)\mathbf{x},\ldots,\mathrm{diag}(\mathbf{c}_C)\mathbf{x})$. $\mathbf{F}\colon \mathbb{C}^{C\times D}\to\mathbb{C}^{C\times D}$ is the 3D discrete Fourier transform applied to each coil-weighted image in parallel. $\mathbf{M}_t$ is a block-diagonal matrix with C identical binary acquisition masks $\mathrm{diag}(\mathbf{m}_t)$ along the diagonal. Here, $\mathbf{m}_t \in \{0,1\}^D$ indicates if a particular k-space coefficient has been acquired at time point t. Finally, $\mathcal{W}(\cdot,\mathbf{v}_t)$ warps a static reference image according to motion state $\mathbf{v}_t \in \mathcal{V}$.

Joint reconstruction framework: We formulate the reconstruction problem within a Bayesian framework, where the posterior decomposes into a likelihood and priors for the image p_X, coils p_C, and motion states p_V

$$p(\mathbf{x},\mathbf{c},\mathbf{v}_{1:T}|\mathbf{z}_{1:T}) = p(\mathbf{z}_{1:T}|\mathbf{x},\mathbf{c},\mathbf{v}_{1:T})p_X(\mathbf{x})p_C(\mathbf{c})p_V(\mathbf{v}_{1:T}). \tag{2}$$

Computing the maximum a posteriori (MAP) estimate of (2) leads to the equivalent variational formulation

$$(\hat{\mathbf{x}},\hat{\mathbf{c}},\hat{\mathbf{v}}_{1:T}) \in \operatorname*{argmin}_{\mathbf{x},\mathbf{c},\mathbf{v}_{1:T}}\left\{\mathcal{D}\left(\mathcal{A}\left(\mathbf{x},\mathbf{c},\mathbf{v}_{1:T}\right),\mathbf{z}_{1:T}\right) + \mathcal{R}_{\mathbf{x}}(\mathbf{x}) + \mathcal{R}_{\mathbf{v}}(\mathbf{v}) + \mathcal{R}_{\mathbf{c}}(\mathbf{c})\right\}, \tag{3}$$

where $\mathcal{D}$ denotes the quadratic data fidelity term, The image prior ($\mathcal{R}_{\mathbf{x}}$) is defined implicitly by a diffusion model, while motion ($\mathcal{R}_{\mathbf{v}}$) and coil priors ($\mathcal{R}_{\mathbf{c}}$) are explicit regularizers.

Diffusion posterior sampling for motion compensation: We solve the joint reconstruction problem in (3) using a hybrid optimization strategy that combines *diffusion posterior sampling* (DPS) [6] for image estimation with *inertial proximal alternating linearized minimization* (iPALM) [8] for coil sensitivity and motion estimation. The resulting algorithm, termed *MotionDPS*, alternates between three coupled subproblems: (i) diffusion-guided image updates accounting for the posterior distribution conditioned on current coil sensitivities and motion states, (ii) proximal gradient updates for the coil sensitivity maps, and (iii) proximal gradient updates for the motion states. This alternating structure enables the integration of a strong learned image prior with physics-based data consistency and explicit motion modeling.

Algorithm 1 summarizes the proposed MotionDPS framework. Starting from an initial noise realization, coil sensitivity estimate, and zero motion, the algorithm iteratively progresses through a predefined diffusion noise schedule. At each diffusion step, the image is updated via a DPS step, followed by coil sensitivity and motion updates using proximal gradient steps derived from the data fidelity and corresponding priors. After completion of the diffusion process, the algorithm returns the final reconstructed image, coil sensitivity maps, and estimated motion trajectories.

3 Results

We evaluate the proposed method under motion simulation experiments using the FastMRI [9] and CC359 [10] datasets, and on real motion-corrupted data from PMoC3D dataset [11]. We report in terms of PSNR and SSIM, and compare against the following retrospective motion-correction methods: AltOpt [3], MotionTTT [5], and StackedUnet [4].

Motion-simulation experiments: We simulate inter-shot motion on fully sampled FastMRI and CC359 data using 52 acquisition shots. Performance is evaluated with and without undersampling under three motion levels – *mild*, *moderate*, and *severe* – with maximum translational and rotational amplitudes of $\pm(3,6,9)$ mm and $\pm(5°,10°,15°)$, respectively. For FastMRI, we apply 3D Cartesian undersampling along the phase-encoding dimensions with acceleration factors $R=4$ and $R=8$ (4% ACL), using a linear circular trajectory $\mathcal{T}$ starting from the k-space center. For CC359, we use an interleaved trajectory acquiring the central 3×3 region first and apply a 3D Cartesian Poisson-disc mask with $R=4.9$. Table 1 reports the quantitative average results. MotionDPS achieves the best performance in most cases and remains robust even under severe motion and strong undersampling. Fig. 1 shows reconstruction results for a scan with severe simulated motion.

Algorithm 1 MotionDPS

Require: k-space data $\mathbf{z}_{1:T}$, noise schedule $\{\sigma^i\}_{i=N}^0$, pre-trained denoiser D_θ.

$\triangleright$ *Initialization*

$\mathbf{x}^N = \sigma^N \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathrm{Id})$

$\mathbf{c}^N = \mathbf{F}^{-1}(\mathbf{z}) / \sqrt{\sum_{c=1}^{C} |\mathbf{F}^{-1}(\mathbf{z}_c)|^2}$

$\mathbf{v}_{1:T}^N = 0$

for $i = N-1, \ldots, 0$ **do**

$\beta = \frac{N-i}{N-i+3}$ $\triangleright$ Momentum

$\triangleright$ *Image subproblem*

$\mathbf{x}^i, \hat{\mathbf{x}}^0 = \mathrm{DPS}(\mathbf{x}^{i+1}, \mathbf{c}^{i+1}, \mathbf{v}_{1:T}^{i+1}, \mathbf{z}_{1:T}, \sigma^i, D_\theta)$ $\triangleright$ Diffusion Posterior Sampling

$\triangleright$ *Coil subproblem*

$\tilde{\mathbf{c}} = \mathbf{c}^{i+1} + \beta(\mathbf{c}^{i+1} - \mathbf{c}^{i+2})$

$\mathbf{c}^i = \mathrm{prox}_{\frac{1}{L_\mathbf{c}}\mathcal{R}_\mathbf{c}} \left(\tilde{\mathbf{c}} - \frac{1}{L_\mathbf{c}} \nabla_\mathbf{c} \mathcal{D} \left(\mathcal{A}(\hat{\mathbf{x}}^0, \tilde{\mathbf{c}}, \mathbf{v}_{1:T}^{i+1}), \mathbf{z}_{1:T} \right) \right)$

$\triangleright$ *Motion subproblem*

$\tilde{\mathbf{v}}_{1:T} = \mathbf{v}_{1:T}^{i+1} + \beta(\mathbf{v}_{1:T}^{i+1} - \mathbf{v}_{1:T}^{i+2})$

$\mathbf{v}_{1:T}^i = \mathrm{prox}_{P_\mathbf{v}^{-1}\mathcal{R}_\mathbf{v}}(\tilde{\mathbf{v}}_{1:T} - P_\mathbf{v}^{-1} \nabla_\mathbf{v} \mathcal{D}\mathcal{A}(\hat{\mathbf{x}}^0, \mathbf{c}^i, \tilde{\mathbf{v}}_{1:T}), \mathbf{z}_{1:T}))$

end for

return $(\mathbf{x}_0, \mathbf{c}, \mathbf{v})$

Table 1: Quantitative comparison of motion correction methods on FastMRI, CC359, and PMoC3D. Average PSNR (dB) and SSIM are reported across mild, moderate, and severe motion levels.

Dataset	Motion severity	AltOpt		MotionTTT		StackedUnet		MotionDPS	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
FastMRI	Mild	33.836	0.909	36.437	0.932	30.424	0.909	**38.260**	**0.962**
	Moderate	26.726	0.772	31.348	0.882	27.993	0.875	**36.539**	**0.954**
	Severe	24.982	0.721	26.804	0.803	26.575	0.838	**32.186**	**0.909**
CC359	Mild	**36.611**	**0.943**	32.019	0.892	31.017	0.906	35.720	0.933
	Moderate	32.011	0.866	30.940	0.869	28.008	0.849	**33.820**	**0.916**
	Severe	21.559	0.571	21.910	0.585	26.462	0.811	**27.575**	**0.856**
PMoC3D	Mild	31.595	0.956	32.199	0.958	31.603	0.952	**32.778**	**0.959**
	Moderate	30.718	0.955	31.141	0.958	29.110	0.939	**32.034**	**0.959**
	Severe	29.862	0.948	30.309	0.954	27.386	0.926	**31.277**	**0.955**

Real-motion experiments: We conduct real-motion experiments using the PMoC3D dataset, which comprises brain T1 volumes from eight healthy subjects. Acquisition was performed with an undersampling factor of $R = 4.94$ along the two phase-encoding directions, over 52 shots following a quasi-random k-space sampling trajectory $\mathcal{T}$, in which the central 3×3 region of k-space is acquired first. For quantitative evaluation, we adopt the same motion severity categorization (*mild*, *moderate*, and *severe*) as in Wang *et al.* [11]. Table 1 reports the quantitative reconstruction results, where MotionDPS consistently achieves higher PSNR and SSIM values. Fig. 2 presents qualitative reconstruction results for a scan in the *severe* motion category. MotionDPS more effectively suppresses motion artifacts while preserving anatomical structure, as reflected by lower-magnitude and more spatially uniform residual errors. Additionally, among the physics-based baselines, MotionDPS demonstrates a favorable accuracy–runtime trade-off, requiring on average 25min per volume, compared to 4h and 1h required by AltOpt and MotionTTT, respectively, on an NVIDIA A100 GPU.

4 Conclusion

We introduced MotionDPS, a novel method integrating diffusion-based image priors with explicit motion-aware data consistency to enable joint image and coil reconstruction as well as motion estimation for 3D MRI. Across both simulated and real-motion datasets, the method demonstrates superior robustness to motion and undersampling, consistently achieving higher PSNR and SSIM, generating cleaner reconstructions, while a the same time reducing runtime compared to competitive baselines. These results highlight the effectiveness of MotionDPS as a practical and scalable solution for 3D MRI reconstruction in the presence of patient motion.

Acknowledgments and Disclosure of Funding

AO, LS, and AE acknowledge funding by the German Research Foundation (DFG) Grants EXC2151-390873048 and EXC-2047/1-390685813. EK acknowledges funding from the DFG within the SPP2298 under project 543939932 and from the Austrian Science Fund (FWF) project number 10.55776/COE12.

References

[1] J. M. Slipsager, S. L. Glimberg, J. Søgaard, R. R. Paulsen, H. H. Johannesen, P. C. Martens, A. Seth, L. Marner, O. M. Henriksen, O. V. Olesen, and L. Højgaard, "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," *J. Magn. Reson. Imaging*, vol. 52, no. 3, pp. 731–738, 2020.

[2] A. Loktyushin, H. Nickisch, R. Pohmann, and B. Schölkopf, "Blind retrospective motion correction of MR images," *Magn. Reson. Med.*, vol. 70, no. 6, pp. 1608–1618, 2013.

[3] L. Cordero-Grande, R. P. A. G. Teixeira, E. J. Hughes, J. Hutter, A. N. Price, and J. V. Hajnal, "Sensitivity encoding for aligned multishot magnetic resonance reconstruction," *IEEE Trans. Comput. Imaging*, vol. 2, no. 3, pp. 266–280, 2016.

[4] M. A. Al-Masni, S. Lee, J. Yi, S. Kim, S.-M. Gho, Y. H. Choi, and D.-H. Kim, "Stacked U-nets with self-assisted priors towards robust correction of rigid motion artifact in brain MRI," *NeuroImage*, vol. 259, p. 119411, 2022.

[5] T. Klug, K. Wang, S. Ruschke, and R. Heckel, "MotionTTT: 2D test-time training motion estimation for 3D motion-corrected MRI," *Proc. NeurIPS*, 2025.

[6] H. Chung, J. Kim, M. T. McCann, M. L. Klasky, and J. C. Ye, "Diffusion posterior sampling for general noisy inverse problems," in *Proc. ICLR*, 2023.

[7] L. Chen, X. Tian, J. Wu, R. Feng, G. Lao, Y. Zhang, H. Liao, and H. Wei, "Joint coil sensitivity and motion correction in parallel MRI with a self-calibrating score-based diffusion model," *Med. Image Anal.*, vol. 102, p. 103502, 2025.

[8] T. Pock and S. Sabach, "Inertial proximal alternating linearized minimization (iPALM) for nonconvex and nonsmooth problems," *SIAM J. Imaging Sci.*, vol. 9, no. 4, pp. 1756–1787, 2016.

[9] F. Knoll, J. Zbontar, A. Sriram, M. J. Muckley, M. Bruno, A. Defazio, M. Parente, K. J. Geras, J. Katsnelson, H. Chandarana, Z. Zhang, M. Drozdzal, A. Romero, M. Rabbat, P. Vincent, J. Pinkerton, D. Wang, N. Yakubova, E. Owens, C. L. Zitnick, M. P. Recht, D. K. Sodickson, and Y. W. Lui, "fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning," *Radiol.: Artif. Intell.*, vol. 2, no. 1, p. e190007, 2020.

[10] R. Souza, O. Lucena, J. Garrafa, D. Gobbi, M. Saluzzi, S. Appenzeller, L. Rittner, R. Frayne, and R. Lotufo, "An open, multi-vendor, multi-field-strength brain MR dataset and analysis of publicly available skull stripping methods agreement," *NeuroImage*, vol. 170, pp. 482–494, 2018.

[11] K. Wang, T. Klug, S. Ruschke, J. S. Kirschke, and R. Heckel, "Reliable evaluation of MRI motion correction: Dataset and insights," 2025, arXiv preprint arXiv:2506.05975.