Physics-driven Cold Diffusion with Severity-aware Sampling and Adaptive Degradation Estimation for MRI Motion Artifact Removal

Chuanpu Li[1,2,3], Jinlong Zhang[1,2], Yuan Yang[4], Yueming Jin[3(✉)], and Wei Yang[1,2(✉)]

[1] School of Biomedical Engineering, Southern Medical University, Guangzhou, China
weiyanggm@gmail.com

[2] Guangdong Provincial Key Laboratory of Medical Image Processing, Southern Medical University, Guangzhou, China

[3] Department of Biomedical Engineering and Department of Electrical and Computer Engineering, National University of Singapore (NUS)
ymjin@nus.edu.sg

[4] Department of Radiation Oncology, Affiliated Hangzhou First People' s Hospital, School of Medicine, Westlake University, Hangzhou, China

Abstract. Motion artifacts remain a major challenge in magnetic resonance imaging (MRI) because they degrade image quality and compromise diagnostic reliability. Recently, the strong generative capacity of diffusion models has driven increasing interest in applying these methods to MRI artifact removal. However, conventional noise-driven diffusion relies on injecting Gaussian noise, which does not match the acquisition-induced corruption produced by motion. It also introduces sampling uncertainty and may hallucinate structures. To address these issues, we propose a noise-free cold diffusion for MRI motion artifact removal trained solely on motion-free image. In this framework, the diffusion process is formulated as a physics-driven degradation operator that incorporates an acquisition-consistent masking strategy and a dirichlet-weighted motion state assignment. This design reflects the mechanism through which motion artifacts arise during acquisition and naturally aligns diffusion timesteps with artifact severity. To support inference across diverse motion patterns and artifact levels, we further introduce severity-aware sampling and adaptive degradation estimation to guide the restoration trajectory and improve generalization. Experiments on two public datasets demonstrate that the proposed method achieves superior performance in MRI motion artifact removal.

Keywords: Cold diffusion · MRI motion artifact removal · Physics-driven training.

1 Introduction

Magnetic Resonance Imaging (MRI) is a widely used modality in clinical diagnosis because it provides rich soft-tissue contrast without ionizing radiation. However, the inherently long acquisition time makes MRI highly vulnerable to patient motion. Such motion introduces phase inconsistencies across k-space lines, leading to structured artifacts [6, 9]. These artifacts not only degrade image quality but also obscure pathological details, potentially leading to misdiagnosis or the need for repeated scans [22]. Therefore, effective motion artifact removal is of great clinical value.

Many learning-based methods have been proposed for retrospective MRI motion correction, mainly following supervised or unsupervised paradigms. Supervised methods learn a direct mapping from motion-corrupted images to motion-free targets [7, 10, 13, 23]. However, paired motion-free and motion-corrupted data are rarely available in clinical practice. Unsupervised paradigm relaxes the requirement of paired data and interprets motion artifact removal as CycleGAN-based image-to-image translation [1,3,11,12,15,25,26]. However, unpaired translation implicitly assumes a consistent imaging style between the corrupted and clean domains, whereas actual MRI appearances vary substantially across sequences and acquisition parameters [19]. Lacking an explicit motion-artifact rejection mechanism [17, 18], these methods are sensitive to scanner and protocol shifts, and may erroneously entangle style difference with artifact removal.

To mitigate the data dependency and domain shift issues, an alternative paradigm focuses on learning generative priors from motion-free MRI images [4, 17, 18, 28] or by physics-informed reconstruction that excludes motion-corrupted k-space lines in a self-supervised manner [6]. Within this context, diffusion models have attracted considerable attention due to their strong generative capabilities [4, 5, 17, 20, 28]. For example, Oh et al. proposed an annealed score-based diffusion model that is trained on motion-free data and removes motion artifacts via repeated forward-reverse diffusion with low-frequency consistency guidance during inference [17]. Xu et al. used a pretrained diffusion prior and constructed alternating complementary masks in the reverse process to destroy the artifact structure [28]. Chen et al. integrated a score-based diffusion prior into parallel MRI to jointly correct motion and coil sensitivities [4]. Despite promising results, these approaches still typically define diffusion via a Gaussian noising process, which differs from how motion corrupts MRI measurements through k-space inconsistencies [8]. Moreover, the stochastic noise-based sampling introduces unpredictable variability and risks hallucinating anatomical structures [14], which conflicts with the demand of medical image for a precise and deterministic process [21].

Recently, cold diffusion demonstrated that the generative forward process can be constructed using arbitrary deterministic degradations, moving beyond the sole reliance on stochastic Gaussian noise addition [2]. Related diffusion-bridge formulations have been developed for accelerated MRI reconstruction [16], whereas motion artifact removal remains unexplored. Inspired by this, we propose a noise-free, physics-driven cold diffusion framework trained solely on motion-free images for MRI motion artifact removal. Instead of adding Gaussian noise, we formulate the diffusion process as a physics-inspired degradation operator. The operator integrates an acquisition-consistent masking strategy and a dirichlet-weighted motion state assignment, enabling realistic corruption that

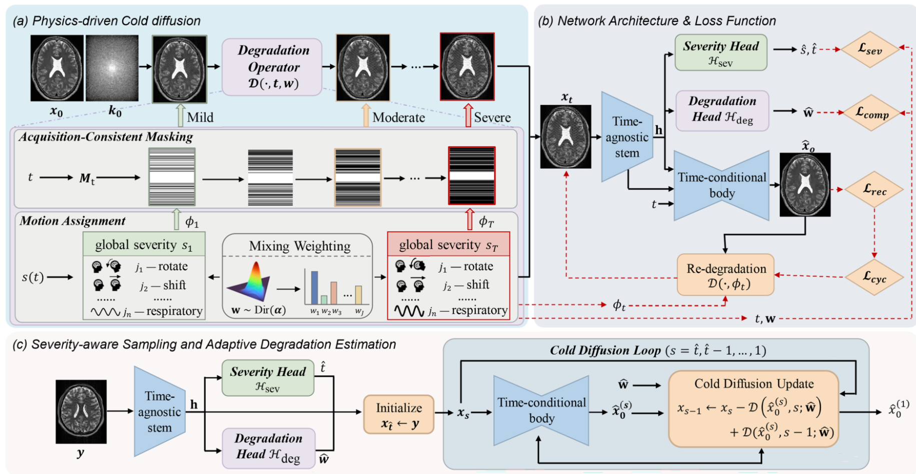

Fig. 1. Overview of the proposed method.

mirrors motion-induced inconsistencies while keeping the diffusion timesteps naturally aligned with artifact severity. Furthermore, to support robust inference across diverse motion patterns and artifact levels, we introduce a severity-aware sampling strategy and adaptive degradation estimation. By automatically guiding the deterministic restoration trajectory, our method ensures that no uncertainty factors arise during the deterministic reverse process, avoiding random sampling issues and enabling the network to better comprehend the physical artifact distribution. Extensive experiments on two public datasets demonstrate that our proposed method achieves superior performance.

2 Method

2.1 Cold Diffusion

As shown in Fig. 1(a), cold diffusion learns to invert a deterministic degradation process rather than additive Gaussian noise. Given a motion-free target $\mathbf{x}_0 \sim p_{\text{data}}$, a degradation family $\{\mathcal{D}_t\}_{t=0}^{T}$ defines a forward trajectory:

$$\mathbf{x}_t = \mathcal{D}(\mathbf{x}_0, t), \quad t \in \{1, \dots, T\}, \quad \text{s.t. } \mathcal{D}(\mathbf{x}_0, 0) = \mathbf{x}_0. \tag{1}$$

The goal is to train a time-conditioned restoration network $\mathcal{R}_\theta$ to invert this process and recover the clean image from any intermediate state:

$$\hat{\mathbf{x}}_0 = \mathcal{R}_\theta(\mathbf{x}_t, t). \tag{2}$$

In MRI images, motion artifacts are induced by k-space inconsistency whose strength increases with motion frequency and amplitude [22, 29]. We therefore construct $\mathcal{D}_t$ as a physics-inspired degradation operator that progressively increases k-space inconsistency and state variation with t. As a result, the cold diffusion timestep t becomes naturally aligned with artifact severity, enabling severity-aware initialization and stable iterative inversion over a wide range of artifact severities.

2.2 Physics-inspired Degradation Operator

Let $\mathbf{k}_0 = \mathcal{F}(\mathbf{x}_0)$ be the k-space image of $\mathbf{x}_0$, where $\mathcal{F}$ denotes the Fourier Transform. We assume the degradation process is corrupted by J motion states (e.g., rotation, translation, respiratory, etc). The corrupted k-space $\mathbf{k}_t$ at degradation severity t is modeled as a mix of the coherent base signal and motion-perturbed signals:

$$\mathbf{k}_t = (\mathbf{1} - \mathbf{M}_t) \odot \mathbf{k}_0 + \sum_{j=1}^{J} \mathbf{M}_t^{(j)} \odot \mathcal{F}\left(\mathcal{T}(\mathbf{x}_0; \phi_j)\right), \tag{3}$$

where $\odot$ denotes the element-wise Hadamard product. The term $\mathcal{T}(\cdot;\phi_j)$ represents a motion transformation parameterized by ϕ_j, corresponding to the j-th motion state. $\mathbf{M}_t = \sum_{j=1}^{J} \mathbf{M}_t^{(j)}$ denotes the union of all corrupted k-space trajectories at severity t, serving as a frequency-domain selection mask. The resulting degraded image $\mathbf{x}_t$ is obtained via the inverse Fourier transform:

$$\mathbf{x}_t = \mathcal{D}_t(\mathbf{x}_0) \triangleq \left|\mathrm{F}^{-1}(\mathbf{k}_t)\right|. \tag{4}$$

Acquisition-Consistent Masking Strategy. In standard Cartesian MRI, the central k-space lines which determine the global anatomy are often sampled first. Patients typically remain stationary during this initial phase, preserving the integrity of low-frequency data. However, as the scanning continues, the likelihood of involuntary motion increases. Consequently, motion-induced errors mainly affect the high-frequency lines acquired later in the process [17, 28]. Guided by this physical observation, we specify the distribution of $\mathbf{M}_t$ using an acquisition-consistent strategy. Specifically, let u denote the phase index across the total H trajectories, and u_c be the center frequency index. We first define the protected central region Ω_{clean} with a bandwidth controlled by the center preservation fraction $c\,(t) = c_{min} + \ (c_{max} - c_{min})(1 \ - \frac{t}{T})$:

$$\Omega_{\text{clean}}^{(t)} = \left\{ u \;\middle|\; |u - u_c| \leq \frac{c(t)\,H}{2} \right\}. \tag{5}$$

Trajectories within $\Omega_{\text{clean}}^{(t)}$ are strictly preserved. For the remaining candidate trajectories in the peripheral region $\Omega_{\text{cand}}^{(t)}$, we define a severity-dependent sampling weight $w_t(u)$ for each $u \in \Omega_{\text{cand}}^{(t)}$:

$$w_t(u) = \left[\sigma \left(\frac{|u - u_c| - \frac{c(t)H}{2}}{\tau} \right) \right]^p, \tag{6}$$

where $\sigma(\cdot)$ is the Sigmoid function. τ controls the softness of the frequency transition, and $p \geq 1$ is a power factor that strengthens the bias towards higher frequencies. Finally, $\mathbf{M}_t$ is generated via weighted sampling without replacement. Specifically, the subset of corrupted phase index $\{u \mid \mathbf{m}_t(u) = 1\}$ is drawn from $\Omega_{\mathrm{cand}}^{(t)}$:

$$N_{pick} = \left\lfloor \rho(t) \cdot \left| \Omega_{\mathrm{cand}}^{(t)} \right| \right\rfloor, \tag{7}$$

where $\rho\left(t\right) = \rho_{max} \cdot \frac{t}{T}$ is the corruption ratio, scaling the artifact intensity linearly with the severity timestep t. To maintain a fully deterministic restoration trajectory at inference, we replace the stochastic mask sampling in Eq. (7) with its continuous relaxation.

Dirichlet-Weighted Motion State Assignment. At a fixed severity level t, the relative weights of different motion patterns can vary depending on patients and scanned anatomy. For instance, brain imaging typically exhibits higher probabilities of rotation and translation, whereas abdominal MRI is predominantly affected by respiratory motion artifacts [22]. To model this variability without altering the global severity, we propose a dirichlet-weighted mechanism that generalizes to J motion degrees of freedom. Let $\Phi = \{\phi^{(1)}, \ldots, \phi^{(J)}\}$ denote the set of possible motion parameters (e.g., rotation angle, translation shift, etc). We sample a mixing weight vector $\mathbf{w} \in \Delta^{J-1}$ from a dirichlet distribution:

$$\mathbf{w} \sim \mathrm{Dir}(\boldsymbol{\alpha}), \quad \text{s.t.} \sum_{j=1}^{J} w_j = 1, \tag{8}$$

where $\boldsymbol{\alpha} \in \mathbb{R}^J$ is the concentration hyperparameter governing the sparsity of motion types. The parameter value $\phi_t^{(j)}$ for the j-th degrees of freedom at severity t is instantiated as:

$$\phi_t^{(j)} = \mathcal{S}_j \left(s(t) \cdot \lambda_{\max} \cdot w_j \right), \tag{9}$$

where $s(t) = \frac{t}{T}$ is the severity scalar, and $\lambda_{\max}$ defines the physical upper bound for the j-th motion type. The operator $\mathcal{S}_j(\cdot)$ represents a stochastic direction sampling appropriate for those specific motion states.

2.3 Severity-aware Sampling and Adaptive Degradation Estimation

At test time, motion-corrupted MR images differ in both artifact severity and motion composition. Accurately estimating severity is crucial as it determines the starting point for diffusion inversion, while identifying the composition is essential for instantiating the appropriate degradation operator. Therefore, we introduce two lightweight auxiliary heads to predict the severity $\hat{s}$ and the composition $\hat{\mathbf{w}}$ from the input. Specifically, as shown in Fig. 1(b), we adapt the U-Net based architecture with a timestep-agnostic stem to extract features $\mathbf{h}$ without time embeddings. The severity head $\mathcal{H}_{\mathrm{sev}}$ maps the globally pooled features $\bar{\mathbf{h}} = \mathrm{GAP}(\mathbf{h})$ to a normalized severity index $\hat{s}$ and a timestep $\hat{t}$:

$$\hat{s} = \sigma\left(\mathcal{H}_{\mathrm{sev}}(\bar{\mathbf{h}})\right), \quad \hat{t} = \lfloor \hat{s} \cdot T \rceil. \tag{10}$$

Simultaneously, the degradation estimation head $\mathcal{H}_{\text{deg}}$ identifies the composition of the motion artifacts by predicting a directional composition vector $\hat{\mathbf{w}}$:

$$z = \tanh\big(\mathcal{H}_{\text{deg}}(\bar{\mathbf{h}})\big), \quad \hat{\mathbf{w}} = \frac{z}{\|z\|_1}. \tag{11}$$

With $\hat{s}$ and $\hat{\mathbf{w}}$, we can realize severity-aware sampling and adaptive degradation estimation. Given a corrupted input $\mathbf{y}$, we first obtain $(\hat{t}, \hat{w})$ to instantiate the degradation operator $\mathcal{D}(\cdot,\, t;\, \hat{w})$ in Eq. (3)-(7), and set $\mathbf{x}_{\hat{t}} \leftarrow \mathbf{y}$. Following the cold diffusion [2], we iteratively solve for $s = \hat{t}, \hat{t}-1, \ldots, 1$. At each step, the network predicts the clean image $\hat{\mathbf{x}}_0^{(s)}$, followed by the deterministic update:

$$\hat{\mathbf{x}}_0^{(s)} \leftarrow R_\theta\left(\mathbf{x}_s, s\right), \quad \mathbf{x}_{s-1} \leftarrow \mathbf{x}_s - \mathcal{D}\left(\hat{\mathbf{x}}_0^{(s)}, s; \hat{\mathbf{w}}\right) + \mathcal{D}\left(\hat{\mathbf{x}}_0^{(s)}, s-1; \hat{\mathbf{w}}\right), \tag{12}$$

where $\mathcal{D}\left(\cdot, s; \hat{\mathbf{w}}\right)$ denotes the adaptive degradation operator instantiated dynamically by $\hat{\mathbf{w}}$.

2.4 Loss Function

Following the cold diffusion paradigm, the primary reconstruction loss L_{rec} enforces pixel-level fidelity between the restored image $\hat{\mathbf{x}}_0$ and the ground truth $\mathbf{x}_0$:

$$\mathcal{L}_{rec} = \|\hat{\mathbf{x}}_0 - \mathbf{x}_0\|_1. \tag{13}$$

The severity regression is optimized via a smooth ℓ_1 loss, while the composition weights are aligned using the squared ℓ_2 norm:

$$\mathcal{L}_{sev} = \ell_{\text{smooth}}(\hat{t}, t), \quad \mathcal{L}_{comp} = \sum_{j=1}^{J} \left(\hat{w}_j - \mathcal{S}_j \cdot w_j\right)^2. \tag{14}$$

To guarantee physical validity, we re-degrade the restoration output with the ground-truth parameters ϕ_t:

$$\mathcal{L}_{cyc} = \|\mathcal{D}(\hat{\mathbf{x}}_0; \phi_t) - \mathbf{x}_t\|_1. \tag{15}$$

The total objective is the weighted sum:

$$\mathcal{L}_{total} = \mathcal{L}_{rec} + \lambda_{sev}\mathcal{L}_{sev} + \lambda_{comp}\mathcal{L}_{comp} + \lambda_{cyc}\mathcal{L}_{cyc}. \tag{16}$$

3 Experiments

3.1 Experimental Setup.

Datasets. We evaluated our method on two public MRI datasets. The first dataset is the KMAR-50K dataset [27], which comprises multi-view, multi-sequence knee MRIs featuring real-world clinical motion artifacts and paired ground truths. The second dataset is the Human Connectome Project (HCP) data [24], which contains T2-weighted images of the human brain. For evaluation and validation, motion artifacts were simulated in this dataset by introducing phase errors into the k-space data, following established methods in [26, 28].

Table 1. Quantitative comparison of different methods on the KMAR-50K dataset and HCP brain dataset.

Method	KMAR-50K Dataset		HCP Brain Dataset	
	PSNR (dB) ↑	SSIM (%) ↑	PSNR (dB) ↑	SSIM (%) ↑
MARC [23]	28.34 ± 1.35	80.22 ± 3.12	33.48 ± 4.00	94.81 ± 2.44
CycleMedGAN [1]	26.67 ± 1.23	77.45 ± 3.49	29.61 ± 3.57	91.25 ± 2.59
DUNCAN [15]	27.21 ± 1.40	77.92 ± 3.45	30.80 ± 4.88	89.44 ± 2.22
UDDN [26]	27.29 ± 1.41	78.09 ± 3.41	31.44 ± 4.53	90.16 ± 2.92
Bootstrap [18]	27.05 ± 1.36	77.99 ± 3.72	31.72 ± 4.60	90.91 ± 2.07
ASBD [17]	28.32 ± 1.36	81.03 ± 3.12	32.12 ± 4.01	92.94 ± 2.42
PFAD [28]	28.71 ± 1.35	81.82 ± 3.06	32.28 ± 4.13	93.00 ± 2.43
Ours	**29.12 ± 1.36**	**82.32 ± 3.00**	**33.63 ± 3.96**	**95.43 ± 2.42**

Implementation Details. Our framework was implemented in PyTorch and trained on an NVIDIA A5000 GPU with a batch size of 4. To capture complex anatomical movements, the motion-state manifold encompasses diverse spatial transformations, where the motion transformation $\mathcal{T}$ is instantiated following [22], with physical upper bounds $\lambda_{\max}$ for shifts and rotations set to 10.0 and 8.0, as artifacts beyond these limits irrecoverably destroy anatomical structures [29]. For the degradation operator, we set the minimum center preservation $c_{\min} = 0.15$ to protect low-frequency k-space data, ensuring the maintenance of fundamental image contrast [22]. Furthermore, we configured the maximum high-frequency corruption $\rho_{\max} = 0.9$. The network was optimized for 400 epochs using Adam with a learning rate of 2×10^{-4}. The loss weights were set to $\lambda_{sev} = 0.005$, $\lambda_{comp} = 0.05$, $\lambda_{cycle} = 0.5$, respectively.

3.2 Results

Comparisons with State-of-the-Art. We compared our method with several MRI artifact removal methods including supervised methods such as MARC [23], unsupervised methods such as CycleMedGAN [1], DUNCAN [15], UDDN [26], and motion-free training methods such as GAN-based Bootstrap [18], diffusion-based ASBD [17] and PFAD [28]. As illustrated in Table 1, our proposed method demonstrates superior performance, exhibiting a PSNR of 29.12 ± 1.36 dB and a SSIM of 89.22 ± 2.83 % for real knee MRI motion artifact removal, and a PSNR of 33.63 ± 3.96 dB and a SSIM of 95.43 ± 2.42 % for simulated brain MRI artifact removal. Visual comparisons in Fig. 2 demonstrates that while noise-based diffusion models like ASBD [17] and PFAD [28] effectively suppress artifacts, our method better preserves fine anatomical details (red boxes) to achieve superior structural fidelity.

Ablation Study. Table 2 shows the ablation study in the validation set. The baseline (M1) uses standard cold diffusion with masking strategy. Adding the dirichlet-weighted motion state assignment (M2) brings the about 1 dB gain. Next, the severity-aware sampling (M3) further improves the model's stabil-

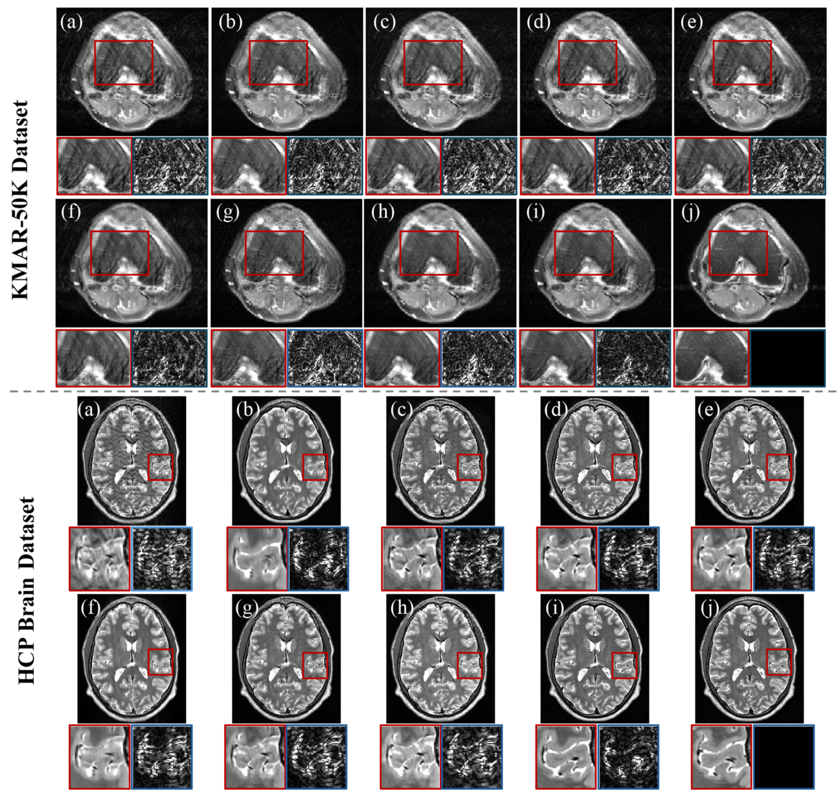

Fig. 2. Visual comparison of motion artifact removal results. (a) Artifact image, (b) MARC, (c) CycleMedGAN, (d) UDDN, (e) DUNCAN, (f) Bootstrap, (g) ASBD, (h) PFAD, (i) ours , (j) Ground truth.

ity across different artifact levels. Finally, using adaptive degradation estimation (M4) provides the correct motion composition for input, achieving the best structural details. This improvement proves that combining physical priors with severity-aware sampling and adaptive degradation estimation is effective.

4 Conclusion

This paper proposes a noise-free, physics-driven cold diffusion framework for MRI motion artifact removal. By formulating the forward process as a physical-inspired degradation operator that mirrors actual k-space inconsistencies, our method avoids the inherent stochasticity of Gaussian noise. To handle diverse artifact distributions during inference, we introduce a severity-aware sampling and adaptive degradation estimation mechanism, which automatically initializes and guides the deterministic restoration trajectory, and enables the network to effectively recover artifact-free anatomical structures. Extensive experiments on

Table 2. Ablation study of critical components of the proposed method in the validation set.

Model	KMAR-50K Dataset		HCP Brain Dataset	
	PSNR (dB) ↑	SSIM (%) ↑	PSNR (dB) ↑	SSIM (%) ↑
M1: Baseline	27.33 ± 1.27	80.04 ± 3.01	32.02 ± 4.00	92.88 ± 2.33
M2: M1+weighting	28.58 ± 1.36	81.24 ± 3.03	32.93 ± 3.98	94.52 ± 2.34
M3: M2+severity	28.89 ± 1.36	81.56 ± 3.01	33.12 ± 3.99	95.12 ± 2.36
M4: M3+estimation	29.35 ± 1.35	81.96 ± 3.03	33.69 ± 4.01	95.45 ± 2.37

two public datasets demonstrate that our approach outperforms existing state-of-the-art methods.

Acknowledgments. This work was partially supported by grants from the National Natural Science Foundation of China (No. 82472050), Ministry of Education Tier 1 grant, Singapore (24-1250-P0001),and Outstanding Young Talents of Guangdong Special Support Program (Health Commission of Guangdong Province) (0820250210).

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

References

1. Armanious, K., Tanwar, A., Abdulatif, S., Küstner, T., Gatidis, S., Yang, B.: Unsupervised adversarial correction of rigid mr motion artifacts. In: 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI). pp. 1494–1498. IEEE (2020)
2. Bansal, A., Borgnia, E., Chu, H.M., Li, J., Kazemi, H., Huang, F., Goldblum, M., Geiping, J., Goldstein, T.: Cold diffusion: Inverting arbitrary image transforms without noise. Advances in Neural Information Processing Systems **36**, 41259–41282 (2023)
3. Chen, G., Xie, H., Rao, X., Liu, X., Otikovs, M., Frydman, L., Sun, P., Zhang, Z., Pan, F., Yang, L., Zhou, X., Liu, M., Bao, Q., Liu, C.: Mri motion correction through disentangled cyclegan based on multi-mask k-space subsampling. IEEE Transactions on Medical Imaging **44**(4), 1907–1921 (2025)
4. Chen, L., Tian, X., Wu, J., Feng, R., Lao, G., Zhang, Y., Liao, H., Wei, H.: Joint coil sensitivity and motion correction in parallel mri with a self-calibrating score-based diffusion model. Medical Image Analysis **102**, 103502 (2025)
5. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Advances in neural information processing systems **34**, 8780–8794 (2021)
6. Eichhorn, H., Spieker, V., Hammernik, K., Saks, E., Weiss, K., Preibisch, C., Schnabel, J.A.: Physics-informed deep learning for motion-corrected reconstruction of quantitative brain mri. In: Linguraru, M.G., Dou, Q., Feragen, A., Giannarou, S., Glocker, B., Lekadir, K., Schnabel, J.A. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. pp. 562–571. Springer Nature Switzerland, Cham (2024)
7. 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)
8. Huang, J., Aviles-Rivero, A.I., Schönlieb, C.B., Yang, G.: Cdiffmr: Can we replace the gaussian noise with k-space undersampling for fast mri? In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 3–12. Springer (2023)
9. Jin, K.H., Um, J.Y., Lee, D., Lee, J., Park, S.H., Ye, J.C.: Mri artifact correction using sparse+ low-rank decomposition of annihilating filter-based hankel matrix. Magnetic resonance in medicine **78**(1), 327–340 (2017)
10. 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)
11. Jung, S., Choi, Y., Al-Masni, M.A., Jung, M., Kim, D.H.: Deformation-aware segmentation network robust to motion artifacts for brain tissue segmentation using disentanglement learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 213–222. Springer (2024)
12. Kim, S., Al-masni, M.A., Lee, S., Jung, S., Jung, K.J., Cui, C., Gho, S.M., Choi, Y.H., Kim, D.H.: Unsupervised learning for motion correction and assessment in brain magnetic resonance imaging using severity-based regularized cycle consistency. Engineering Applications of Artificial Intelligence **142**, 109978 (2025)
13. 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)
14. Li, F., Zhou, Z., Fang, Y., Cai, J., Wang, Q.: Moco-diff: adaptive conditional prior on diffusion network for mri motion correction. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 411–421. Springer (2024)
15. Liu, S., Thung, K.H., Qu, L., Lin, W., Shen, D., Yap, P.T.: Learning mri artefact removal with unpaired data. Nature Machine Intelligence **3**(1), 60–67 (2021)
16. Mirza, M.U., Dalmaz, O., Bedel, H.A., Elmas, G., Korkmaz, Y., Gungor, A., Dar, S.U.H., Oguz, K.K., Çukur, T.: Learning fourier-constrained diffusion bridges for mri reconstruction. IEEE Transactions on Medical Imaging **45**(6), 3229–3245 (2026)
17. Oh, G., Jung, S., Lee, J.E., Ye, J.C.: Annealed score-based diffusion model for mr motion artifact reduction. IEEE Transactions on Computational Imaging **10**, 43–53 (2023)
18. 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)
19. Scholz, D., Erdur, A.C., Holland, R., Ehm, V., Peeken, J.C., Wiestler, B., Rueckert, D.: Contrastive anatomy-contrast disentanglement: A domain-general mri harmonization method. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 100–110. Springer (2025)
20. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
21. Song, T., Nie, F., Guo, Y., Xu, F., Zhang, S.: Filterdiff: Noise-free frequency-domain diffusion models for accelerated mri reconstruction. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 205–215. Springer (2025)

22. 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)
23. Tamada, D., Kromrey, M.L., Ichikawa, S., Onishi, H., Motosugi, U.: Motion artifact reduction using a convolutional neural network for dynamic contrast enhanced mr imaging of the liver. Magnetic Resonance in Medical Sciences **19**(1), 64–76 (2020)
24. Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T.E., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S.W., et al.: The human connectome project: a data acquisition perspective. Neuroimage **62**(4), 2222–2231 (2012)
25. Wang, J., Yang, H., Yang, Y., Sun, J.: Mri motion artifact correction via frequency-assisted artifact disentanglement and confidence-guided knowledge distillation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 392–401. Springer (2025)
26. Wu, B., Li, C., Zhang, J., Lai, H., Feng, Q., Huang, M.: Unsupervised dual-domain disentangled network for removal of rigid motion artifacts in mri. Computers in Biology and Medicine **165**, 107373 (2023)
27. Xi, Y., Wang, F., Shi, F., Zhou, Q., Li, R., Li, Y., Wang, Y.: A multi-view open-access dataset of paired knee mri for motion artifact removal. Scientific Data **12**(1), 1173 (2025)
28. Xu, J., Zhou, D., Hu, L., Guo, J., Yang, F., Liu, Z., Wang, N., Gao, X.: Motion artifact removal in pixel-frequency domain via alternate masks and diffusion model. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 8878–8886 (2025)
29. 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)