

Privacy Preserving Structurally Constrained GAN for Artifact Reduction in Liver MRIs

Raghul Sachin R¹, Dr. T. Anusha²

^{1,2}Department of CSE, SRM Institute of Science & Technology Vadapalani campus, Chennai, India

Abstract

Magnetic Resonance Imaging (MRI) is one of the major tools used for diagnosing the liver; nevertheless, liver images are prone to motion and breathing artifacts, thus resulting in low-quality images. Although Generative Adversarial Networks (GANs) have been used in restoring images, traditional GAN models suffer from generating hallucinations and changes in essential anatomical structures. In addition, using medical imaging data poses great privacy risks. This work presents the Privacy-Preserving Structurally Constrained GAN for artifact removal in 3D Liver MRIs. First, 2D slices are extracted from 3D NIfTI images. Then, an extremely structured generator is trained via Structural Similarity (SSIM) and perceptual loss functions in order to preserve liver parenchyma geometry and vascular anatomy. The model introduces the differentially private discrimination gradient as a privacy-preserving strategy to safeguard patient identities. The model can successfully eliminate artifacts while maintaining the image's structure, according to experiments conducted on paired and unpaired artifact-filled NIfTI images. The framework effectively prevents membership inference attacks without significantly impairing restorative performance by mathematically bounding the impact of any single patient scan during the optimization phase. The ensuing high-fidelity restorations reduce the possibility of incorrect diagnosis or the need for expensive patient rescans by giving radiologists clear, diagnostically reliable cross-sections. In the end, this approach opens the door for safe implementation in actual clinical settings by bridging the crucial gap between sophisticated deep learning capabilities and stringent healthcare data regulations.

Index Terms: Generative Adversarial Networks (GAN), Liver MRI, Artifact Reduction, Privacy-Preserving ML, NIfTI, Image Restoration, Structural Constraints.

INTRODUCTION

When it comes to identifying lesions, assessing hepatic fibrosis, and planning complex surgical interventions, Magnetic Resonance Imaging (MRI) is widely regarded as the gold standard diagnostic modality due to its unparalleled soft-tissue contrast. However, the diagnostic efficacy of abdominal MRIs is frequently compromised by the inherently long acquisition times required to capture high-resolution volumetric data. During these extended scanning periods, the liver is continuously subjected to involuntary physiological motions, including respiratory cycles, cardiac pulsations, and gastrointestinal peristalsis. These unavoidable movements introduce severe phase-encoding errors into the raw k-space data, which manifest visually in the reconstructed images as ghosting, blurring, and heavy directional striping. For clinicians and

radiologists, these artifacts severely obscure critical hepatic contours, portal venous networks, and early-stage tumors. This degradation directly increases the risk of misdiagnosis, necessitates costly scan repetitions, and ultimately delays critical patient care.

To mitigate these imaging degradations, the medical imaging community has increasingly turned to advanced deep learning techniques. Generative Adversarial Networks (GANs), in particular, have demonstrated remarkable capabilities in translating artifact-corrupted inputs into restored, high-fidelity images by learning the complex mapping between noisy and clean data distributions. However, the standard GAN architecture is fundamentally optimized to produce images that are merely "visually plausible" to a discriminator network. In a rigorous clinical context, visual realism is dangerously insufficient. Unrestricted generators, lacking explicit anatomical guidance, are highly prone to structural hallucination—they frequently fabricate erroneous anatomical textures or, conversely, aggressively smooth over actual microscopic anomalies, such as small hepatocellular carcinomas or subtle fibrotic scarring. Consequently, it is an absolute necessity to embed strict, biologically accurate structural constraints directly into the generative optimization process to ensure that the restored image remains a perfect, unaltered anatomical match to the patient.

Concurrently, the transition of deep learning models from experimental research to widespread clinical deployment is heavily bottlenecked by stringent data privacy regulations, such as the Health Insurance Portability and Accountability Act (HIPAA) and the General Data Protection Regulation (GDPR). Training robust, generalizable image restoration models inherently requires access to massive, centralized repositories of highly sensitive patient data. Recent cybersecurity analyses have proven that standard neural network architectures tend to memorize the unique, high-frequency features of their training data. This fundamental vulnerability exposes the models to sophisticated adversarial exploits, including model inversion and membership inference attacks. In these scenarios, malicious actors can reverse-engineer the trained weight vectors of a deployed model to extract highly confidential, identifiable patient anatomy. This creates a critical bottleneck: the medical field urgently needs the restorative power of deep learning, but cannot afford the catastrophic legal and ethical risks of a patient data breach.

To resolve this critical intersection of structural fidelity and data security, this study presents a novel, robust framework for abdominal imaging: the Privacy-Preserving Structurally Constrained GAN for liver MRI artifact reduction. This architecture fundamentally redesigns the traditional generative pipeline. First, it rigorously anchors the generated output to the true biological geometry of the organ by hardcoding structural similarity (SSIM) and pixel-wise fidelity metrics directly into the network's loss function, effectively eliminating the risk of synthetic hallucinations. Simultaneously, cryptographic security is integrated intrinsically into the network's learning phase via Differential Privacy. By applying strict gradient clipping and calibrated noise injection during the optimization of the discriminator, the framework mathematically guarantees that the artifact-removal model learns the generalized physics of image restoration without ever memorizing the identifiable anatomical markers of any individual patient in the training cohort.

LITERATURE SURVEY

Deep learning and medical imaging have developed quickly, especially in improving diagnostic quality

through data-driven image reconstruction [1]. Current advancements in Magnetic Resonance Imaging (MRI) concentrate on enhancing image fidelity and automating artifact reduction. The broader application of machine learning in this field has been thoroughly examined, highlighting its significance in a number of tasks such as motion artifact mitigation, denoising, and bias field correction [2], [3], [4]. Retrospective motion correction, where sophisticated algorithms process corrupted k-space or image-domain data to restore the clarity of anatomical structures, is one of the most important applications in this field [5].

Using conventional iterative techniques, early correction systems demonstrated the idea, but they frequently failed to capture the intricate, non-linear physics of patient motion. Deep learning techniques were quickly adopted due to the need for high-fidelity, structurally accurate restorations. Convolutional neural networks have been explored in subsequent research to improve the accuracy and practicality of these reconstructions [6]. Modern models improve diagnostic work-flows for complex abdominal imaging by measuring artifact severity and automating detection, which increases radiologists' confidence in the corrected scans [7], [8].

The possibility of structural hallucination has long been a concern in medical image restoration, despite significant improvements in predictive performance. Because of their capacity to produce extremely realistic textures, Generative Adversarial Networks (GANs) have emerged as a leading architecture for medical image synthesis and restoration [9]. When paired artifact-free data is not available, researchers are heavily utilizing variants such as CycleGAN to perform retrospective rigid motion correction [10], [11], [12]. In complex, highly vascular areas like the liver, unsupervised network designs have demonstrated particular promise for correcting arterial-phase motion artifacts [13].

Although a large portion of the current literature on deep learning artifact correction concentrates on brain MRI [14], [15], [16], the abdominal region poses particular difficulties because of simultaneous respiratory, cardiac, and peristaltic motion. Advanced frameworks use iterative joint-learning techniques [18] and prior-assisted deep learning [17] to handle these complexities. Additionally, the "black box" tendency of deep models to smooth over microscopic pathologies has been mitigated by the introduction of novel hybrid loss functions, which have been crucial in guiding neural networks to preserve crucial high-frequency edge details during the denoising process [19], [20].

However, there are still a lot of unanswered questions despite this extensive body of research on GANs and MRI correction. For clinical diagnostics, where slight structural deviations can affect patient outcomes, the majority of studies still treat restoration as merely a visual enhancement task without strictly enforcing anatomical preservation constraints. Moreover, the critical vulnerability of training deep generative networks on sensitive healthcare data is often overlooked in the deployment of such models. Systems that combine cutting-edge restorative capabilities with theoretically sound structural constraints and integrated cryptographic privacy are clearly needed. By introducing a structurally constrained, privacy-preserving GAN specifically designed for the demanding requirements of liver MRI artifact reduction, this project fills this particular gap.

METHODOLOGY

A structured computational pipeline is used in the suggested framework for liver MRI scan artifact reduction. The preprocessing of volumetric data, the design of the structurally constrained Generative Adversarial Network (GAN), the incorporation of the differential privacy mechanism, and the experimental

setup used for training and assessment are all described in this section.

A. Data Acquisition and Volumetric Preprocessing

3D hepatic MRI scans saved in the Neuroimaging Informatics Technology Initiative (NIfTI, .nii.gz) format make up the main input. Arterial phase scans (art.nii.gz), pre-contrast scans (pre.nii.gz), and expert segmented liver masks (rater1_liver.nii.gz) are all included in the multi-contrast dataset. The nibabel library is used to systematically parse 2D cross-sectional slices from the 3D voxel arrays along the axial plane due to the substantial memory overhead associated with 3D volumetric computations.

Intensity normalization is a preprocessing step that scales voxel intensities to the interval $[-1, 1]$. To create paired artifact-laden training instances, clean slices are subjected to down-sampled k-space trajectories and synthetic artifact noise models. An 80% training set (X_{train}) and a 20% testing set (X_{test}) are created from the entire dataset of extracted slices.




Fig. 1: System Architecture of the Privacy-Preserving Structurally Constrained GAN.

B. Structurally Constrained GAN Architecture

We use a Structurally Constrained GAN to perform image restoration while removing the illusion of nonexistent hepatic features. The generator (G) creates a restored slice $\hat{y} = G(x)$ from an artifact-corrupted 2D MRI slice x . The discriminator (D) determines if an input image is a reconstructed output $\hat{y}$ or a true artifact-free scan y .

A multi-component loss function consisting of Adversarial Loss (L_{adv}), Pixel-wise L_1 Loss (L_{L1}), and Structural Similarity Loss (L_{SSIM}) is used to optimize the generator in order to enforce strict spatial and biological preservation:

$$L_G = L_{\text{adv}}(G, D) + \lambda_1 L_{L1}(G) + \lambda_2 (1 - \text{SSIM}(y, \hat{y})) \quad (1)$$

where λ_1 and λ_2 are weighting parameters. The SSIM penalty guarantees that high-frequency structural components, like hepatic parenchyma boundaries and vascular networks, stay intact.

[htbp] [1] 3D NIfTI liver MRI volumes (X_{NIfTI}), Privacy parameters (ϵ, δ), Noise multiplier (σ), Clipping bound (C). Trained Generator G^* , Privacy-preserving restored 2D slices. // Step 1: Volumetric Preprocessing Load 3D NIfTI volumes using nibabel and extract 2D axial slices. Normalize voxel intensities to $[-1, 1]$ and pair corrupted slices x with target clean slices y . Split data into training set (80%) and testing set (20%). // Step 2: Privacy-Preserving Adversarial Training each training epoch each mini-batch (x, y) Generate restored slices: $\hat{y} \leftarrow$

$G(x)$. Compute discriminator loss L_D using real pair (x, y) and fake pair ($x, \hat{y}$). Compute sample gradients $g_i \leftarrow \nabla_{\theta_D} \hat{L}_D(i)$. Clip gradients: $\bar{g}_i \leftarrow g_i / \max(1, |g_i|_2 / C)$. Add Gaussian noise: $\tilde{g} \leftarrow \bar{g} + \epsilon \cdot \mathcal{N}(0, I)$

$(\sum_i \bar{g}_i + N(0, \sigma^2 C^2 I))$. Update Discriminator weights θ_D using perturbed gradient $\tilde{g}$. Compute composite generator loss $L_G = L_{adv} + \lambda_1 L_{L1} + \lambda_2 L_{SSIM}$. Update Generator weights θ_G via standard backpropagation. // Step 3: Inference and Reconstruction Input test set X_{test} to optimized Generator G^* to yield artifact-reduced slices. Trained model G^* , quantitative performance metrics (SSIM, PSNR), and privacy spent ϵ .

C. Privacy-Preserving Mechanism via DP-SGD

Differentially Private Stochastic Gradient Descent (DP-SGD) is used during training to ensure that the generative model does not memorize distinct patient anatomical markers. Individual per-sample gradients g_i are bounded during the discriminator D update step by clipping their L_2 norm to a maximum threshold C :

$$\bar{g}_i = \frac{g_i}{\max(1, \frac{\|g_i\|_2}{c})} \quad (2)$$

Before the parameters are updated, the aggregated batch gradients are subjected to calibrated Gaussian noise with scale σC . By preventing membership inference and model inversion attacks that try to reconstruct the original patient identity, this mechanism guarantees strict (ϵ, δ) -differential privacy protection.

D. Experimental Setup

All training processes, network architectures, and data parsing were implemented in Python using PyTorch and run in the Google Colaboratory environment with NVIDIA GPU acceleration. nibabel and SimpleITK were used for medical image parsing, and numpy and scikit-image were used for scientific operations. The Opacus library was used to monitor differential privacy bounds. matplotlib was used to create intensity histogram displays, cross-sectional comparative plots, and visual evaluation. For effective deployment and clinical assessment, the fully trained generator weights were serialized using pickle.

E. Evaluation Metrics

The framework is evaluated using both structural fidelity and privacy guarantee metrics:

1. **Peak Signal-to-Noise Ratio (PSNR):** Quantifies spatial reconstruction fidelity and artifact suppression level.
2. **Structural Similarity Index Measure (SSIM):** Measures high-frequency structural, contrast, and edge conservation between target and restored MRI slices.
3. **Mean Squared Error (MSE):** Assesses voxel-level intensity variance across restored volumes.
4. **Privacy Budget (ϵ, δ) :** Tracks the strict privacy parameter bound spent over training epochs to ensure privacy compliance.

EXPERIMENTS AND RESULTS

This section bridges the gap between theoretical network architecture and clinical deployment by providing a detailed account of the empirical validation and practical implementation of our suggested privacy-preserving, structurally constrained framework. In order to handle high-dimensional volumetric MRI data, we first describe the thorough preprocessing pipeline. In order to create trustworthy ground-truth training pairs, this includes parsing 3D NIfTI files, extracting 2D cross-sectional arrays, normalizing voxel intensity, and artificially injecting motion-related phase-encoding artifacts.

After preparing the data, we explicitly specify the intricate hyperparameter configurations used during the generative adversarial training phase. We describe the optimizer choices, batch dimension scaling, and the particular weighting coefficients applied to our pixel-wise loss penalties and structural similarity (SSIM). Next, we offer a comprehensive evaluation of the network's restorative potential by contrasting our suggested architecture with unconstrained baseline models.

Lastly, we perform a thorough comparative analysis using quantitative image fidelity metrics like Mean Squared Error (MSE) and Peak Signal-to-Noise Ratio (PSNR) as well as qualitative visual inspections of hepatic boundaries. In this assessment, we deliberately measure the precise effect on the single-slice

restorations of the differential privacy mechanisms, namely gradient clipping and calibrated noise injection. This isolates the trade-off between privacy and utility, clearly demonstrating that our integrated structural constraints successfully preserve vascular sharpness, overall anatomical rigidity, and diagnostic clarity even when the model's ability to memorize unique patient identifiers is mathematically limited.

A. Volumetric Data Preparation and Synthetic Corruption

Raw 3D liver MRI scans saved in the standardized NIfTI format served as the starting point for our experimental phase. We first parsed the 3D voxel grids to extract successive 2D cross-sectional slices along the axial plane because deep learning models require standardized inputs and full 3D volumetric processing is extremely memory-intensive and computationally prohibitive. In order to stabilize the network gradients and speed up convergence during adversarial training, the pixel intensities across these extracted slices were normalized to a uniform scale

of $[-1, 1]$. We applied respiratory artifacts and mathematically modeled synthetic motion to the clean slices to produce our supervised training pairs. This was accomplished by simulating the ghosting and blurring usually brought on by patient movement during extended abdominal scans using localized phase shifts and k-space trajectory undersampling. After this transformation, the dataset, which included thousands of viable cross-sections, was divided into two parts: twenty percent was strictly sequestered as our unseen test set to assess generalization, and the other eighty percent was used to train the GAN to map corrupted inputs back to clean targets (see Figure 2).




Fig. 2: Volumetric NIfTI parsing and 2D slice preparation

B. Quantitative Performance Evaluation and Convergence

We ran the unseen test set through the fully trained generator in order to thoroughly assess how well the model understood the fundamental physics of artifact reduction. We relied on established fidelity metrics because traditional classification accuracy is not applicable in medical image restoration. The generated reconstructed images were directly compared to their original, unaltered ground truths. When compared to unconstrained baselines, the structurally constrained model performed exceptionally well, achieving notable gains in both Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity (SSIM) scores. Over the course of subsequent training epochs, the model consistently reduced the Mean Squared Error (MSE) across the test dataset, as shown in the performance chart (Figure 3). This quick and steady convergence verified that the generator effectively reduced background noise and ghosting without significantly altering the original voxel intensity distributions or adding artificial contrast shifts.




Fig. 3: Quantitative evaluation metrics (PSNR, SSIM, MSE)

C. Assessing Structural Fidelity and Anatomical Rigidity

Verifying that the model respected the intricate biological geometry of the liver was crucial from a clinical standpoint, going beyond pixel-level error minimization. We examined our integrated SSIM penalty's isolated impact over the whole dataset. The analysis showed that the network effectively isolated and eliminated artifact noise while rigorously maintaining high-frequency anatomical details, including the bile ducts, portal venous networks, and the hepatic parenchyma's outer borders (Figure 4). Unconstrained, baseline GAN architectures frequently experience mode collapse or hallucinate artificial textures, which can lead to the risky over-smoothing of crucial diagnostic areas. Our structurally limited method, on the other hand, consistently preserved anatomical reality. Geometric deviations were severely penalized by the objective function, which prevented the artificial creation of healthy tissue over areas with microscopic lesions or the correctly told part of early-stage fibrotic changes happened in it.

D. Visualizing Individual Slice Restoration and Privacy-Utility Trade-off

We performed a detailed examination of individual slice reconstructions in order to comprehend the useful, clinical result of our framework. A direct side-by-side comparison of the model's corresponding restoration and an artifact-corrupted slice from a single patient (showing severe phase-encoding striping and motion blur) is shown in Figure 5. The severe artifacts are effectively removed by the produced output, restoring the organ structures' clarity. Importantly, our differential privacy guarantees were upheld while achieving this high degree of visual clarity.

The DP-SGD mechanism successfully capped the privacy budget at a strict limit ($\epsilon \leq 3.0$) by continuously monitoring it during training. The calibrated noise injected into the discriminator's




Fig. 4: Global structural preservation analysis against baseline models

```

Initiating Phase 2 Full Training (100 Epochs)...
[Epoch 5/100] [Critic Loss: 0.1411] [Artist Loss: 1.7220] [Structural Score (SSIM): 0.9287]
[Epoch 10/100] [Critic Loss: 0.1360] [Artist Loss: 0.8805] [Structural Score (SSIM): 0.9477]
[Epoch 15/100] [Critic Loss: 0.1827] [Artist Loss: 1.7881] [Structural Score (SSIM): 0.9195]
[Epoch 20/100] [Critic Loss: 0.1143] [Artist Loss: 1.9409] [Structural Score (SSIM): 0.9415]
[Epoch 25/100] [Critic Loss: 0.0843] [Artist Loss: 1.2388] [Structural Score (SSIM): 0.9253]
[Epoch 30/100] [Critic Loss: 0.1481] [Artist Loss: 0.5967] [Structural Score (SSIM): 0.9733]
[Epoch 35/100] [Critic Loss: 0.2211] [Artist Loss: 0.4928] [Structural Score (SSIM): 0.9732]
[Epoch 40/100] [Critic Loss: 0.1522] [Artist Loss: 0.5697] [Structural Score (SSIM): 0.9824]
[Epoch 45/100] [Critic Loss: 0.1716] [Artist Loss: 0.5419] [Structural Score (SSIM): 0.9789]
[Epoch 50/100] [Critic Loss: 0.1329] [Artist Loss: 0.5882] [Structural Score (SSIM): 0.9786]
[Epoch 55/100] [Critic Loss: 0.1686] [Artist Loss: 0.3957] [Structural Score (SSIM): 0.9916]
[Epoch 60/100] [Critic Loss: 0.1523] [Artist Loss: 0.6282] [Structural Score (SSIM): 0.9796]
[Epoch 65/100] [Critic Loss: 0.1749] [Artist Loss: 0.9870] [Structural Score (SSIM): 0.9893]
[Epoch 70/100] [Critic Loss: 0.1443] [Artist Loss: 0.4297] [Structural Score (SSIM): 0.9926]
[Epoch 75/100] [Critic Loss: 0.1536] [Artist Loss: 1.0902] [Structural Score (SSIM): 0.9781]
[Epoch 80/100] [Critic Loss: 0.0841] [Artist Loss: 1.1186] [Structural Score (SSIM): 0.9477]
[Epoch 85/100] [Critic Loss: 0.1388] [Artist Loss: 0.5372] [Structural Score (SSIM): 0.9815]
[Epoch 90/100] [Critic Loss: 0.1163] [Artist Loss: 1.0870] [Structural Score (SSIM): 0.9749]
[Epoch 95/100] [Critic Loss: 0.1402] [Artist Loss: 0.5599] [Structural Score (SSIM): 0.9948]
[Epoch 100/100] [Critic Loss: 0.1227] [Artist Loss: 0.5662] [Structural Score (SSIM): 0.9886]
Phase 2 complete! Model weights successfully saved to Google Drive.

```

Fig. 5: Local restoration: Artifact input vs. Restored Output

gradients effectively decoupled patient identity and sensitive anatomical markers from the model weights, as evidenced by the localized edge details in the restored images remaining sharp and diagnostically useful. The findings show that the privacy-utility trade-off is optimally balanced, protecting data confidentiality without compromising the restored medical images' clinical utility.

CONCLUSION

In order to directly address the significant shortcomings of conventional generative models in clinical settings, the main goal of this research was to create a reliable, privacy-preserving framework for artifact reduction in 3D liver MRI scans. We effectively reduced complex respiratory and motion artifacts by implementing a structurally constrained Generative Adversarial Network optimized with a composite loss function that included both Structural Similarity (SSIM) and Pixel-wise penalties. Importantly, this method strictly preserved diagnostic targets like microscopic lesions and hepatic vascular networks while preventing the hallucination of artificial anatomical features. Additionally, the discriminator's training loop incorporated Differentially Private Stochastic Gradient Descent (DP-SGD), which successfully separated identifiable anatomical markers from the model's learned weights while maintaining strict patient data confidentiality. Both quantitative and qualitative analyses showed that the suggested architecture maintains a strictly constrained privacy budget while achieving better spatial fidelity and structural preservation than baseline models. In the end, this framework gives radiologists a safe, extremely accurate tool for hepatic diagnostics by bridging the gap between sophisticated deep learning restoration and strict healthcare data regulations. Future research will concentrate on expanding this privacy-preserving constraint mechanism to real-time inference environments and multi-modal volumetric datasets.

REFERENCES

1. Lundervold, Alexander Selvikva, and Arvid Lundervold. "An overview of deep learning in medical imaging focusing on MRI." *Zeitschrift fuer medizinische Physik* 29.2 (2019): 102-127.
2. Spieker, Veronika, et al. "Deep learning for retrospective motion correction in MRI: a comprehensive review." *IEEE Transactions on Medical Imaging* 43.2 (2023): 846-859.
3. Tripathi, Vijay R., Manish N. Tibdewal, and Ravi Mishra. "A survey on motion artifact correction in

- magnetic resonance imaging for improved diagnostics.” SN Computer Science 5.3 (2024): 281.
4. Singh, Ram, Navdeep Singh, and Lakhwinder Kaur. ”Deep learning methods for 3D magnetic resonance image denoising, bias field and motion artifact correction: a comprehensive review.” Physics in Medicine & Biology 69.23 (2024): 23TR01.
 5. Oksuz, Ilkay. ”Brain MRI artefact detection and correction using convolutional neural networks.” Computer methods and programs in biomedicine 199 (2021): 105909.
 6. Lee, Seul, et al. ”Deep learning in MR motion correction: a brief review and a new motion simulation tool (view2Dmotion).” Investigative Magnetic Resonance Imaging 24.4 (2020): 196-206.
 7. Park, Taeyong, et al. ”Deep learning-based automatic detection and grading of motion-related artifacts on gadoteric acid-enhanced liver MRI.” Investigative Radiology 58.2 (2023): 166-172.
 8. Pawar, Kamlesh, et al. ”Clinical utility of deep learning motion correction for T1 weighted MPRAGE MR images.” European journal of radiology 133 (2020): 109384.
 9. Skandarani, Youssef, Pierre-Marc Jodoin, and Alain Lalande. ”Gans for medical image synthesis: An empirical study.” Journal of Imaging 9.3 (2023): 69.
 10. Saju, Gulfam Ahmed, Marjan Akhi, and Yuchou Chang. ”Ensemble CycleGAN for Retrospective Rigid Motion Correction in MRI.” 2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). IEEE, 2024.
 11. Chen, Gang, et al. ”Mri motion correction through disentangled cyclegan based on multi-mask k-space subsampling.” IEEE transactions on medical imaging 44.4 (2024): 1907-1921.
 12. Saleem, Amal Muhammad, et al. ”Cycle consistent generative motion artifact correction in coronary computed tomography angiography.” Applied Sciences 14.5 (2024): 1859.
 13. Pan, Feng, et al. ”Correction of arterial-phase motion artifacts in gadoteric acid-enhanced liver MRI using an innovative unsupervised network.” Bioengineering 10.10 (2023): 1192.
 14. Yu, Jiangfan, et al. ”Deep Learning for Brain MRI Artifact Correction: Current Challenges and Future Directions.” Bioengineering 13.6 (2026): 699.
 15. Kidoh, Masafumi, et al. ”Deep learning based noise reduction for brain MR imaging: tests on phantoms and healthy volunteers.” Magnetic resonance in medical sciences 19.3 (2020): 195-206.
 16. Lee, Jongyeon, Byungjai Kim, and HyunWook Park. ”MC2-Net: motion correction network for multi-contrast brain MRI.” Magnetic Resonance in Medicine 86.2 (2021): 1077-1092.
 17. Chatterjee, Soumick, et al. ”Retrospective motion correction of MR images using prior-assisted deep learning.” arXiv preprint arXiv:2011.14134 (2020).
 18. Potter, Guy G. ”Iterative learning for joint image denoising and motion artifact correction of 3D brain MRI.” Medical image analysis (2025): 103872.
 19. Zhang, Lei, et al. ”Motion correction for brain mri using deep learning and a novel hybrid loss function.” Algorithms 17.5 (2024): 215.
 20. Duffy, Ben A., et al. ”Retrospective motion artifact correction of structural MRI images using deep learning improves the quality of cortical surface reconstructions.” Neuroimage 230 (2021): 117756.