

MPFlow: Multi-modal Posterior-Guided Flow Matching for Zero-Shot MRI Reconstruction

Seunghoi Kim^{1*} Chen Jin² Henry F. J. Tregidgo¹ Matteo Figini¹ Daniel C. Alexander¹

¹ University College London, UK ² Centre for AI, DS&AI, AstraZeneca, UK

Abstract. Zero-shot MRI reconstruction relies on generative priors, but single-modality unconditional priors produce hallucinations under severe ill-posedness. In many clinical workflows, complementary MRI acquisitions (e.g. high-quality structural scans) are routinely available, yet existing reconstruction methods lack mechanisms to leverage this additional information. We propose *MPFlow*, a zero-shot multi-modal reconstruction framework built on rectified flow that incorporates auxiliary MRI modalities at inference time without retraining the generative prior to improve anatomical fidelity. Cross-modal guidance is enabled by our proposed self-supervised pretraining strategy, *Patch-level Multi-modal MR Image Pretraining (PAMRI)*, which learns shared representations across modalities. Sampling is jointly guided by data consistency and cross-modal feature alignment using pre-trained *PAMRI*, systematically suppressing intrinsic and extrinsic hallucinations. Extensive experiments on HCP and BraTS show that MPFlow matches diffusion baselines on image quality using only 20% of sampling steps while reducing tumor hallucinations by more than 15% (segmentation dice score). This demonstrates that cross-modal guidance enables more reliable and efficient zero-shot MRI reconstruction. Code is available at <https://github.com/edshkim98/MPFlow>.

Keywords: MRI reconstruction · Generative models · Multi-modal

1 Introduction

Magnetic resonance imaging (MRI) reconstruction from low-quality measurements is an important inverse problem in medical imaging. MRI data are often sub-sampled or acquired with thick slices to reduce acquisition time or improve signal-to-noise ratio, resulting in substantial information loss.

Recent deep learning methods have achieved strong performance in MRI reconstruction [1,8,10,17,19,28,26,31], with diffusion models emerging as powerful generative priors for inverse problems in medical imaging. These diffusion-based approaches have been extended to zero-shot reconstruction [6,12,29,15,20], where a learned prior is guided by data consistency during sampling, eliminating the need for paired supervision. More recently, flow matching has been explored to

* Correspondence to: seunghoi.kim.17@ucl.ac.uk




Fig. 1: **Schematic diagram of the proposed framework.** (a) *PAMRI*: cross-modal feature alignment by learning a shared representation space across different MRI contrasts, (b) *MPFlow*: flow-matching prior is guided by both data consistency and *PAMRI*, to reduce intrinsic and extrinsic hallucinations.

improve sampling efficiency within the natural image domain [23,13]. However, these methods rely on unconditional priors, which may produce anatomically plausible yet incorrect reconstructions under severe ill-posedness.

This limitation manifests as *hallucinations*, which can be categorized as *intrinsic*, violating measurement consistency, or *extrinsic*, remaining measurement-consistent but unsupported by the ground truth (e.g. within the measurement null space) [2,15]. Importantly, extrinsic hallucinations may persist even when data-consistency errors are small. Single-modality priors lack a mechanism to resolve ambiguity among multiple plausible null-space solutions without updating its prior or providing additional information. In clinical practice, multi-sequence protocols, such as multi-parametric imaging (diffusion, perfusion MRI) or multi-contrast structural scans, are routinely acquired and provide complementary anatomical information. Although conditional approaches have explored multi-modal MRI reconstruction [5,11,18,7], they do not offer a principled mechanism to incorporate auxiliary modalities when the generative prior is unconditional.

To address this limitation, we introduce *MPFlow*, a zero-shot multi-modal MRI reconstruction framework based on flow matching that incorporates auxiliary MRI acquisitions at inference time without updating its unconditional prior. As illustrated in Fig. 1, we first train *PAMRI*, a self-supervised framework that learns patch-level cross-modal correspondences between target (e.g. T2) and auxiliary (e.g. T1) modalities via a patch-wise adaptive InfoNCE loss. The temperature is dynamically chosen by normalized mutual information of paired patches to relax the contrastive penalty for augmentation-distorted positive pairs, preserving modality-specific structural details. At inference, the rectified flow prior is jointly guided by data consistency and *PAMRI*, reducing *intrinsic and extrinsic hallucinations*. To the best of our knowledge, this is the first zero-shot MRI reconstruction framework incorporating auxiliary imaging modalities at inference time.

Our contributions are threefold: (1) We formulate multi-modal zero-shot MRI reconstruction, where an unconditional prior leverages auxiliary modalities at inference without modifying the prior, and show theoretically and empirically that

such conditioning reduces hallucinations. (2) We propose *MPFlow*, a zero-shot multi-modal flow-matching framework, integrating *PAMRI*, our self-supervised patch-level cross-modal alignment strategy, into a posterior update with noise optimization to mitigate poor trajectory initialization. (3) We demonstrate that *MPFlow* matches diffusion baselines in image quality with only 20% of the sampling steps, while reducing hallucinations by 15% in tumor segmentation (Dice score) and 26% in hallucination score (SHAFE) on BraTS and HCP.

2 Method

2.1 Problem Formulation

Let $\mathcal{X} \in \mathbb{R}^N$ for spatial domain tasks (e.g. super-resolution) and $\mathcal{X} \in \mathbb{C}^N$ for k-space tasks (e.g. subsampling), and let $x_{\text{true}} \in \mathcal{X}$ denote the target high-quality MRI scan. We consider the general degradation model:

$$y = \mathcal{F}(x_{\text{true}}) + \boldsymbol{\eta}, \quad (1)$$

where $y \in \mathcal{Y}$ are the acquired measurements, $\boldsymbol{\eta} \sim \mathcal{N}(0, \sigma^2 I)$ is additive Gaussian measurement noise, and $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}$ is the forward operator (e.g. k-space subsampling, spatial downsampling, or blurring). Generative priors such as diffusion or flow models regularize reconstruction via posterior inference, $\hat{x} \sim p(x | y)$, combining data consistency with a learned prior. However, when $\mathcal{F}$ is highly ill-posed, the posterior may assign probability mass to anatomically plausible but incorrect solutions, leading to hallucinations. While data consistency reduces intrinsic hallucinations, extrinsic hallucinations can still persist without prior-update or injecting additional information.

2.2 Hypothesis and Justification

We hypothesize that when registered modalities encode the same anatomy with complementary information, joint conditioning on measurements and an auxiliary modality reduces both intrinsic and extrinsic hallucinations.

Justification Let x denote the target-modality image, and x_{aux} a high-quality auxiliary modality. In a unimodal setting, the posterior $p(x | y) \propto p(y | x) p(x)$ is typically broad due to the weakly informative likelihood (under highly ill-posed setting), leading to extrinsic hallucinations.

From an information-theoretic perspective, the remaining uncertainty is quantified by the conditional entropy $H(x | y)$. When additionally conditioning on an auxiliary modality, the posterior becomes $p(x | y, x_{\text{aux}}) \propto p(y | x) p(x | x_{\text{aux}})$, and the uncertainty satisfies:

$$H(x | y, x_{\text{aux}}) = H(x | y) - I(x; x_{\text{aux}} | y), \quad (2)$$

where $I(x; x_{\text{aux}} | y)$ denotes the conditional mutual information provided by the auxiliary modality beyond the measurement. Since registered MR modalities share overlapping but non-identical anatomical information, $I(x; x_{\text{aux}} | y) > 0$, implying $H(x | y, x_{\text{aux}}) < H(x | y)$. Thus, data consistency and auxiliary conditioning can suppress intrinsic and extrinsic hallucinations, respectively.

2.3 Rectified Flow

We use rectified flow [21] from flow-matching as the generative prior. Rectified flow defines a continuous-time probability flow from a simple base distribution (e.g. Gaussian noise) to the data distribution via the ODE:

$$\frac{dx_t}{dt} = v_\theta(x_t, t), \quad t \in [0, 1], \quad (3)$$

where v_θ is a learned velocity field. We adopt a straight-line interpolation between noise $z \sim \mathcal{N}(0, I)$ at $t = 0$ and data x_1 at $t = 1$:

$$x_t = (1 - t)z + tx_1, \quad v(x_t, t) = \frac{dx_t}{dt} = x_1 - z. \quad (4)$$

The rectified flow objective minimizes

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{t,z,x_1} [\|v_\theta(x_t, t) - (x_1 - z)\|_2^2]. \quad (5)$$

Through this objective, rectified flow learns to approximate linear transport from noise to data. The near-linear and deterministic trajectories enable high-quality image generation with substantially fewer sampling steps than diffusion models.

2.4 Patch-level Multi-modal MR Image Pretraining (PAMRI)

We establish *PAMRI*, a contrastive feature alignment framework for multi-modal MR images. *PAMRI* uses independent encoders, $\phi(\cdot)$ for the target and $\psi(\cdot)$ for the auxiliary modality, to map them into a shared latent space. This independence disentangles modality-specific appearance from shared anatomical semantics. We adopt a lightweight yet effective ResNet18 [9] architecture for both encoders to minimize computational overhead at inference.

MRI reconstruction is a dense task requiring fine-grained structural information, but conventional contrastive learning discards spatial detail. To preserve it, we operate on patches (e.g. 32×32 pixels) and introduce an auxiliary reconstruction task. Given a batch of B paired patches $\{(p_i^{\text{tar}}, p_i^{\text{aux}})\}_{i=1}^B$, we compute normalized latent embeddings $u_i = \phi(p_i^{\text{tar}})/\|\phi(p_i^{\text{tar}})\|_2$ and $w_i = \psi(p_i^{\text{aux}})/\|\psi(p_i^{\text{aux}})\|_2$. This alignment is driven by our proposed adaptive contrastive objective. Building upon InfoNCE loss [4], our loss evaluates both inter- and intra-modal negatives: $\mathcal{L}_{\text{NCE}} = \frac{1}{2B} \sum_{i=1}^B (\ell(u_i, w_i) + \ell(w_i, u_i))$. The loss for a specific positive pair (e_i, e_j) , where $e_i \in \{u_i, w_i\}$, is defined as:

$$\ell(e_i, e_j) = -\log \frac{\exp(e_i^\top e_j / \tau_{i,j})}{\sum_{k=1}^B \exp(e_i^\top u_k / \tau_{i,k}) + \sum_{k=1}^B \exp(e_i^\top w_k / \tau_{i,k}) - \exp(e_i^\top e_i / \tau_{i,i})}. \quad (6)$$

Here, $\tau_{i,j}$ is an adaptive temperature scaled by the Normalized Mutual Information (NMI) between their corresponding raw patches: $\tau_{i,j} = \tau_{\min} + (\tau_{\max} - \tau_{\min})(1 - \text{NMI}(p_i, p_j))$. This dynamically adjusts the contrastive penalty based on the inherent multi-modal information overlap. Specifically, for positive pairs

substantially distorted by augmentations, the temperature increases, relaxing the pull to preserve modality-specific structural details crucial for downstream dense reconstruction.

Lightweight decoders $d_{\text{tar}}, d_{\text{aux}}$ are attached to regularize the representations via patch reconstruction:

$$\mathcal{L}_{\text{rec}} = \frac{1}{B} \sum_{i=1}^B \left(\|d_{\text{tar}}(\phi(p_i^{\text{tar}})) - p_i^{\text{tar}}\|_1 + \|d_{\text{aux}}(\psi(p_i^{\text{aux}})) - p_i^{\text{aux}}\|_1 \right). \quad (7)$$

The final objective becomes $\mathcal{L}_{\text{SSL}} = \mathcal{L}_{\text{NCE}} + \lambda_{\text{rec}} \mathcal{L}_{\text{rec}}$. Rather than using mutual information as a direct pixel-level alignment loss, which struggles with multi-modal intensity variations (e.g. a tumor is bright in FLAIR but nearly isointense in T1), *PAMRI* elegantly uses it to weight the latent contrastive space. By dynamically adjusting the temperature for naturally dissimilar patches, *PAMRI* learns robust, modality-invariant representations that enable spatially structured and semantically meaningful posterior guidance.

2.5 Multi-modal Posterior-Guided Flow Matching

We perform posterior inference by guiding rectified flow with data consistency (DC) and cross-modal alignment using pre-trained *PAMRI* encoders from Sec. 2.4. At each timestep t , we project the current state to the clean image manifold via the estimate $\hat{x}_{1|t} = x_t - (1-t)v_\theta(x_t, t)$ to drive the posterior update. Over T sampling steps, the update in each time step becomes:

$$v(x_t|y) = \underbrace{v_\theta(x_t, t)}_{\text{Prior}} - \alpha_t \nabla_{x_t} \left(\underbrace{\|\mathcal{F}(\hat{x}_{1|t}) - y\|_2^2}_{\text{Intrinsic Reduction}} + \underbrace{\lambda_P \mathcal{L}_P(\phi(\mathcal{P}(\hat{x}_{1|t})), \psi(\mathcal{P}(x_{\text{aux}})))}_{\text{Extrinsic Reduction}} \right), \quad (8)$$

where $v_\theta(x_t, t)$ and α_t denote the prior drift and step size. λ_P controls the $\mathcal{L}_P$'s strength. $\mathcal{L}_P$ is MSE between normalized latent features of patches extracted (excluding background) via $\mathcal{P}(\cdot)$ from the current estimate $\hat{x}_{1|t}$ and the auxiliary x_{aux} . DC reduces measurement space deviation, and $\mathcal{L}_P$ minimizes deviations from the auxiliary image in the latent space, systematically tackling both intrinsic and extrinsic hallucinations.

Initial noise optimization The initial noise z can influence the reconstruction quality of flow-matching. To mitigate poor initializations, we sample S candidate seeds $\{z^{(s)}\}_{s=1}^S \sim \mathcal{N}(0, I)$ and perform a short warm-start of t_{noise} steps in parallel. Each candidate is evaluated using the composite objective $\Phi(x) = \|\mathcal{F}(x) - y\|_2^2 + \lambda_P \mathcal{L}_P(x, x_{\text{aux}})$, and we select:

$$s^* = \arg \min_{s \in \{1, \dots, S\}} \Phi(\hat{x}_{1|t_{\text{noise}}}^{(s)}). \quad (9)$$

Sampling then continues only with the selected seed. Unlike previous method [30] that optimize via DC only, $\Phi(\cdot)$ jointly uses DC and *PAMRI* to select the candidate with lowest intrinsic and extrinsic hallucinations, yielding better posterior trajectories with minimal overhead.

3 Experiments

3.1 Experimental Setup

Datasets We evaluate our method on two widely used datasets: the Human Connectome Project (HCP) [25] and BraTS [22]. We perform T2 super-resolution ($4\times$) on HCP, and FLAIR k-space reconstruction ($8\times$) on BraTS, both using fully sampled T1 as the auxiliary modality. For HCP, we follow prior work [15] for the testing ($N=200$) and for BraTS, we use the official validation set ($N=250$) to follow the standard protocol. Experiments are conducted on axial slices.

Implementation Details Diffusion and rectified flows are trained on same dataset for 100k iterations. For *MPFlow*, we set $\lambda_P = 0.1$, $\lambda_{\text{rec}} = 0.5$, $S = 8$, and $t_{\text{noise}} = 0.2T$. Patch size 32×32 is used for *PAMRI* pretraining.

Baselines & Metrics We compare against DIP [27], DPS [6], DiffDeuR [20], and DynamicDPS [15] (without warm-start). Since no existing zero-shot method incorporates auxiliary imaging modalities, all baselines are unimodal. Reconstruction quality is evaluated using PSNR, SSIM, and LPIPS [32]. Hallucinations are assessed via (i) measurement-space error, (ii) tumor segmentation Dice score, and (iii) SHAFE (Semantic Hallucination Assessment via Feature Evaluation) [16]. For Dice score, a Swin-UNet [3] trained on BraTS FLAIR is applied to both ground-truth and reconstructed images. SHAFE measures semantic feature discrepancies via a pre-trained vision encoder (DINO [24]) with exponentially weighted aggregation to emphasize hallucinated structures.

Table 1: **Quantitative comparisons on HCP ($\text{SR} \times 4$) and BraTS (Acc. $\times 8$).** T denotes total time steps; DDIM used for $T = 100$. p-values are calculated against the second-best baseline for each metric.

Method	T	HCP ($\text{SR} \times 4$)			BraTS (Acc. $\times 8$)		
		PSNR ($\uparrow$)	SSIM ($\uparrow$)	LPIPS ($\downarrow$)	PSNR ($\uparrow$)	SSIM ($\uparrow$)	LPIPS ($\downarrow$)
DIP [27]	3K	24 $\pm$.7	.66 $\pm$.03	.30 $\pm$.03	25 $\pm$ 2.4	.57 $\pm$.10	.19 $\pm$.01
DiffDeuR [20]	.5K	23 $\pm$ 1.6	.78 $\pm$.04	.23 $\pm$.04	30 $\pm$ 1.9	.88 $\pm$.02	.12 $\pm$.02
DPS [6]	.1K	23 $\pm$ 1.4	.69 $\pm$.16	.20 $\pm$.08	23 $\pm$ 5.8	.65 $\pm$.27	.23 $\pm$.12
DPS [6]	.5K	24 $\pm$ 1.0	.79 $\pm$.03	.15 $\pm$.01	25 $\pm$ 2.4	.89 $\pm$.01	.09 $\pm$.02
DynamicDPS [15]	.1K	23 $\pm$ 1.8	.69 $\pm$.18	.21 $\pm$.09	24 $\pm$ 5.6	.68 $\pm$.16	.22 $\pm$.12
DynamicDPS [15]	.5K	24 $\pm$.9	.79 $\pm$.03	.15 $\pm$.02	31 $\pm$ 1.9	.89 $\pm$.02	.09 $\pm$.02
MPFlow (Ours)	.1K	23 $\pm$.9	.78 $\pm$.03	.16 $\pm$.02	31 $\pm$ 1.8	.90 $\pm$.02	.08 $\pm$.01
MPFlow (Ours)	.5K	25 $\pm$.9	.82 $\pm$.03	.14 $\pm$.02	31 $\pm$.9	.91 $\pm$.02	.07 $\pm$.01
p-value	-	<.001	<.001	.031	.016	.001	<.001

3.2 Main Results

Tab. 1 presents quantitative comparisons under varying sampling budgets. Across all tasks and metrics, *MPFlow* consistently outperforms zero-shot baselines, with

all improvements statistically significant ($p < 0.05$). At $T = 500$, *MPFlow* outperforms the second best baseline (DynamicDPS) by 2–4% in SSIM and 6–22% in LPIPS, despite DynamicDPS requiring multiple inner steps per timestep, demonstrating *MPFlow*’s efficiency. The improvements are more pronounced in the time-constrained setting ($T = 100$). Diffusion-based methods suffer large degradation with SSIM and LPIPS dropping 15–83% and increased variance, indicating unstable inference due to large discretization errors from DDIM sampling. In contrast, *MPFlow* exhibits marginal performance reduction (2.5–14%), demonstrating superior efficiency and robustness as the near-linear trajectories of rectified flow permit larger step sizes without significant discretization error. Remarkably, *MPFlow* matches the reconstruction quality of diffusion-based baselines at $T = 500$ using only $T = 100$ steps.

Fig. 2 presents visual comparisons of *MPFlow* against baseline methods. As highlighted in the red boxes, baselines introduce hallucinated structures including distorted sulci and incorrect tumor morphology, which can directly compromise clinical decisions such as surgical planning and radiotherapy contouring. In contrast, *MPFlow* preserves anatomically faithful structures with notably sharper tumor boundaries, demonstrating the benefit of cross-modal alignment via *PAMRI*. These qualitative results further demonstrate *MPFlow*’s ability to reduce hallucinations and improve structural fidelity in clinically critical regions.




Fig. 2: **Visual comparisons on HCP (top), BraTS (bottom).** Highlighted red boxes show *MPFlow* clearly reduces hallucinations compared to the baselines.

3.3 Further Analysis

Global image quality metrics (e.g. PSNR, SSIM) often fail to capture hallucinations, as they prioritize global fidelity while overlooking sparse, localized failures [14,16]. Instead, we evaluate hallucinations in Tab. 2 using measurement-space loss, SHAFE, and tumor segmentation Dice score. Across both datasets, *MPFlow* substantially outperforms the baselines in all hallucination metrics. Across both datasets, *MPFlow* (Full) substantially outperforms all baselines: it reduces measurement-space loss by over 80% on HCP and 88% on BraTS relative to DPS, while improving SHAFE by 38% and Dice by 16%, confirming that jointly guiding the prior with data consistency and *PAMRI* suppresses both

Table 2: **Quantitative hallucination evaluation on HCP and BraTS (T=100)**. We evaluate hallucination across baselines using measurement-space loss, SHAFE (HCP) and dice score on tumor segmentation (BraTS).

(a) Hallucination Analysis (HCP)			(b) Tumor Segmentation (BraTS)		
Method	Meas. L $\downarrow$	SHAFE $\downarrow$	Method	Meas. L $\downarrow$	Dice $\uparrow$
DPS	3.84×10^{-3}	0.278	DPS	4.20×10^{-2}	0.618
DynamicDPS	3.44×10^{-3}	0.274	DynamicDPS	4.51×10^{-2}	0.639
MPFlow (Base)	6.84×10^{-3}	0.233	MPFlow (Base)	4.36×10^{-3}	0.704
MPFlow (PAMRI)	1.43×10^{-3}	0.171	MPFlow (PAMRI)	2.45×10^{-3}	0.741
MPFlow (Noise)	1.25×10^{-3}	0.204	MPFlow (Noise)	4.90×10^{-4}	0.735
MPFlow (Full)	1.30×10^{-3}	0.170	MPFlow (Full)	4.89×10^{-4}	0.743


intrinsic and extrinsic hallucinations. To isolate the contribution of each component, we ablate *PAMRI* and noise optimization separately. Adding *PAMRI* to the vanilla baseline (Base $\rightarrow$ PAMRI) yields a 27% SHAFE reduction on HCP and 5% Dice improvement on BraTS, with comparatively smaller measurement-loss change, demonstrating that cross-modal guidance primarily suppresses *extrinsic* hallucinations by anchoring null-space content to the auxiliary anatomy. Conversely, noise optimization alone (Base $\rightarrow$ Noise) reduces measurement-space loss by 82% on HCP and 89% on BraTS with smaller SHAFE/Dice gains, indicating that better seed selection primarily suppresses *intrinsic* hallucinations by initializing trajectories that converge to measurement-consistent solutions. This confirms that the two modules address complementary failure modes.

Fig. 3a compares PAMRI against simpler auxiliary guidance losses; normalized mutual information (NMI), Canny edge, and pixel-MSE. PAMRI outperforms these baselines on both image quality and hallucination score. Notably, Canny edge and pixel-MSE hallucinate more than the DC-only baseline, indicating that naive heuristic guidance can adversely affect reconstruction quality. To assess *PAMRI*’s role in reducing extrinsic hallucinations, we hypothesize that its benefit should scale with task severity, since greater ill-posedness enlarges the null-space ambiguity. Fig. 3b confirms this: Δ SSIM grows from 3.81% at $4\times$ SR to 6.02% at $8\times$ SR, showing that cross-modal guidance becomes increasingly important as data consistency alone cannot resolve the null-space ambiguity.

4 Conclusion

We propose *MPFlow*, a multi-modal posterior-guided flow-matching framework for zero-shot MRI reconstruction. The auxiliary imaging modality is leveraged at inference time via *PAMRI*, our self-supervised contrastive pretraining strategy that learns shared cross-modal representations. At inference, the unconditional prior is guided by data consistency and *PAMRI*, suppressing both intrinsic and extrinsic hallucinations without updating the prior. Experiments on HCP and BraTS across super-resolution and k-space subsampling tasks show improvements in reconstruction fidelity and hallucination reduction. Future work explores extending the framework to non-imaging modalities and adaptive guid-

Fig. 3: **Further analysis of MPFlow on HCP (T=100).** (a) PAMRI vs. intensity-based auxiliary losses on image quality (SSIM $\uparrow$) and hallucination (SHAFÉ $\downarrow$). (b) Effect of PAMRI under increasing super-resolution difficulty.



ance schemes. In overall, our results show that auxiliary modalities can reshape posterior geometry at inference time under unconditional generative priors, enabling robust multi-modal inverse reconstruction without paired supervision.

Acknowledgments. This work is supported by the EPSRC-funded UCL Centre for Doctoral Training in Intelligent, Integrated Imaging in Healthcare (i4health) under grant number EP/S021930/1. The work of DCA is supported by the Wellcome Trust awards 221915 and 317797, and EPSRC grant EP/Y028856/1. The work of MF and HG is supported by the MRC (award MR/W031566/1), and the NIHR UCLH Biomedical Research Centre.

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

References

1. Alexander, D.C., Darko, Z., Ghosh, A., Tanno, R., Wottschel, V., et al.: Image quality transfer and applications in diffusion mri. *NeuroImage* **152**, 283–298 (2017)
2. Bhadra, S., Kelkar, V.A., Brooks, F.J., Anastasio, M.A.: On hallucinations in tomographic image reconstruction. *IEEE Transactions on Medical Imaging* **40**(11), 3249–3260 (2021)
3. Cao, H., Wang, Y., Chen, J., Jiang, D., Zhang, X., Tian, Q., Wang, M.: Swin-unet: Unet-like pure transformer for medical image segmentation. In: *Proceedings of the European Conference on Computer Vision Workshops(ECCVW)* (2022)
4. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709* (2020)
5. Choi, Y., Jung, S., Al-masni, M.A., Yang, M.H., Kim, D.H.: TESLA: Test-time Reference-free Through-plane Super-resolution for Multi-contrast Brain MRI . In: *proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2025* (2025)
6. Chung, H., Kim, J., Mccann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. In: *The Eleventh International Conference on Learning Representations* (2023)

7. Cicimen, A.G., Tregidgo, H.F.J., Figini, M., Messaritaki, E., B. McNabb, C., Palombo, M., Evans, C.J., Cercignani, M., Jones, D.K., C. Alexander, D.: Image quality transfer of diffusion mri guided by high-resolution structural mri. In: *Computational Diffusion MRI* (2025)
8. Gopinath, K., Hoopes, A., Alexander, D.C., Arnold, S.E., Balbastre, Y., et al.: Synthetic data in generalizable, learning-based neuroimaging. *Imaging Neuroscience* **2**, 1–22 (2024)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385* (2015)
10. Iglesias, J.E., Billot, B., Balbastre, Y., Magdamo, C., Arnold, S.E., et al.: SynthsR: A public ai tool to turn heterogeneous clinical brain scans into high-resolution t1-weighted images for 3d morphometry. *Science Advances* **9**(5) (2023)
11. Kang, L., Tang, B., Huang, J., Li, J.: 3d-mri super-resolution reconstruction using multi-modality based on multi-resolution cnn. *Computer Methods and Programs in Biomedicine* **248**, 108110 (2024)
12. Kavar, B., Elad, M., Ermon, S., Song, J.: Denoising diffusion restoration models. In: *Advances in Neural Information Processing Systems* (2022)
13. Kim, J., Kim, B.S., Ye, J.C.: FlowDPS: Flow-driven posterior sampling for inverse problems (2025)
14. Kim, S., Jin, C., Diethe, T., Figini, M., Tregidgo, H.F.J., et al.: Tackling structural hallucination in image translation with local diffusion. In: *European Conference on Computer Vision (ECCV)* (2024)
15. Kim, S., Tregidgo, H.F.J., Figini, M., Jin, C., Joshi, S., Alexander, D.C.: Tackling Hallucination from Conditional Models for Medical Image Reconstruction with DynamicDPS . In: *proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2025* (2025)
16. Kim, S., Tregidgo, H.F.J., Jin, C., Figini, M., Alexander, D.C.: Hallugen: Synthesizing realistic and controllable hallucinations for evaluating image restoration (2025), <https://arxiv.org/abs/2512.03345>
17. Kim, S., Tregidgo, H.F., Eldaly, A.K., Figini, M., Alexander, D.C.: A 3d conditional diffusion model for image quality transfer—an application to low-field mri. *arXiv preprint arXiv:2311.06631* (2023)
18. Li, G., Zhao, L., Sun, J., Lan, Z., Zhang, Z., Chen, J., Lin, Z., Lin, H., Xing, W.: Rethinking multi-contrast mri super-resolution: Rectangle-window cross-attention transformer and arbitrary-scale upsampling. In: *2023 IEEE/CVF International Conference on Computer Vision (ICCV)* (2023)
19. Lin, H., Figini, M., D’Arco, F., Ogbole, G., Tanno, R., et al.: Low-field magnetic resonance image enhancement via stochastic image quality transfer. *Medical Image Analysis* **87**, 102807 (2023)
20. Lin, X., Du, C., Wu, Q., Tian, X., Yu, J., et al.: Zero-shot Low-field MRI Enhancement via Denoising Diffusion Driven Neural Representation . In: *proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024* (2024)
21. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003* (2022)
22. Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K.e.a.: The multimodal brain tumor image segmentation benchmark (brats). *IEEE Transactions on Medical Imaging* **34**(10), 1993–2024 (2015)
23. Patel, M., Wen, S., Metaxas, D.N., Yang, Y.: Steering rectified flow models in the vector field for controlled image generation. *arXiv preprint arXiv:2412.00100* (2024)

24. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., Massa, F., Haziza, D., Wehrstedt, L., Wang, J., Darcet, T., Moutakanni, T., Sentana, L., Roberts, C., Vedaldi, A., Tolan, J., Brandt, J., Couprie, C., Mairal, J., Jégou, H., Labatut, P., Bojanowski, P.: DINOv3 (2025)
25. Sotiropoulos, S.N., Jbabdi, S., Xu, J., Andersson, J.L.R., Moeller, S., et al.: Advances in diffusion mri acquisition and processing in the human connectome project. *NeuroImage* **80**, 125–143 (2013)
26. Sriram, A., Zbontar, J., Murrell, T., Defazio, A., Zitnick, C.L., Yakubova, N., Knoll, F., Johnson, P.M.: End-to-end variational networks for accelerated mri reconstruction. *ArXiv* (2020)
27. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep image prior. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2018)
28. Wang, J., Chen, Y., Wu, Y., Shi, J., Gee, J.: Enhanced generative adversarial network for 3d brain mri super-resolution. In: *IEEE/CVF Winter Conference (WACV)* (2020)
29. Wang, Y., Yu, J., Zhang, J.: Zero-shot image restoration using denoising diffusion null-space model. *The Eleventh International Conference on Learning Representations* (2023)
30. Wu, H., He, L., Zhang, M., Chen, D., Luo, K., Luo, M., Zhou, J.Z., Chen, H., Lv, J.: Diffusion posterior proximal sampling for image restoration. In: *Proceedings of the 32nd ACM International Conference on Multimedia* (2024)
31. Zhang, K., Hu, H., Philbrick, K.A., Conte, G.M., Sobek, J.D., et al.: Soup-gan: Super-resolution mri using generative adversarial networks. *Tomography* **8**, 905 – 919 (2021)
32. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)