BiSA-DiT: A Bidirectional Spatially-Aware Diffusion Transformer for Sparse Tomographic Magnetic Particle Imaging

Xun Zhang*, Wenxuan Zou*, Yixiang Zhou, Ziwei Chen, Zeyu Zhang†, and Jie Tian†

School of Engineering Medicine, Beihang University, Beijing, China
zhang_zeyu@buaa.edu.cn, tian@ieee.org

Abstract. Magnetic Particle Imaging (MPI) is a promising tomographic medical imaging. However, System-Matrix (SM) calibration in MPI is prohibitively time-consuming. As a result, MPI image reconstruction using a sparse SM has attracted increasing attention. Nevertheless, reconstructions with a sparse SM often suffer from degraded in-plane details and structural discontinuities along the axial direction. To this end, we propose the BiSA-DiT framework, which improves tomographic reconstruction quality by alternately applying a powerful spatially aware generative prior with a physics-driven data consistency module. Specifically, by incorporating bidirectional synergistic Axial and Planar Spatial Transformer blocks, the generative-prior operator captures rich in-plane details while maintaining robust axial consistency. During the iterative optimization process, we employ an ADMM solver to ensure that the reconstructed results remain consistent with the MPI physical forward model. We validated BiSA-DiT on clinical coronary images using 4× sparse SM. Experimental results demonstrated that BiSA-DiT achieved high-quality tomographic MPI reconstruction with excellent image details and z-axis continuity. It significantly outperformed existing methods, achieving a substantial PSNR gain of 4.05 dB and a 2.72% increase in SSIM over state-of-the-art approaches.

Keywords: Magnetic Particle Imaging · Image Reconstruction · Diffusion Model.

1 Introduction

Magnetic Particle Imaging (MPI) [6] is an emerging functional medical imaging modality that exhibits promising advantages in sensitivity and spatial resolution. It has been applied in a range of biomedical applications [1, 8], including vascular imaging [9] and cell tracking [18]. Reconstructing MPI images from the measured signal relies heavily on the system matrix (SM), which is calibrated point by

* indicates equal contribution.
† indicates the corresponding author.

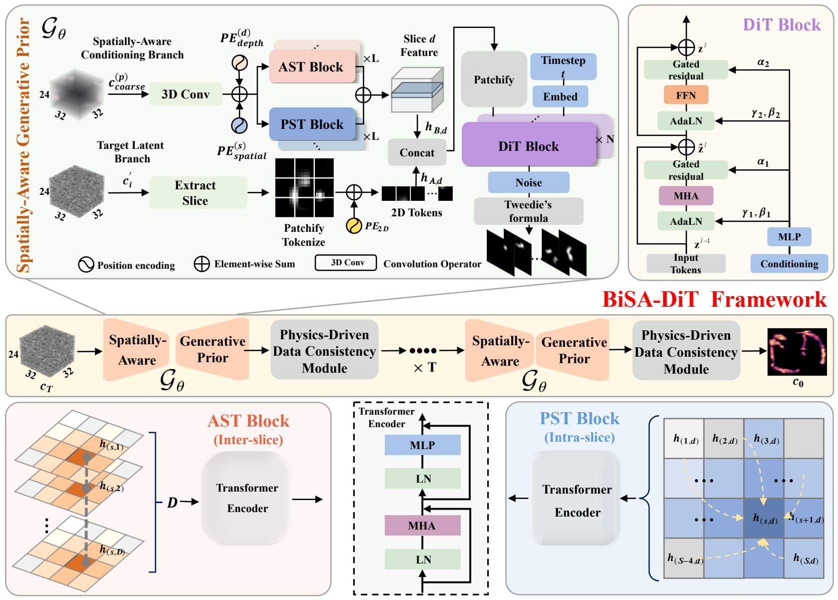

Fig. 1. Overview of the proposed BiSA-DiT framework.

point across the field of view. However, acquiring a full 3D SM is difficult and time-consuming because it requires tedious calibration scans [13]. This limitation restricts the application of tomographic MPI in clinical research. To address this challenge, developing MPI reconstruction algorithms based on sparse SM has received increasing attention [19, 16].

Medical image reconstruction from sparse space is inheranntly an ill-posed inverse problem. Current approaches are mainly categorized into model-based and learning-based algorithms [21]. Traditional model-based methods, such as Alternating Direction Method of Multipliers algorithm (ADMM)[2] or Kaczmarz [12], struggle to recover image structures from highly undersampled data [7]. Recently, deep learning algorithms, particularly Denoising Diffusion Probabilistic Models (DDPM) [10], have shown great potential in medical image reconstruction. They can capture complex data distribution and provide a powerful generative prior. Nevertheless, constrained by computational costs, such models are still difficult to directly learn 3D volumetric data distributions. To solve highly ill-posed 3D inverse problems, a prevailing research trend is to integrate model-based and learning-based algorithms. The Decomposed Diffusion Sampler (DDS) framework [3] utilizes pretrained 2D diffusion models to provide data-consistency priors for Conjugate Gradient (CG) iterations. Since DDS processes each slice independently, it ignores the axial continuity of medical images, leading to severe structural discontinuities in the reconstruction results. To improve upon this, DiffusionMBIR [4] incorporates z-axis Total Variation (TV) regularization term, attempting to reduce inter-slice inconsistency through handcrafted physical constraints. However, due to their simplified mathematical forms, these handcrafted priors have limited success to effectively recover axial structures from sparse data. Although the DiffusionBlend++ [15] to enhance performance by leveraging 3D patch context, its focus on local inter-slice relationships is still insufficient to guarantee global consistency along the z-axis. Furthermore, all the aforementioned methods overlook the synergy between intra-slice structural dependencies and inter-slice axial continuity.

To this end, we propose the BiSA-DiT framework, which integrates a Spatially-Aware Generative Prior operator $\mathcal{G}_\theta$ and a physics-driven data consistency module. The core of operator $\mathcal{G}_\theta$ lies in using parallel Planar Spatial Transformer (PST) and Axial Spatial Transformer (AST) blocks to establish bidirectional spatial awareness across both planar and axial dimensions. The spatial features output from these two modules serve as conditional guidance for the reverse denoising process of the DiT block [17], thereby enabling the generated priors to possess both planar details and robust axial consistency. Meanwhile, by applying the ADMM algorithm, the data consistency module ensures that the generated results remain aligned with the measured signals, thereby guaranteeing the physical fidelity of the reconstructed images. Experimental results demonstrate that in sparse SM reconstruction tasks, BiSA-DiT remains capable of reconstructing high-quality tomographic MPI volumes with superior z-axis continuity under undersampling conditions, significantly outperforming existing benchmark methods.

2 Methodology

2.1 Preliminary Knowledge

From a Bayesian perspective, the MPI reconstruction of c can be formulated as a Maximum A Posteriori (MAP) estimation. Therefore, the negative log-likelihood leads to the least-squares data fidelity term, while the negative log-prior acts as the regularization term $\mathcal{R}(c)$ to stabilize the ill-posed problem. This yields the standard regularized optimization task:

$$\hat{c} = \arg\min_{c} \left\{ \frac{1}{2} \|Sc - u\|_2^2 + \lambda \mathcal{R}(c) \right\}, \tag{1}$$

where $u \in \mathbb{C}^M$ is the measured signal, $c \in \mathbb{R}^N$ is the particle distribution image to be reconstructed, $S \in \mathbb{C}^{M \times N}$ is the SM characterizing the magnetic characteristics. To solve this problem, the ADMM is typically employed to decouple the physical constraint and the learned prior, iteratively refining the solution to achieve high fidelity while maintaining physical consistency.

2.2 Framework Overview

In this work, we propose BiSA-DiT, an iterative optimization framework for reconstructing high-fidelity and spatially continuous 3D MNP concentration distributions c from sparse system matrices S and measurements u, as illustrated in Fig. 1. First, we perform one-step sampling with the spatially-aware generative prior operator $\mathcal{G}_\theta$ to obtain an in-distribution generative prior. Next, a physics-driven data-consistency step constrains the prior toward the measurement hyperplane $u = Sc$. Then, the overall framework proceeds in an alternating manner, switching between the generative prior and physical consistency across.

2.3 Iterative Reconstruction Strategy

Specifically, at each timestep i, the BiSA-DiT framework coordinates the generative prior and physical optimization through the following two core steps:

Spatially-Aware Generative Prior. This stage generates a prior proposal c'_i by applying a one-step reverse denoising process [10]. It takes the current noisy state c_{i+1} as input to $\mathcal{G}_\theta$. Mathematically, this process follows Tweedie's formula [2] to estimate the posterior mean $\mathbb{E}[c_0|c_{i+1}]$:

$$\mathbb{E}[c_0|c_{i+1}] \approx c'_i = \mathcal{G}_\theta(c_{i+1}, i). \tag{2}$$

As an initial structural estimate, c'_i provides both axial and planar references. This acts as a high-quality starting point for the subsequent optimization.

Physics-Driven Data Consistency Update. This stage addresses the data-consistency constraint via a least-squares objective coupled with a z-axis TV term. It constructs the objective using the sparse S, measurements u, and variables from the previous iteration (z_{i+1}, w_{i+1}). Since c'_i provides a preliminary statistical estimate of c_0, we use it as the starting point for the Krylov subspace method [3] to accelerate convergence. By introducing the axial derivative operator D_z, the ADMM module suppresses reconstruction noise while maintaining a certain degree of axial continuity. The updated solution c_i effectively combines the generative prior with physical feasibility:

$$c_i = \mathrm{CG}(S^T S + \rho D_z^T D_z,\ S^T u + \rho D_z^T (z_{i+1} - w_{i+1}),\ c'_i). \tag{3}$$

Following the update of c_i, the algorithm further updates the auxiliary variable z_i and the dual variable w_i. Specifically, z_i is updated via the proximal operator of the ℓ_1 norm. Its closed-form solution is implemented via the soft-thresholding operator $\mathcal{P}_{\lambda/\rho}$. While the dual variable w_i is updated based on the residual term to ensure the convergence of the iterative process:

$$z_i = \mathrm{prox}_{|\cdot|_1, \lambda/\rho}(D_z c_i + w_{i+1}) = \mathcal{P}_{\lambda/\rho}(D_z c_i + w_{i+1}). \tag{4}$$

$$w_i = w_{i+1} + D_z c_i - z_i. \tag{5}$$

2.4 Generative Prior Operator $\mathcal{G}_\theta$ Architecture

In this section, we introduce the architectural details of the generative prior operator $\mathcal{G}_\theta$. The backbone of this operator is based on Conditional DiT. To implement conditional guidance within our iterative optimization framework, $\mathcal{G}_\theta$ employs a dual-branch input strategy:

Target Latent Branch. The Target Latent Branch is responsible for transforming the noisy target 2D slice into latent tokens for the DiT backbone. Before entering this branch, the target i-th slice of patient p, denoted as $c^{(p)}_{coarse,i}$, is prepared as a noisy representation $c^{*(p)}_{t,i}$ according to the forward diffusion process. The branch then processes this input through the following pipeline:

1) Tokenization and Positional Encoding: The branch performs spatial tokenization by splitting the noisy slice $c^{*(p)}_{t,i}$ into a series of non-overlapping patches via linear projection, mapping them into the hidden embedding space to form tokens. To incorporate spatial context, 2D sinusoidal positional encodings (PE_{2D}) are added to provide explicit coordinates for each patch:

$$h^{(p)}_{A,i} = \text{Linear}(\text{Patchify}(c^{*(p)}_{t,i})) + PE_{2D}. \tag{6}$$

The resulting sequence $h^{(p)}_{A}$ serves as the primary input for DiT, awaiting spatial prior guidance from Spatially-Aware Conditioning Branch for high-fidelity reconstruction.

Spatially-Aware Conditioning Branch. Unlike the Target Latent Branch, which only focuses on localized slices, the Spatially-Aware Conditioning Branch aims to extract a global spatial prior from the 3D coarse reconstruction volume $c^{(p)}_{coarse} \in \mathbb{R}^{1\times D\times H\times W}$.To efficiently capture comprehensive 3D context, we design a parallel structure extractor consisting of the PST and AST Modules.

1) Initial Coarse Reconstruction: Before entering the denoising stage, we construct an initial 3D coarse reconstruction volume, c_{coarse}, to serve as the foundational input. Given the sparse SM, we first use bilinear interpolation to upsample its spatial resolution to match that of the target image. Subsequently, we run a few iterations of an ADMM solver to obtain a coarse reconstruction.

2) 3D Tokenization and Embedding: We first map the 3D volume into tokens using 3D convolution and add decoupled positional encodings. For the p-th patient, the initial token $z^{(p,0)}_{(s,d)}$ at spatial location s and depth d is defined as:

$$z^{(p,0)}_{(s,d)} = \text{Conv3D}(c^{(p)}_{coarse})_{(s,d)} + PE^{(s)}_{spatial} + PE^{(d)}_{depth}, \tag{7}$$

where $s \in \{1,\ldots,S\}$ is the flattened 2D spatial index and $d \in \{1,\ldots,D\}$ denotes the axial depth index. $PE_{spatial}$ and PE_{depth} respectively represent the independently learned spatial and depth positional encodings. Subsequently, we duplicate this sequence as input for both PST and AST modules.

3) Bidirectional Spatial Transformer Blocks: To capture both intra-slice details and inter-slice continuity, we employ PST and AST modules. PST focuses on capturing anatomical details within each slice. In the l-th layer, self-attention is restricted to the same depth d, only operating across the spatial dimension S. The attention weights α^{PST} and feature updates are computed as:

$$\alpha^{PST}_{(s,s',d)} = \text{Softmax}\left(\frac{q^{PST}_{(s,d)} \cdot (k^{PST}_{(s',d)})^T}{\sqrt{C_h}}\right), \tag{8}$$

$$h_{(s,d)}^{PST,(p,l)} = h_{(s,d)}^{PST,(p,l-1)} + \sum_{s'=1}^{S} \alpha_{(s,s',d)}^{PST} v_{(s',d)}^{PST}, \tag{9}$$

where s' iterates over all spatial positions in the slice.

In parallel, the AST module captures inter-slice continuity by restricting self-attention to the same spatial location s and operating across the depth dimension D. Similar to PST, AST computes axial attention weights α^{AST} to aggregate features from all depth slices d', ensuring robust structural consistency along the z-axis.

4) Parallel Feature Fusion. Following L independent layers, we introduce learnable scalars γ_{PST} and γ_{AST} to fuse parallel features through the weighted element-wise sum [5]:

$$h_{(s,d)}^{B,(p)} = \text{LayerNorm}\left(\gamma_{PST} h_{(s,d)}^{PST,(p,L)} + \gamma_{AST} h_{(s,d)}^{AST,(p,L)}\right). \tag{10}$$

To align with the target slice from the Target Latent Branch, we extract the 2D feature sequence $h_{B,i}^{(p)} = \{h_{(s,i)}^{B,(p)}\}_{s=1}^{S}$ at depth $d = i$. Following the unified sequence strategy in OminiControl [17], this global spatial prior is concatenated with the noised latent $h_{A,i}^{(p)}$ along the channel dimension to form a conditional input. This fused representation $[h_{A,i}^{(p)}; h_{B,i}^{(p)}]$ is then fed into the DiT backbone ϵ_θ to guide the reverse denoising process.

Training Strategy. The Generative Prior Operator $\mathcal{G}_\theta$ is trained using the standard forward diffusion denoising objective. The parameters are updated by minimizing the Mean Squared Error (MSE) between the ground-truth noise ϵ and model prediction. Once trained, the operator $\mathcal{G}_\theta$ is embedded into the iterative framework to perform high-fidelity 3D reconstruction.

3 Experiments and Results

3.1 Dataset and Evaluation Metrics

Following the DEQ-MPI methodology [7], we utilized 1,000 clinical 3D coronary CTA images from the ImageCAS dataset [20] for our MPI dataset simulations. After extracting regions of interest (ROIs) via vessel segmentation bounding boxes, we applied uniform splitting and mean-intensity projection (MIP) techniques along the axial dimension to generate 800 training and 200 testing ground truth (GT) volumes, all with a resolution of $32 \times 32 \times 24$. In order to construct the sparse reconstruction dataset, we spatially downsampled the SM with a stride of 4 and simulated signals via the forward model. Specifically, we used the Open-MPI Experiment #7 SM [22], which was calibrated with Synomag-D tracers (Micromod GmbH), resulting in 19,200 training and 4,800 testing measurement-SM pairs. Reconstruction performance was quantitatively evaluated using Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), and normalized Root Mean Square Error (nRMSE) [22].

Table 1. Quantitative comparison of 3D MPI reconstruction methods on clinical coronary images using 4× sparse SM.

Method	PSNR (dB) ↑	SSIM ↑	nRMSE ↓
ADMM + TV [11]	19.47 ± 0.51	0.5960 ± 0.0609	0.1064 ± 0.0061
Kaczmarz [12]	18.59 ± 0.49	0.5717 ± 0.0583	0.1178 ± 0.0066
DiffusionMBIR [4]	19.71 ± 0.94	0.7580 ± 0.0397	0.1040 ± 0.0112
DDS - CG [3]	21.94 ± 0.76	0.6121 ± 0.0407	0.0803 ± 0.0070
DDS - ADMM + TV [3]	22.61 ± 0.80	0.7953 ± 0.0299	0.0744 ± 0.0068
DiffusionBlend++ [15]	23.01 ± 1.10	0.8606 ± 0.0331	0.0712 ± 0.0090
BiSA-DiT(Ours)	$\mathbf{27.06 \pm 1.13}$	$\mathbf{0.8878 \pm 0.0167}$	$\mathbf{0.0447 \pm 0.0059}$

Table 2. Ablation study of the AST and PST modules on 3D MPI reconstruction performance.

Settings	Module		Overall		
	PST	AST	PSNR (dB) ↑	SSIM ↑	nRMSE ↓
1 (Pure DiT) [14]			21.55 ± 0.95	0.7652 ± 0.0410	0.0842 ± 0.0091
2	✓		24.83 ± 0.90	0.8318 ± 0.0208	0.0576 ± 0.0059
3		✓	26.31 ± 1.10	0.8734 ± 0.0182	0.0488 ± 0.0061
4 (Ours)	✓	✓	$\mathbf{27.06 \pm 1.13}$	$\mathbf{0.8878 \pm 0.0167}$	$\mathbf{0.0447 \pm 0.0059}$

3.2 Implementation Details

All experiments were implemented using PyTorch 2.4.1 and conducted on 4 NVIDIA RTX 4090 (24GB) GPUs. The BiSA-DiT model was built upon the DiT-S/4 backbone. Specifically, the Spatially-Aware Conditioning Branch consists of 4 PST and 4 AST blocks. The model was trained for 400 epochs using the AdamW optimizer with a fixed learning rate of 1×10^{-4}, and the global batch size was set to 128. A linear noise schedule was applied with $T = 1,000$ and $\beta_t \in [10^{-4}, 0.02]$. For inference, a 200-step DDIM sampler was used for acceleration.We run ADMM for 5 iterations (1 CG step), with $\rho = 1.0$ and $\lambda = 1 \times 10^{-5}$.

3.3 Results

In our experiments, we benchmarked BiSA-DiT against six different types of reconstruction methods (including both model-based and learning-based iterative approaches). Table 1 shows the quantitative results on the test set, where BiSA-DiT demonstrates superior performance over all baselines. It achieves the highest PSNR (27.06 ± 1.13 dB), SSIM (0.8878 ± 0.0167), and the lowest nRMSE (0.0447 ± 0.0059). Compared to the strongest baseline, DiffusionBlend++ [15], our approach outperforms it by a substantial margin, with a PSNR gain of approximately 4.05 dB and a 2.72% increase in SSIM. Furthermore, we observe that BiSA-DiT exhibits significantly lower standard deviations in both SSIM and nRMSE. This phenomenon underscores the model's exceptional reconstruction stability and robustness when processing data from various subjects.

Fig. 2 (a) further supports the quantitative results discussed above. For Patient 881 and 941, traditional methods only reconstruct fragmented, blocky shapes because they lack structural priors. Although newer diffusion methods [4, 3, 15] offer some improvements, their error maps reveal that they still have

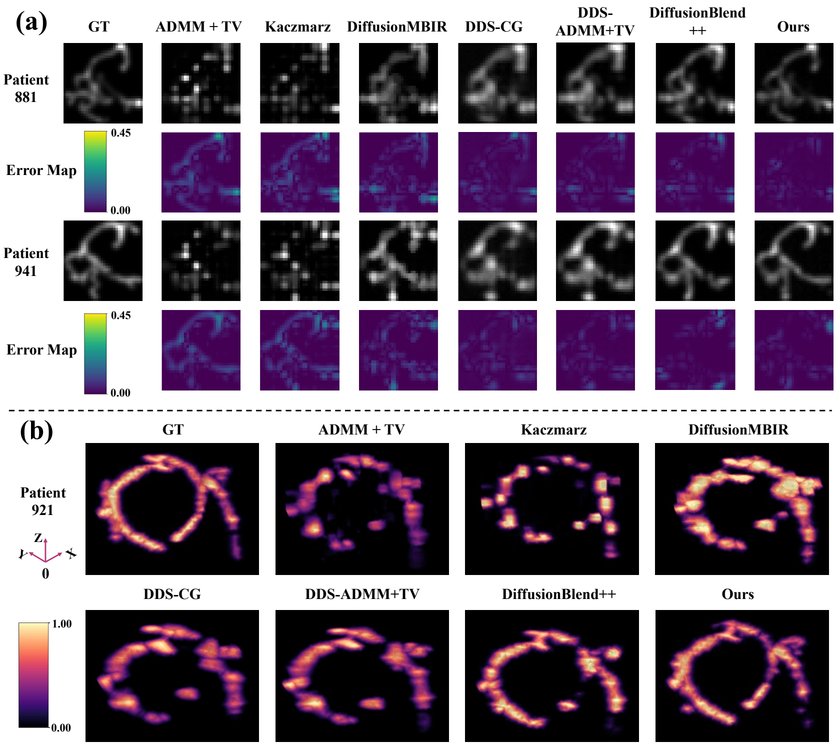

Fig. 2. Visualization of different 3D MPI reconstruction methods under Sparse SM. (a) Mean intensity projection results and corresponding error maps for Patients 881 and 941. (b) 3D MPI visualizations for Patient 921.

some structural deviations. In contrast, BiSA-DiT shows the closest agreement with the GT. Its error maps are mostly dark purple, reflecting minimal residuals. These results clearly show that the PST and AST modules in Branch B provide accurate spatial correlation guidance. This design effectively addresses the challenges caused by the extreme undersampling of signals and SM. Additionally, the 3D MPI reconstructions from different methods for Patient 921 are shown in Fig. 2 (b). Visual inspection reveals that BiSA-DiT exhibits superior z-axis continuity compared to other methods. It provides smoother connections between slices, creating more continuous 3D representations of vascular concentration.

As shown in Table 2, we conducted the ablation study to evaluate the contribution of each module in our framework. The results show that adding either the PST or AST module leads to clear improvements over the Pure DiT baseline. When these two modules are combined to form the complete BiSA-DiT framework, the model achieves the best performance. This confirms that both PST and AST modules are essential for high-quality 3D MPI sparse reconstruction.

4 Conclusion

In this paper, we introduced BiSA-DiT, a novel iterative optimization framework for 3D MPI reconstruction that combines a dual-branch spatially-aware generative prior operator with an ADMM-based data consistency module. Experiments demonstrate that BiSA-DiT achieves SOTA performance and significantly outperforms existing methods by establishing bidirectional spatial awareness and mitigating axial structural discontinuities. Future work will compare the direct solution of the sparse-SM inverse problem using generative priors against two-stage methods based on explicit SM super-resolution to further clarify the optimal strategy for maintaining physical fidelity.

Acknowledgements. This work was supported in part by the National Natural Science Foundation of China under Grants 625B2018, 62471019, 62027901, 92359301, U23A20487; in part by the Beijing Natural Science Foundation under Grant L241074; in part by Noncommunicable Chronic Diseases-National Science and Technology Major Project under Grant 2024ZD0525800; in part by Sichuan Clinical Research Center for Medical Imaging (Dazhou) under Grant YXYX2416.

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

References

1. Bulte, J.W.: Superparamagnetic iron oxides as mpi tracers: A primer and review of early applications. Adv. Drug Deliv. Rev. **138**, 293–301 (2019)
2. Chung, H., Kim, J., Mccann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. In: Int. Conf. Learn. Represent. (ICLR) (2023)
3. Chung, H., Lee, S., Ye, J.C.: Decomposed diffusion sampler for accelerating large-scale inverse problems. In: Int. Conf. Learn. Represent. (ICLR) (2024)
4. Chung, H., Ryu, D., et al.: Solving 3d inverse problems using pre-trained 2d diffusion models. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR). pp. 22542–22551 (2023)
5. Fu, J., Liu, J., et al.: Dual attention network for scene segmentation. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR). pp. 3146–3154 (2019)
6. Gleich, B., Weizenecker, J.: Tomographic imaging using the nonlinear response of magnetic particles. Nature **435**(7046), 1214–1217 (2005)
7. Güngör, A.: Deq-mpi: A deep equilibrium reconstruction with learned consistency for magnetic particle imaging. IEEE Trans. Med. Imaging **43**(1), 321–334 (2023)
8. Haegele, J., Rahmer, J., Gleich, B., Borgert, J., Wojtczyk, H., Panagiotopoulos, N., Buzug, T.M., Barkhausen, J., Vogt, F.M.: Magnetic particle imaging: visualization of instruments for cardiovascular intervention. Radiology **265**(3), 933–938 (2012)
9. Herz, S., Vogel, P., Kampf, T., Rückert, M., Veldhoen, S., Behr, V.C., Bley, T.A.: Magnetic particle imaging for quantification of vascular stenoses: a phantom study. IEEE Trans. Med. Imaging **37**(1), 61–67 (2017)

10. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Adv. Neural Inf. Process. Syst. (NeurIPS). vol. 33, pp. 6840–6851 (2020)
11. Ilbey, S., Top, C.B., Güngör, A., Çukur, T., Saritas, E.U., Güven, H.E.: Comparison of system-matrix-based and projection-based reconstructions for field free line magnetic particle imaging. Int. J. Magn. Part. Imaging **3**(1) (2017)
12. Knopp, T., Rahmer, J., et al.: Weighted iterative reconstruction for magnetic particle imaging. Phys. Med. Biol. **55**(6), 1577–1589 (2010)
13. Knopp, T., Weber, A.: Sparse reconstruction of the magnetic particle imaging system matrix. IEEE Trans. Med. Imaging **32**(8) (2013)
14. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: IEEE Int. Conf. Comput. Vis. (ICCV). pp. 4195–4205 (2023)
15. Song, B., Hu, J., et al.: Diffusionblend: Learning 3d image prior through position-aware diffusion score blending for 3d computed tomography reconstruction. In: Adv. Neural Inf. Process. Syst. (NeurIPS). vol. 37, pp. 89584–89611 (2024)
16. Sun, S., Zhong, J., et al.: Image reconstruction for magnetic particle imaging based on sparse representation and deep learning. IEEE Trans. Instrum. Meas. **73**, 1–9 (2024)
17. Tan, Z., Liu, S., Yang, X., Xue, Q., Wang, X.: Ominicontrol: Minimal and universal control for diffusion transformer. In: IEEE Int. Conf. Comput. Vis. (ICCV). pp. 14940–14950 (2025)
18. Them, K., Salamon, J., Szwargulski, P., Sequeira, S., Kaul, M., Lange, C., Ittrich, H., Knopp, T.: Increasing the sensitivity for stem cell monitoring in system-function based magnetic particle imaging. Phys. Med. Biol. **61**(9), 3279 (2016)
19. Von Gladiß, A., Knopp, T., et al.: Compressed sensing of the system matrix and sparse reconstruction of the particle concentration in magnetic particle imaging. IEEE Trans. Magn. **51**(2), 1–4 (2015)
20. Zeng, A., Wu, C., et al.: ImageCAS: A large-scale dataset and benchmark for coronary artery segmentation based on computed tomography angiography images. Comput. Med. Imaging Graph. **109**, 102287 (2023)
21. Zhang, K., Li, Y., et al.: Plug-and-play image restoration with deep denoiser prior. IEEE Trans. Pattern Anal. Mach. Intell. **44**(10), 6360–6376 (2021)
22. Zou, W., Tian, J., et al.: U-N2C: a dual memory-guided disentanglement framework for unsupervised system matrix denoising in magnetic particle imaging. IEEE Trans. Image Process. (2025)