PP-MPI: A Deep Plug-and-Play Prior for Magnetic Particle Imaging Reconstruction

Baris Askin[1(✉)], Alper Güngör[1,2], Damla Alptekin Soydan[2], Emine Ulku Saritas[1], Can Barış Top[2], and Tolga Cukur[1]

[1] Department of Electrical and Electronics Engineering, Bilkent University, Ankara, Turkey
barisaskin99@gmail.com, {alperg,saritas,cukur}@ee.bilkent.edu.tr
[2] Aselsan Research Center, Ankara, Turkey
{alpergungor,dasoydan,cbtop}@aselsan.com.tr

Abstract. Magnetic particle imaging (MPI) is a recent modality that enables high contrast and frame-rate imaging of the magnetic nanoparticle (MNP) distribution. Based on a measured system matrix, MPI reconstruction can be cast as an inverse problem that is commonly solved via regularized iterative optimization. Yet, hand-crafted regularization terms can elicit suboptimal performance. Here, we propose a novel MPI reconstruction "PP-MPI" based on a deep plug-and-play (PP) prior embedded in a model-based iterative optimization. We propose to pre-train the PP prior based on a residual dense convolutional neural network (CNN) on an MPI-friendly dataset derived from magnetic resonance angiograms. The PP prior is then embedded into an alternating direction method of multiplier (ADMM) optimizer for reconstruction. A fast implementation is devised for 3D image reconstruction by fusing the predictions from 2D priors in separate rectilinear orientations. Our demonstrations show that PP-MPI outperforms state-of-the-art iterative techniques with hand-crafted regularizers on both simulated and experimental data. In particular, PP-MPI achieves on average 3.10 dB higher peak signal-to-noise ratio than the top-performing baseline under variable noise levels, and can process 12 frames/sec to permit real-time 3D imaging.

Keywords: Magnetic particle imaging · Reconstruction · Plug and play · Deep learning

1 Introduction

Magnetic Particle Imaging (MPI) is a recent imaging modality that allows high contrast imaging of magnetic nanoparticles (MNP) with high frame rate. Important applications include cancer imaging, stem cell tracking, angiography and

N. Haq et al. (Eds.): MLMIR 2022, LNCS 13587, pp. 105–114, 2022.
https://doi.org/10.1007/978-3-031-17247-2_11

targeted drug delivery [14,15,17]. In MPI, a field free region (FFR) is generated to later measure a signal that reflects the total MNP concentration within the targeted region. Since the signal response is influenced by properties of the imaging system and MNP characteristics, a common procedure is to measure a system matrix (SM) that characterizes the forward signal model. Image formation from measured MPI signals can then be cast as an inverse problem based on the SM. As the inverse problem is ill-posed, reconstructions characteristically embody regularization terms that reflect prior information on the characteristics of MNP distribution.

A common reconstruction method is based on ℓ_2-norm regularization during algebraic reconstruction technique (ART) optimization [10]. Compressed sensing methods were also proposed for image regularization via ℓ_1-norm and/or total variation (TV) terms during an alternating direction method of multipliers (ADMM) optimization [8]. While prominent results have been reported, handcrafted regularization terms can induce well-known blurring, or blocking artifacts in the reconstructed image. As an alternative, several recent studies have proposed purely learning-based approaches. A deep image prior (DIP) method has been considered in [5] that performs reconstruction based on an untrained network via inference optimization, which may limit real-time applicability. End-to-end methods trained to map MPI signals to images have also been introduced [3,7,12]. Yet, end-to-end methods require retraining under changes to the SM, and their success relies on training datasets that match the characteristics of MPI images. Given the absence of a dedicated database of MPI images, previous studies have primarily used training sets containing simulated vessel phantoms or numerical digit images [4]. This may restrict generalization performance for experimental data and 3D imaging scenarios.

Here, we introduce a novel deep plug-and-play (PP) image prior for model-based MPI reconstruction. Inspired by the success of task-agnostic PP priors in medical imaging [1,13], we propose to train a denoising PP prior for MPI based on a training set derived from time-of-flight cerebral magnetic resonance angiograms (MRA) to closely mimic the vascular structures targeted in MPI scans. We build the prior on a residual dense convolutional neural network (CNN) [16]. During inference, we embed trained prior into a model-based ADMM optimization that includes measured SM. The PP prior enhances the quality of reconstructed images while ADMM offers a fast implementation suitable for real-time imaging. To utilize 2D dataset in 3D reconstruction and improve computational efficiency, we propose to fuse the predictions from denoised cross-sections of separate rectilinear orientations. Our main contributions are: (1) We introduce the first PP method for MPI reconstruction, (2) We propose an efficient algorithm based on ADMM and fused 2D priors for 3D reconstruction, (3) We validate the proposed method on simulated and experimental datasets [11].

2 Background

2.1 MPI Signal Model

During an MPI scan, drive field oscillations occurring at a fundamental frequency excite the MNPs, while a spatially encoded selection field creates an FFR. MNPs in FFR respond to the drive field, and responses are recorded as voltage waveforms induced on a receive coil. The MNP response is apparent in frequency bands centered on the harmonics of the fundamental frequency. The received signal can be compressed by filtering out low intensity bands in the frequency domain. The filtered signal is complex valued, and commonly represented by concatenating real and imaginary parts of the signal in separate rows of the measurement vector. The forward imaging model is:

$$\mathrm{A}\mathbf{x} + \mathbf{n} = \mathbf{r}, \tag{1}$$

where $\mathrm{A} \in \mathbb{R}^{M \times N}$ is the SM (M denotes twice the number of frequency components and N is the number of points on the imaging grid), $\mathbf{x} \in \mathbb{R}^N$ is the underlying MNP distribution, $\mathbf{n} \in \mathbb{R}^M$ is the measurement noise, and $\mathbf{r} \in \mathbb{R}^M$ denotes the measured MNP responses. Given the SM A, the forward model expressed in Eq. (1) can be inverted to estimate the MPI image $\mathbf{x}$ given the measurement $\mathbf{r}$. However, non-idealities regarding the imaging system and MNP characteristics prohibit accurate analytical calculation of the SM. Thus, it is common practice to experimentally measure A via calibration measurements [6]. Calibration is performed while a voxel-sized MNP sample is located at each grid point across the imaging Field of View (FOV), separately for each frequency component.

2.2 MPI Image Reconstruction

MPI reconstruction involves solution of the inverse problem characterized with the forward model in Eq. (1). Yet, since the problem is ill-conditioned, regularization terms are typically enforced to guide the reconstruction:

$$\arg\min_{\mathbf{x} \geq 0} \sum_i \alpha_i f_i(\mathbf{x}) \;\; s.t \;\; \|\mathrm{A}\mathbf{x} - \mathbf{r}\|_2 < \epsilon, \tag{2}$$

where α_i is the regularization weight for the i^{th} regularization function $f_i(\cdot)$, and ϵ is the upper bound on noise level reflecting measurement and SM errors. Note that there is a non-negativity constraint on $\mathbf{x}$ as MNP distribution cannot assume negative values. The most common approach to solve the model-based reconstruction problem in Eq. (2) is via iterative optimization methods. Kluth et al. proposed the ART method with ℓ_2-norm regularization ($f_1(\mathbf{x}) = ||\mathbf{x}||_2^2$), which assumes that the underlying MNP distribution is spatially smooth [9]. Ilbey et al. proposed the ADMM method with mixed ℓ_1-norm and TV regularization ($f_1(\mathbf{x}) = ||\mathbf{x}||_1$ and $f_2(\mathbf{x}) = \mathrm{TV}(\mathbf{x})$), which assumes that the MNP distribution is sparse in image and finite-differences domains [8]. Despite their pervasive use,

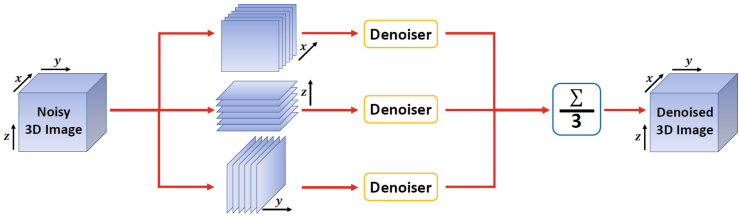

Fig. 1. Multi-orientation 2D PP priors to achieve efficient processing of 3D MPI data. Volumetric data are split into cross-sections across x, y and $z-$axes. The same denoiser processes the cross-sections in each orientation. Resultant volumes are averaged across orientations to produce the denoised 3D data.

Algorithm 1. An ADMM algorithm for PP-MPI

Initialize $z_0^{(i)}$ and $d_0^{(i)}$ for $i = 0, 1$, choose μ, set $n \leftarrow 0$
while Stopping criterion is not satisfied **do**
$\mathbf{x}_{n+1} \leftarrow (m\mathrm{I} + \mathrm{A}^H\mathrm{A})^{-1}(\mathrm{A}^H(\mathbf{z}_n^{(0)} + \mathbf{d}_n^{(0)}) + \mathbf{z}_n^{(1)} + \mathbf{d}_n^{(1)})$
$\mathbf{z}_{n+1}^{(1)} \leftarrow f_{PP}(\mathbf{x}_{n+1} - \mathbf{d}_n^{(1)}; \alpha_1, \mu)$ ▷ Variable update for PP subproblem
$\mathbf{d}_{n+1}^{(1)} \leftarrow \mathbf{d}_n^{(1)} + \mathbf{z}_{n+1}^{(1)} - \mathbf{x}_{n+1}$ ▷ Lagrange multiplier update for PP subproblem
$\mathbf{z}_{n+1}^{(0)} \leftarrow \Psi_{l_{E(\epsilon,I,\mathbf{b})}}(\mathrm{A}\mathbf{x}_{n+1} - \mathbf{d}_n^{(0)})$ ▷ Projection for DF subproblem
$\mathbf{d}_{n+1}^{(0)} \leftarrow \mathbf{d}_n^{(0)} + \mathbf{z}_{n+1}^{(0)} - \mathrm{A}\mathbf{x}_{n+1}$ ▷ Lagrange multiplier update for DF subproblem
$n \leftarrow n + 1$
end while

methods that rely on hand-constructed regularizers can yield suboptimal performance when underlying assumptions are not satisfied on experimental data. Several recent studies have adopted purely learning-based reconstruction methods. In a DIP approach, an untrained neural network was optimized for reconstruction at test time based on a data-fidelity objective [5]. Because DIP requires prolonged inference, it is not ideally suited for real-time MPI. Multi-layer perceptron (MLP) architectures were proposed for MPI reconstruction [3,7]. The MLP model was trained and tested on simulated phantoms, so experimental utility remains to be demonstrated. A CNN was proposed that was similarly trained and validated on simulated vessel phantoms and numerical digit images [4,12]. While data-driven priors promise to mitigate the limitations of hand-constructed regularizers, end-to-end networks capture a prior conditioned on the SM. Therefore, they have to be retrained for notable changes in the MPI system such as scan trajectories or MNP properties.

3 Methods

3.1 Plug-and-Play MPI Reconstruction (PP-MPI)

Here, we propose a novel plug-and-play method, PP-MPI, for reconstruction of MPI images. To improve reliability against changes in the system properties, PP-MPI decouples the forward imaging model characterized by the SM from training of the deep prior for data-driven regularization. As such, a task-agnostic deep prior is first trained for image denoising on an MPI-friendly dataset. During inference, the trained prior is then adapted to perform MPI reconstruction by incorporating the forward imaging model in Eq. (1). For fast inference, we devised an ADMM-based algorithm comprising two subproblems: a first problem for data-driven regularization based on the PP prior, and a second problem for data-fidelity projection based on the SM. Taken together, these design elements enable PP-MPI to offer reliable and rapid reconstruction of MPI data.

MPI-Friendly Dataset: In many MPI studies, a limited number of simulated vessel phantoms are used for demonstration of reconstruction methods as vascular imaging is a key application. However, it is difficult to generate a broad diversity of manually designed phantoms for training deep networks. We instead propose to generate MPI-like training images from abundant time-of-flight MRA images in the public "ITKTubeTK - Bullitt - Healthy MR Database" by CASILab at the University of North Carolina-CH. Data from 95 healthy subjects are split into non-overlapping training (77 subjects), validation (9) and test (9) sets. Multiple 3D patches of size $10 \times 64 \times 64$ are randomly cropped, a thin-slab maximum-intensity projection is performed across the first dimension, and the resulting projection is downsampled to 32×32 MPI-like images. To avoid of

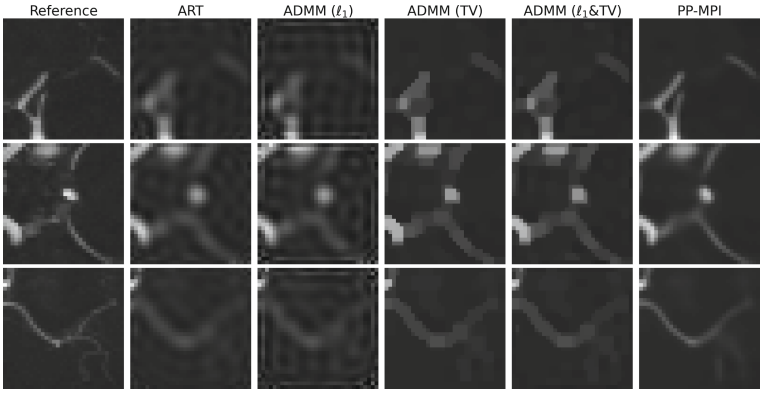

Fig. 2. 2D reconstructions for simulated vessel phantoms under 40 dB measurement SNR. Images from competing methods are shown along with reference phantom images for three separate phantoms (each phantom depicted on a separate row).

Table 1. Average PSNR (dB) for reconstructed simulated phantom images with varying measurement SNRs (20/30/40 dB SNRs).

Method	20 dB SNR	30 dB SNR	40 dB SNR
ART	24.57	26.79	27.87
ADMM (ℓ_1)	15.53	18.83	24.48
ADMM (TV)	25.99	27.04	27.64
ADMM (ℓ_1&TV)	25.89	27.13	28.26
PP-MPI	**29.01**	**30.49**	**31.18**

empty patches, patches with higher ℓ_2-norm of pixel intensities were retained. Resultant 2D images are normalized to a maximum intensity of 1.

Training the PP Prior: We leverage a Residual Dense Network (RDN) as backbone architecture for the PP prior given its success in image restoration tasks in computer vision studies [16]. As MPI images are substantially smaller than natural images, we devise a specialized RDN architecture with a compact set of parameters in order to match the characteristics of MPI data. We also employ rectified linear unit at the output layer to enforce non-negativity of MPI image intensities. The resulting model is trained for image denoising on the MPI-friendly dataset. Here, we consider both 2D and 3D reconstruction. For computational efficiency, we propose to employ 2D priors for both reconstruction scenarios. Because denoising cross-sections of one rectilinear orientation will offer suboptimal capture of context along the through-plane dimension, we introduce a volumetrization strategy for 3D reconstruction. During reconstruction, we first denoise cross-sections of three orientations using the same 2D prior. Then, we fuse their predictions via averaging (Fig. 1).

ADMM-Based Inference: The trained PP prior is integrated with the forward imaging model during inference on actual MPI data. An inference optimization is performed to find a reconstruction that is consistent with the PP prior and satisfies data-fidelity (DF) to MPI measurements. For fast inference, an ADMM based implementation is developed [8], where classical proximal operators in Eq. (2) for regularization functions are replaced with projections through the PP prior. The proposed algorithm splits the overall inference optimization into two easier subproblems as outlined in Algorithm 1, where $f_{PP}(\cdot)$ denotes the PP prior trained for image denoising, $\Psi_{l_{E(\epsilon,I,\mathbf{b})}}(\cdot)$ is the proximal mapping for indicator function $l_{E(\epsilon,I,\mathbf{b})}$ of being the element of set $E(\epsilon, I, \mathbf{b})$ as described in [8], and μ indicates the step size of the algorithm. μ and α_1 were set to 1. The maximum number of iterations n_{max} was selected as a stopping criterion via cross-validation. Note that large training datasets containing 3D imaging data may not be broadly available, and inference optimization with a 3D prior might elevate computational burden. Thus, here, we instead adopted a volumetrization strategy based on multi-orientation 2D priors.

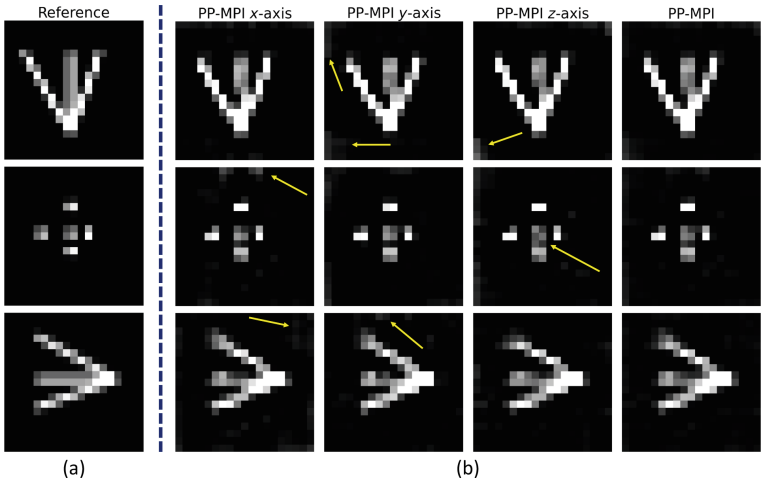

Fig. 3. 3D reconstructions of the OpenMPI resolution phantom with PP-MPI are shown along with the phantom's CAD model as visual reference. Central cross-sections along x, y, z axes are shown in separate rows. Arrows indicate notable artifacts.

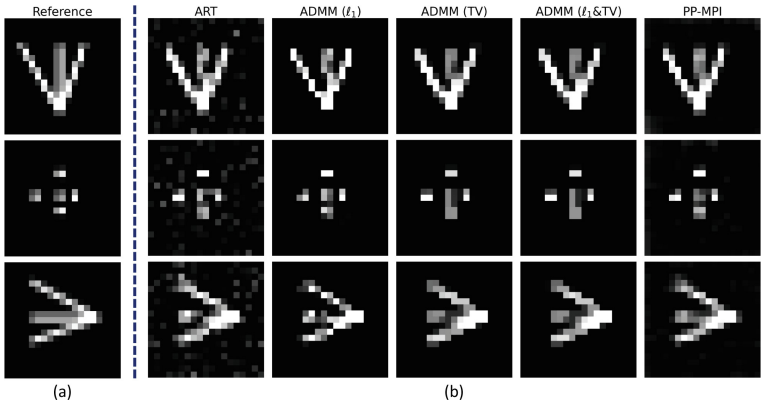

Fig. 4. 3D reconstructions of the OpenMPI resolution phantom are shown along with the CAD model of the phantom as visual reference. Central cross-sections along x, y, z axes are shown in separate rows.

3.2 Analyses

We demonstrated PP-MPI for reconstructing images from a 2D simulated dataset and a 3D experimental dataset (Open MPI) [11]. First, 2D PP priors were trained for denoising on the MPI-friendly dataset. Hyperparameters of the RDN architecture were tuned based on denoising performance on the validation set. Accordingly, RDN had 3 residual blocks, with 5 convolutional layers and 8 channels each. Original MPI-friendly images were taken as clean references, and noisy versions were then generated by adding white Gaussian noise with

a standard deviation of 0.05. The RDN model was then trained to predict the clean images provided as input noise variants. Model training was performed via the ADAM optimizer and with ℓ_1-loss between predicted and reference images.

Simulated Dataset: Field-free-line (FFL) based SM was simulated for an FOV of 32×32 mm^2, imaging grid of 64×64. For reconstruction, the SM was then downsampled to 32×32 to avoid inverse crime [2]. MNP saturation, selection field and drive field were $0.55/\mu_0$, 2.5 T/m, 40 mT. The 2nd to 20th harmonics of the received signal were retained. White Gaussian noise was added to the measurement vector simulated for a numerical vessel phantom. To improve SNR, singular value truncation was performed on the SM and the measurement vector ($c = 220$ singular values out of $N = 1024$). Reconstruction performance was examined visually and via peak signal-to-noise ratio (PSNR).

Open MPI: Field-free point (FFP) data for the resolution phantom measured on a 3D scanner (Bruker, Ettlingen) with a 3D Lissajous trajectory and Perimag (Micromod Partikeltechnologie GmbH, Germany) were used. The FOV was $38 \times 38 \times 19$ mm^3, and the grid size was $19 \times 19 \times 19$. The MNP saturation, selection field and drive field were 50 mmol/L, $-0.5/-0.5/1$ T/m, 12 mT. Singular value truncation was used ($c = 2200$ singular values out of $N = 6859$). As no ground truth image is available, performance was assessed visually against the computer aided design (CAD) model of the phantom.

Competing Methods: We comparatively demonstrated PP-MPI against methods that use task-agnostic priors for flexible generalization to changes in system properties that would be captured by the SM. Accordingly, baselines included ART with ℓ_2-norm regularization [9], and ADMM with ℓ_1-norm, TV-norm and mixed ℓ_1&TV norm regularization [8]. Note that we did not consider end-to-end learning-based methods that require retraining under changes to the SM [3], and DIP methods that have long inference times that prohibit real-time processing [5]. All methods were implemented in PyTorch on a Tesla V100 GPU.

4 Results

PP-MPI was first demonstrated for 2D reconstruction of simulated vessel phantoms. PSNR metrics from competing methods are listed in Table 1. On average, PP-MPI outperforms the closest competitor by 3.10 dB PSNR across the test set. Representative reconstructions are displayed in Fig. 2. ART suffers from smoothing and ADMM variants have relatively poor vessel localization with occasional block artifacts. In contrast, PP-MPI alleviates reconstruction artifacts and maintains sharp vessel localization.

We then demonstrated PP-MPI for 3D reconstruction of experimental data. In this case, we first compared results from the multi-orientation 2D priors against 2D priors along singular axes. Representative results are displayed in Fig. 3. While priors in singular directions can suffer from reconstruction artifacts in other orientations containing the through-plane axis, PP-MPI with multi-orientation priors achieves uniformly high quality across all orientations. Afterwards, PP-MPI with multi-orientation priors was compared against competing

baselines as illustrated in Fig. 4. ART has prominent background artifacts and vessel smoothing, whereas ADMM variants produce block artifacts particularly in the vicinity of areas with high MNP concentration. In contrast, PP-MPI achieves the highest visual quality with sharp vessel depiction and low artifacts.

Lastly, we examined 3D reconstruction times to assess suitability for real-time imaging. All methods were executed until convergence for fair comparison. ART with 5 iterations yielded a run time of 1.380 s due to its sequential nature of processing. ADMM variants with 200 iterations had run times of 0.076 (ℓ_1-norm), 0.633 (TV-norm) and 0.648 (ℓ_1-TV-norm). In comparison, PP-MPI achieves run times of 0.078 (single-orientation 2D priors), 0.082 (multi-orientation 2D priors) sec. These run times indicate that PP-MPI can be adopted for real-time reconstruction of a $19 \times 19 \times 19$ volume at 12 frames/sec. In contrast, a variant based on 3D priors could only run at 6 frames/sec.

5 Discussion

In this study, we introduced the first PP approach for MPI reconstruction for improved flexibility in coping with variations in system properties that inevitably alter the forward imaging model as captured by the system matrix. PP-MPI leverages a task-agnostic prior trained on an MPI-friendly dataset for denoising. This deep denoiser is then embedded into a fast model-based reconstruction implemented via an ADMM algorithm. Demonstrations were performed against state-of-the-art MPI reconstruction methods based on hand-crafted regularizers. Quantitative and qualitative improvements were achieved with PP-MPI on both simulated and experimental data. The performance and computational efficiency of PP-MPI render it a promising candidate for real-world MPI applications.

References

1. Ahmad, R., et al.: Plug-and-play methods for magnetic resonance imaging: Using denoisers for image recovery. IEEE Signal Process. Mag. **37**(1), 105–116 (2020). https://doi.org/10.1109/MSP.2019.2949470
2. Bathke, C., Kluth, T., Brandt, C., Maass, P.: Improved image reconstruction in magnetic particle imaging using structural a priori information. Int. J. Magn. Part. Imaging. **3**, 1703015 (2017). https://journal.iwmpi.org/index.php/iwmpi/article/view/64
3. Chae, B.G.: Neural network image reconstruction for magnetic particle imaging (2017). https://doi.org/10.48550/ARXIV.1709.07560, https://arxiv.org/abs/1709.0756
4. Deng, L.: The MNIST database of handwritten digit images for machine learning research. IEEE Signal Process. Mag. **29**(6), 141–142 (2012)
5. Dittmer, S., Kluth, T., Baguer, D.O., Maass, P.: A deep prior approach to magnetic particle imaging. In: Deeba, F., Johnson, P., Würfl, T., Ye, J.C. (eds.) MLMIR 2020. LNCS, vol. 12450, pp. 113–122. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-61598-7_11

6. Güngör, A., et al.: TranSMS: Transformers for super-resolution calibration in magnetic particle imaging. IEEE Trans. Med. Imaging. 1 (2022). https://doi.org/10.1109/TMI.2022.3189693
7. Hatsuda, T., Shimizu, S., Tsuchiya, H., Takagi, T., Noguchi, T., Ishihara, Y.: A basic study of an image reconstruction method using neural networks for magnetic particle imaging. In: 2015 5th International Workshop on Magnetic Particle Imaging (IWMPI), p. 1 (2015). https://doi.org/10.1109/IWMPI.2015.7107046
8. Ilbey, S., et al.: Comparison of system-matrix-based and projection-based reconstructions for field free line magnetic particle imaging. Int. J. Magn. Part. Imaging. **3**, 1703022 (2017). https://doi.org/10.18416/IJMPI.2017.1703022, https://journal.iwmpi.org/index.php/iwmpi/article/view/81
9. Kluth, T., Jin, B.: Enhanced reconstruction in magnetic particle imaging by whitening and randomized SVD approximation. Phys. Med. Biol. **64**(12), 125026 (2019). https://doi.org/10.1088/1361-6560/ab1a4f
10. Knopp, T., et al.: Weighted iterative reconstruction for magnetic particle imaging. Phys. Med. Biol. **55**(6), 1577–1589 (2010). https://doi.org/10.1088/0031-9155/55/6/003
11. Knopp, T., Szwargulski, P., Griese, F., Graser, M.: OpenMPIData: An initiative for freely accessible magnetic particle imaging data. Data Brief **28**, 104971 (2020). https://doi.org/10.1016/j.dib.2019.104971, https://www.sciencedirect.com/science/article/pii/S2352340919313265
12. Koch, P., et al.: Neural network for reconstruction of MPI images. In: 9th International Workshop on Magnetic Particle Imaging, pp. 39–40 (2019)
13. Li, J., Li, J., Xie, Z., Zou, J.: Plug-and-play ADMM for MRI reconstruction with convex nonconvex sparse regularization. IEEE Access **9**, 148315–148324 (2021). https://doi.org/10.1109/ACCESS.2021.3124600
14. McCarthy, J.R., Weissleder, R.: Multifunctional magnetic nanoparticles for targeted imaging and therapy. Adv. Drug Deliv. Rev. **60**(11), 1241–1251 (2008)
15. Zhang, X., Le, T.A., Yoon, J.: Development of a real time imaging-based guidance system of magnetic nanoparticles for targeted drug delivery. J. Magn. Magn. Mater. **427**, 345–351 (2017)
16. Zhang, Y., Tian, Y., Kong, Y., Zhong, B., Fu, Y.: Residual dense network for image restoration. IEEE Trans. Pattern Anal. Mach. Intell. **43**(7), 2480–2495 (2021). https://doi.org/10.1109/tpami.2020.2968521
17. Zheng, B., et al.: Magnetic particle imaging tracks the long-term fate of in vivo neural cell implants with high image contrast. Sci. Rep. **5**(1), 1–9 (2015)