

Deep Unrolled Networks in Representation Space Applied to MRI Reconstruction

Efe Ilıcak, Baris Imre, Chloé Najac, Ruben van den Broek, Beatrice Lena,
Andrew Webb, Marius Staring

Department of Radiology, Leiden University Medical Center, Leiden, Netherlands
`e.ilicak@lumc.nl`

Abstract. Deep unrolled networks (DUNs) integrate physical forward models with learned regularization in cascaded network architectures, achieving exceptional performance in inverse problems while maintaining interpretability. While most DUNs operate in the object domain (e.g., image space), recent variants explored representation spaces for improved information flow. However, these methods rely on heuristic methods for data consistency (DC), sacrificing fidelity with measurements.

In this work, we introduce DUNE (Deep Unrolled Networks in rEpresentation space), a framework that maintains exact adherence to physical measurements while operating in learned representation spaces. By deriving the DC gradient via the chain rule and implementing it through the Vector-Jacobian Product (VJP), we enable exact backpropagation of measurement residuals into the representation space. This formulation supports diverse architectural backbones, including pre-trained encoders to guide the iterative process.

We assess DUNE against state-of-the-art baselines on accelerated MRI reconstruction tasks, demonstrating that exact VJP-based gradients yield superior reconstruction quality and structural fidelity across both single-channel portable low-field and multi-channel high-field MRI acquisitions.

Keywords: Unrolled Networks · Feature Space · Latent Space.

1 Introduction

Inverse problems are central to medical imaging, requiring reconstruction of unknown signals from incomplete or noisy measurements governed by known physical processes [5]. The ill-posed nature of these problems necessitates regularization to incorporate prior knowledge and ensure stable reconstructions [5].

Recently, DUNs have emerged as a powerful paradigm for solving inverse problems, combining the interpretability of model-based iterative algorithms with the representational power of deep learning [1,17]. By mapping iterative optimization algorithms into cascaded network architectures that embed physical forward models [19], DUNs provide enhanced interpretability with superior performance. Their effectiveness has been demonstrated across diverse modalities including magnetic resonance imaging (MRI) [1,17], low-dose computed tomography [20], ultrasound imaging [15,4], and magnetic particle imaging [7]. Yet,

traditional DUNs operate directly in the object domain (e.g., image or measurement space), where iterative updates compress high-dimensional learned features back into the original signal space at each iteration, creating an information bottleneck that limits representational capacity [21].

To alleviate this bottleneck, recent works have explored DUNs in learned representation spaces, including compressed latent spaces [14] and expanded feature spaces [21,6,10]. However, these approaches rely on heuristic approximations for enforcing data consistency, effectively decoupling the optimization from the underlying physics and sacrificing measurement fidelity.

We introduce DUNE (Deep Unrolled Networks in rEpresentation space), which maintains exact adherence to physical measurements while operating in learned representation spaces. Instead of heuristic encoder projections, we derive the exact gradient via the chain rule and implement it using the Vector-Jacobian Product (VJP). This formulation decouples the encoder from gradient computation, enabling diverse architectural choices including pre-trained models.

Our contributions are: (1) derivation of exact data consistency gradients in representation space via VJP, replacing heuristic approximations; (2) encoder decoupling, enabling the use of pre-trained architectures; (3) validation of our approach on challenging scenarios including accelerated single-channel portable low-field MRI (0.047T) and multi-channel clinical MRI (1.5T-3T) acquisitions.

2 Background

2.1 Problem Formulation

We consider recovering an unknown image $\mathbf{x} \in \mathbb{C}^N$ from noisy measurements $\mathbf{y} \in \mathbb{C}^M$ governed by:

$$\mathbf{y} = \mathcal{A}(\mathbf{x}) + \mathbf{n}, \quad (1)$$

where $\mathcal{A} : \mathbb{C}^N \rightarrow \mathbb{C}^M$ is the forward operator and $\mathbf{n}$ is additive noise. This ill-posed inverse problem is typically solved via regularized optimization:

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathcal{A}(\mathbf{x}) - \mathbf{y}\|_2^2 + \lambda \mathcal{R}(\mathbf{x}), \quad (2)$$

where $\mathcal{R}(\mathbf{x})$ is a regularizer (e.g., total variation) that incorporates priors [3].

2.2 Standard Deep Unrolled Networks

This optimization problem can be solved iteratively via gradient descent:

$$\mathbf{x}^{(i+1)} = \mathbf{x}^{(i)} - \eta^{(i)} \left[\mathcal{A}^H(\mathcal{A}(\mathbf{x}^{(i)}) - \mathbf{y}) + \nabla \mathcal{R}(\mathbf{x}^{(i)}) \right], \quad (3)$$

where $\eta^{(i)}$ is the step size and $\mathcal{A}^H$ is the adjoint operator.

Alternatively, modern unrolled networks replace the hand-crafted $\mathcal{R}$ with a learned regularization $\mathcal{R}_\theta$ [1,17]:

$$\mathbf{x}^{(i+1)} = \mathbf{x}^{(i)} - \lambda^{(i)} \mathcal{A}^H(\mathcal{A}\mathbf{x}^{(i)} - \mathbf{y}) - \mathcal{R}_\theta(\mathbf{x}^{(i)}), \quad (4)$$




Fig. 1. Overview of DUN architectures. (a) Standard DUNs use learned regularization $\mathcal{R}_\theta$, but compress features back to object space at each iteration. (b) Heuristic representation space DUNs unroll in representation space, but approximate the data consistency (DC) gradients using encoder ($\mathcal{E}$) projections. (c) DUNE unrolls in representation space with the DC gradients computed via Vector-Jacobian Product (VJP) through decoder ($\mathcal{D}$), decoupling $\mathcal{E}$ (used only for initialization) from iterative updates.

where $\lambda^{(i)}$ is the data consistency weight.

While effective, this formulation compresses learned features back to the object domain (e.g., two channels for real and imaginary components in MRI reconstruction) at each iteration, creating an information bottleneck [21,6]. Importantly, hybrid-domain variants that operate in both image and measurement domains (e.g., dual image & k-space networks [22]) suffer from the same limitation, as they still revert to the object domain for computing data consistency updates.

2.3 Deep Unrolled Networks in Representation Space

In this work, we reformulate the inverse problem to perform optimization directly in a learned representation domain. The objective function in this case becomes:

$$\mathbf{f}^* = \arg \min_{\mathbf{f}} \underbrace{\frac{1}{2} \|\mathcal{A}(\mathcal{D}(\mathbf{f})) - \mathbf{y}\|_2^2}_{\mathcal{L}_{\text{DC}}(\mathbf{f})} + \lambda \mathcal{R}(\mathbf{f}), \quad (5)$$

where $\mathbf{f} \in \mathbb{C}^Q$ is the representation variable, which may be a compressed latent space ($Q < N$) or an expanded feature space ($Q > N$). The mapping is defined by an encoder $\mathcal{E} : \mathbb{C}^N \rightarrow \mathbb{C}^Q$ and a decoder $\mathcal{D} : \mathbb{C}^Q \rightarrow \mathbb{C}^N$, such that $\mathbf{f} = \mathcal{E}(\mathbf{x})$

and $\mathbf{x} \approx \mathcal{D}(\mathbf{f})$. The $\mathcal{L}_{\text{DC}}(\mathbf{f})$ enforces data consistency (DC) with measurements $\mathbf{y}$, and $\mathcal{R}(\mathbf{f})$ represents regularization in the representation domain.

By applying the chain rule to $\mathcal{L}_{\text{DC}}(\mathbf{f})$, we obtain:

$$\nabla_{\mathbf{f}} \mathcal{L}_{\text{DC}} = \left(\frac{\partial \mathcal{D}}{\partial \mathbf{f}} \right)^H \nabla_{\mathbf{x}} \left(\frac{1}{2} \|\mathcal{A}(\mathbf{x}) - \mathbf{y}\|_2^2 \right) \Big|_{\mathbf{x}=\mathcal{D}(\mathbf{f})} = J_{\mathcal{D}}(\mathbf{f})^H \mathbf{g}_x, \quad (6)$$

where $\mathbf{g}_x := \mathcal{A}^H(\mathcal{A}(\mathcal{D}(\mathbf{f})) - \mathbf{y})$ is the residual gradient in the object space, and $J_{\mathcal{D}}(\mathbf{f})$ is the Jacobian of the decoder. Thus, the representation space update is driven by the standard object space DC gradient mapped into the feature space through the adjoint of the decoder’s Jacobian.

Heuristic Approximation in DUNs Recent approaches have explored unrolling in both lower-dimensional latent spaces [14] and higher-dimensional feature spaces [21,6,10] to improve reconstruction quality while maintaining interpretability. However, these methods rely on heuristic gradient approximations for DC, using $\mathcal{E}$ to project residuals back into the representation space:

$$\mathbf{f}^{(i+1)} = \mathbf{f}^{(i)} - \eta^{(i)} \mathcal{E} \left(\mathcal{A}^H(\mathcal{A}(\mathcal{D}(\mathbf{f}^{(i)})) - \mathbf{y}) \right) - \mathcal{R}_{\theta}(\mathbf{f}^{(i)}). \quad (7)$$

However, this approach imposes a dual functional requirement on $\mathcal{E}$. While the exact gradient descent direction in representation space is defined by $J_{\mathcal{D}}(\mathbf{f})^H \mathbf{g}_x$, the heuristic update (Eq. 7) substitutes this term with $\mathcal{E}(\mathbf{g}_x)$. Therefore, $\mathcal{E}$ must simultaneously serve as a feature extractor, and as a proxy for $J_{\mathcal{D}}(\mathbf{f})^H$. Some methods [10] attempt to improve this approximation by regularizing the encoder-decoder pair toward invertibility, aiming to make $\mathcal{E} \approx \mathcal{D}^{-1}$. Nonetheless, because feature encoding and adjoint mapping are fundamentally distinct operations, the heuristic approach forces a compromise between representational accuracy and physical fidelity.

Exact Gradient via VJP To avoid the limitations of heuristic projections, DUNE computes the exact representation space gradient via the Vector–Jacobian Product (VJP). By leveraging standard autograd engines, the term $J_{\mathcal{D}}(\mathbf{f})^H \mathbf{g}_x$ is computed efficiently by backpropagating $\mathbf{g}_x$ through $\mathcal{D}$. This yields the DUNE update rule:

$$\mathbf{f}^{(i+1)} = \mathbf{f}^{(i)} - \eta^{(i)} \underbrace{J_{\mathcal{D}}(\mathbf{f}^{(i)})^H \mathbf{g}_x}_{\text{VJP}} - \mathcal{R}_{\theta}(\mathbf{f}^{(i)}). \quad (8)$$

Encoder Decoupling Note that in Eq. 8, the DC gradient depends exclusively on $\mathcal{D}$ and $\mathcal{A}$, while $\mathcal{E}$ serves solely for initialization ($\mathbf{f}^{(0)} = \mathcal{E}(\mathbf{x}^{(0)})$). This structural decoupling removes the need for $\mathcal{E}$ to serve as a gradient projector or an inverse. This allows the integration of high-capacity, pre-trained models to guide the reconstruction process without compromising exact data consistency updates. This is especially valuable in settings where training data is limited.

3 Experiments

DUNE was evaluated on a prototype single-channel portable low-field and multi-channel clinical high-field MRI datasets. To showcase structural decoupling, we employed a frozen pre-trained $\mathcal{E}$ for low-field reconstructions, and a trainable $\mathcal{E}$ for the high-field setting. All models were optimized for 100 epochs using Huber loss and the AdamW optimizer on an NVIDIA RTX 4000 Ada GPU. Random undersampling masks were generated per batch at each epoch with random monotonic contrast augmentation for robustness. Reconstruction quality was quantified using PSNR, SSIM, and HaarPSI metrics calculated on foreground regions [11]. Statistical significance was assessed using the Wilcoxon signed-rank test between DUNE and competing methods, with * denoting $p < 0.01$.

3.1 Portable Low-Field MRI Reconstruction

Portable MRI systems operating at low-field strengths ($B_0 < 0.1\text{T}$) using permanent magnet arrays offer accessible point-of-care imaging [2]. However, compared to clinical systems, they face severe challenges including several orders of magnitude lower signal-to-noise ratios, and limited measurement redundancy from single-channel receiver coils [16]. While DUNs have recently shown promise for portable MRI reconstruction [18,9], the scarcity of large-scale low-field training data remains a critical bottleneck in training task-specific networks.

DUNE- f -pre: To demonstrate the flexibility of DUNE in such limited-data settings, we propose a feature space model that leverages a frozen pre-trained encoder (DUNE- f -pre). To this end, we employ the stem and first stage of ConvNeXt-Tiny model trained on ImageNet-1k for image classification task [13], as the encoder $\mathcal{E}$. The decoder $\mathcal{D}$ uses 4×4 transposed convolutions to reverse the encoder’s spatial downsampling, followed by two pre-activation residual blocks with GELU activation. For the regularization network $\mathcal{R}_\theta$, we utilize three 3D squeeze-and-excitation residual blocks with LeakyReLU activation. The model is unrolled for 3 iterations with a constant channel dimension of 96.

Portable MRI Experiments We conducted an extensive ablation study comparing DUNE- f -pre against: (i) image domain variant; (ii) feature space variant with heuristic gradient updates; (iii) feature space variant with joint reconstruction and autoencoder loss ($\mathcal{L}_{total} = \mathcal{L}_{recon} + \mathcal{L}_{AE}$). These experiments isolate the benefits of representation space reconstruction, exact VJP gradients, and the impact of enforcing $\mathcal{E} \approx \mathcal{D}^{-1}$.

We also compared DUNE- f -pre against state-of-the-art baselines: a 2D iterative shrinkage-thresholding model (ISTA-NET, 10 iterations) [18], and a 3D dual-domain model (DUN-DD, 5 iterations) with parallel Fourier- and image-domain branches that are fused via an attention U-Net [8].

All models were trained on emulated low-field MRI data [9] from high-field fastMRI brain volumes [12], using 400 volumes for training, and 80 for validation. Pseudo-low-field measurements were undersampled at $R = 2$. The trained




Fig. 2. Representative T2-w reconstructions obtained with a prototype 0.047T portable MRI scanner. Top: ground truth ($R = 1$) and reconstructions at $R = 2$ with zoomed regions of interest; bottom: sampling mask and corresponding error maps.

model and baselines were then tested on 3D T2-w acquisitions from 9 volunteers obtained with a prototype 0.047T portable MRI system [16] at $R = 2$, without any additional training. Volunteer measurements with the prototype system were conducted in accordance with approval from the local ethics committee.

3.2 Clinical High-Field MRI Reconstruction

We assessed DUNE on clinical high-field MRI reconstruction using 2D fastMRI knee data [12], with 1170 slices for training, 388 for validation, and 776 for testing. Multi-coil acquisitions were undersampled at $R = 5$, and all models employed 12 cascades. To evaluate DUNE’s flexibility across different representational domains, we implemented two distinct configurations with learnable $\mathcal{E}$: (i) feature space DUNE (DUNE- f), (ii) latent space DUNE (DUNE- z).

DUNE- f : We designed DUNE- f to closely follow Feature-VarNet [6] for controlled comparison. Following coil sensitivity estimation and combination, images are projected to and from an 18-channel feature space using single 5×5 convolutional layers for both $\mathcal{E}$ and $\mathcal{D}$. This model employs a shared encoder-decoder pair with non-shared $\mathcal{R}_\theta$.

DUNE- z : We further propose DUNE- z as a proof-of-concept to demonstrate unrolling in compressed latent spaces. Multi-coil measurements are projected to and from a 12-channel latent representation using single 1×1 convolutional layers as $\mathcal{E}$ and $\mathcal{D}$. DC is enforced via the forward model with coil sensitivities, while regularization operates in the latent space. This model employs a shared encoder-decoder pair with a shared $\mathcal{R}_\theta$.

Table 1. Reconstruction performance of DUN variants sharing the same $\mathcal{R}_\theta$ architecture, evaluated at $R = 2$ on T2-w measurements from a prototype 0.047T portable MRI. Best results in **bold**.

Ablation Variants	Update Rule	PSNR (dB)	SSIM (%)	HaarPSI (%)
Image Domain	Eq. 4	26.59±1.46*	84.55±1.58*	66.54±2.77*
Heuristic- f -pre	Eq. 7	25.45±1.19*	83.32±2.32*	69.30±2.89*
DUNE- f -pre (+ $\mathcal{L}_{AE}$)	Eq. 8	26.62±1.35*	85.28±2.20*	70.28±2.39*
DUNE- f -pre	Eq. 8	27.25±1.48	86.41±1.89	72.36±2.80

Table 2. Reconstruction performance of competing methods on T2-w acquisitions from the prototype 0.047T portable MRI at $R = 2$, including inference time per volume, and trainable model size. Best results in **bold**.

Method	PSNR (dB)	SSIM (%)	HaarPSI (%)	Time (s)	Size (M)
ISTA-Net	25.26±0.95*	79.96±1.86*	60.99±3.08*	0.60±0.05	0.20
DUN-DD	26.04±1.36*	83.12±1.80*	65.01±3.07*	4.43±1.28	3.53
DUNE- f -pre	27.25±1.48	86.41±1.89	72.36±2.80	1.95±0.20	2.00

Clinical MRI Experiments We compared DUNE against established object-, feature-, and hybrid feature-image domain models: E2E-VarNet [19], Feature-VarNet [6], and FI-VarNet [6]. For fair comparison, all models utilize a 4-level UNet with 18 base channels for $\mathcal{R}_\theta$, except FI-VarNet which combines transformers with UNet. Coil sensitivity maps are estimated using a dedicated UNet [19].

4 Results

4.1 Portable Low-Field MRI Reconstruction

Table 1 presents ablation results on prototype 0.047T portable MRI data, isolating the three key aspects of our approach. DUNE- f -pre significantly outperforms both the image domain and the heuristic gradient variant, validating the benefits of feature space reconstruction and exact VJP-based gradients, respectively. Unlike heuristic updates that force the encoder to double as both a feature projector and a proxy Jacobian, exact VJP decouples representation learning from data consistency optimization. Critically, adding an autoencoder loss to enforce $\mathcal{E} \approx \mathcal{D}^{-1}$ degrades performance, confirming that exact gradients eliminate the need for this constraint, which unnecessarily restricts model flexibility.

Table 2 and Figure 2 compare DUNE- f -pre against baseline methods on T2-w images ($R = 2$, prototype 0.047T scanner). DUNE- f -pre achieves superior reconstruction quality across all metrics with better preservation of anatomical structures, demonstrating effective transfer of pre-trained features to portable low-field MRI. We note, however, that these gains are partially intertwined with specific architecture configurations, warranting further investigations.




Fig. 3. Representative reconstructions at $R = 5$ from fastMRI knee data. Top: ground truth and reconstructions; middle: zoomed regions of interest; bottom: sampling mask and corresponding error maps.

Table 3. Quantitative comparisons on the fastMRI knee data at $R = 5$, including reconstruction metrics, inference time per slice, and model size. Best results in **bold**.

DUN Model	PSNR (dB)	SSIM (%)	HaarPSI (%)	Time (s)	Size (M)
E2E-VarNet [19]	30.27±2.84*	72.64±9.63*	82.16±4.94*	0.09±0.01	29.94
Feature-VarNet [6]	30.56±2.97*	73.44±9.86*	83.54±5.41*	0.08±0.01	30.00
FI-VarNet [6]	30.56±2.92*	73.34±9.77*	83.40±5.21*	0.13±0.01	61.84
DUNE- z	29.37±2.69*	69.69±9.88*	78.79±4.95*	0.08±0.01	2.94
DUNE- f	30.64±3.04	73.62±9.90	83.90±5.38	0.10±0.01	29.98

4.2 Clinical High-Field MRI Reconstruction

Table 3 and Figure 3 present results on fastMRI knee data at $R = 5$. DUNE- f achieves superior performance across all metrics, outperforming object space (E2E-VarNet), feature space (Feature-VarNet), and hybrid (FI-VarNet) baselines. DUNE- z demonstrates competitive performance with substantially reduced model size, illustrating framework flexibility. Although VJP requires additional memory for autograd graphs, inference times are comparable across methods, confirming negligible test-time overhead.

5 Conclusion

We introduced DUNE, a deep unrolled network framework for operating in learned representation spaces with exact physical fidelity. By computing DC gradients via VJP, we replace heuristic encoder approximations while decoupling the encoder from iterative updates. This structural property enables flexible encoder design, demonstrated using frozen pre-trained encoders for low-field MRI where training data is scarce, and trainable encoders for high-field MRI, with strong performance in both scenarios.

Evaluations across various scenarios show that DUNE consistently outperforms baselines, while ablation studies confirm that enforcing encoder-decoder

invertibility degrades performance. DUNE- f achieves superior reconstruction quality across all baselines, while DUNE- z offers competitive performance with substantially reduced parameters. While demonstrated on MRI, the approach extends naturally to other inverse problems with known forward models.

Acknowledgments. This study was supported by the European Innovation Council grant 101136407.

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

References

1. Aggarwal, H.K., Mani, M.P., Jacob, M.: MoDL: Model-Based Deep Learning Architecture for Inverse Problems. *IEEE Transactions on Medical Imaging* **38**(2), 394–405 (2019). <https://doi.org/10.1109/TMI.2018.2865356>
2. Ayde, R., Vornehm, M., Zhao, Y., Knoll, F., Wu, E.X., Sarraçanie, M.: MRI at low field: A review of software solutions for improving SNR. *NMR in Biomedicine* **38**(1), e5268 (2025). <https://doi.org/https://doi.org/10.1002/nbm.5268>
3. Çukur, T., Dar, S.U., Nezhad, V.A., Jun, Y., Kim, T.H., Fujita, S., Bilgic, B.: A Tutorial on MRI Reconstruction: From Modern Methods to Clinical Implications. *IEEE Transactions on Biomedical Engineering* pp. 1–20 (2025). <https://doi.org/10.1109/TBME.2025.3617575>
4. Fan, Y., Wang, H., Gemmeke, H., Hopp, T., Hesser, J.: Model-data-driven image reconstruction with neural networks for ultrasound computed tomography breast imaging. *Neurocomputing* **467**, 10–21 (2022). <https://doi.org/https://doi.org/10.1016/j.neucom.2021.09.035>
5. Fessler, J.A.: Model-Based Image Reconstruction for MRI. *IEEE Signal Processing Magazine* **27**(4), 81–89 (2010). <https://doi.org/10.1109/MSP.2010.936726>
6. Giannakopoulos, I.I., Muckley, M.J., Kim, J., Breen, M., Johnson, P.M., Lui, Y.W., Lattanzi, R.: Accelerated MRI reconstructions via variational network and feature domain learning. *Scientific Reports* **14**, 10991 (5 2024). <https://doi.org/10.1038/s41598-024-59705-0>
7. Güngör, A., Askin, B., Soydan, D.A., Top, C.B., Saritas, E.U., Çukur, T.: DEQ-MPI: A Deep Equilibrium Reconstruction With Learned Consistency for Magnetic Particle Imaging. *IEEE Transactions on Medical Imaging* **43**(1), 321–334 (2024). <https://doi.org/10.1109/TMI.2023.3300704>
8. Ilıcak, E., Imre, B., Najac, C., van den Broek, R., Lena, B., Webb, A., Staring, M.: Physics-Guided Dual-Domain Network with Attention-Based Fusion for Portable MRI Reconstruction. In: 2026 IEEE 23rd International Symposium on Biomedical Imaging (ISBI). pp. 1–4 (2026). <https://doi.org/10.1109/ISBI61048.2026.11515406>
9. Ilıcak, E., Rao, C., Najac, C., Lena, B., Imre, B., Galve, F., Alonso, J., Webb, A., Staring, M.: Physics-Informed Deep Unrolled Network for Portable MR Image Reconstruction (2025), <https://arxiv.org/abs/2509.11790>
10. Jiang, J., He, Z., Quan, Y., Wu, J., Zheng, J.: PGIUN: Physics-Guided Implicit Unrolling Network for Accelerated MRI. *IEEE Transactions on Computational Imaging* **10**, 1055–1068 (2024). <https://doi.org/10.1109/TCI.2024.3422840>

11. Kastrulin, S., Zakirov, J., Pezzotti, N., Dylov, D.V.: Image Quality Assessment for Magnetic Resonance Imaging. *IEEE Access* **11**, 14154–14168 (2023). <https://doi.org/10.1109/ACCESS.2023.3243466>
12. Knoll, F., Zbontar, J., Sriram, A., Muckley, M.J., Bruno, M., Defazio, A., Parente, M., Geras, K.J., Katsnelson, J., Chandarana, H., Zhang, Z., Drozdalv, M., Romero, A., Rabbat, M., Vincent, P., Pinkerton, J., Wang, D., Yakubova, N., Owens, E., Zitnick, C.L., Recht, M.P., Sodickson, D.K., Lui, Y.W.: fastMRI: A Publicly Available Raw k-Space and DICOM Dataset of Knee Images for Accelerated MR Image Reconstruction Using Machine Learning. *Radiology: Artificial Intelligence* **2**(1), e190007 (2020). <https://doi.org/10.1148/ryai.2020190007>
13. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A ConvNet for the 2020s. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11966–11976 (2022). <https://doi.org/10.1109/CVPR52688.2022.01167>
14. Lu, Y., Xie, X., Wang, S., Liu, Q.: Latent-k-space of refinement diffusion model for accelerated MRI reconstruction. *Biomedical Physics & Engineering Express* **11**(5), 055012 (2025). <https://doi.org/10.1088/2057-1976/adf3b4>
15. Mamistvalov, A., Eldar, Y.C.: Deep Unfolded Recovery of Sub-Nyquist Sampled Ultrasound Images. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **68**(12), 3484–3496 (2021). <https://doi.org/10.1109/TUFFC.2021.3093507>
16. O'Reilly, T., Teeuwisse, W.M., de Gans, D., Koolstra, K., Webb, A.G.: In vivo 3D brain and extremity MRI at 50 mT using a permanent magnet Halbach array. *Magnetic Resonance in Medicine* **85**(1), 495–505 (2021). <https://doi.org/https://doi.org/10.1002/mrm.28396>
17. Schlemper, J., Salehi, S.S.M., Kundu, P., Lazarus, C., Dyvorne, H., Rueckert, D., Sofka, M.: Nonuniform variational network: Deep learning for accelerated nonuniform mr image reconstruction. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*. pp. 57–64. Springer International Publishing, Cham (2019). https://doi.org/10.1007/978-3-030-32248-9_7
18. Shimron, E., Shan, S., Grover, J., Koonjoo, N., Shen, S., Boele, T., Sorby-Adams, A.J., Kirsch, J.E., Rosen, M.S., Waddington, D.E.J.: Accelerating Low-field MRI: From Compressed Sensing to Deep Learning Reconstruction with CNNs and Transformers (2025), <https://arxiv.org/abs/2411.06704>
19. Sriram, A., Zbontar, J., Murrell, T., Defazio, A., Zitnick, C.L., Yakubova, N., Knoll, F., Johnson, P.: End-to-end variational networks for accelerated mri reconstruction. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*. pp. 64–73. Springer International Publishing, Cham (2020). https://doi.org/10.1007/978-3-030-59713-9_7
20. Xia, W., Shan, H., Wang, G., Zhang, Y.: Physics-/Model-Based and Data-Driven Methods for Low-Dose Computed Tomography: A survey. *IEEE Signal Processing Magazine* **40**(2), 89–100 (2023). <https://doi.org/10.1109/MSP.2022.3204407>
21. Zhang, J., Zhang, Z., Xie, J., Zhang, Y.: High-Throughput Deep Unfolding Network for Compressive Sensing MRI. *IEEE Journal of Selected Topics in Signal Processing* **16**, 750–761 (6 2022). <https://doi.org/10.1109/JSTSP.2022.3170227>
22. Zhou, B., Zhou, S.K.: DuDoRNet: Learning a Dual-Domain Recurrent Network for Fast MRI Reconstruction With Deep T1 Prior. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4272–4281 (2020). <https://doi.org/10.1109/CVPR42600.2020.00433>