# A Tutorial on MRI Reconstruction: From Modern Methods to Clinical Implications

Tolga Çukur\*, Salman UH. Dar, Valiyeh A. Nezhad, Yohan Jun, Tae Hyung Kim, Shohei Fujita, Berkin Bilgic\*

Abstract - MRI is an indispensable clinical tool, offering a rich variety of tissue contrasts to support broad diagnostic and research applications. Protocols can incorporate multiple structural, functional, diffusion, spectroscopic, or relaxometry sequences to provide complementary information for differential diagnosis, and to capture multidimensional insights into tissue structure and composition. However, these capabilities come at the cost of prolonged scan times, which reduce patient throughput, increase susceptibility to motion artifacts, and may require trade-offs in image quality or diagnostic scope. Over the last two decades, advances in image reconstruction algorithms-alongside improvements in hardware and pulse sequence design—have made it possible to accelerate acquisitions while preserving diagnostic quality. Central to this progress is the ability to incorporate prior information to regularize the solutions to the reconstruction problem. In this tutorial, we overview the basics of MRI reconstruction and highlight state-ofthe-art approaches, beginning with classical methods that rely on explicit hand-crafted priors, and then turning to deep learning methods that leverage a combination of learned and crafted priors to further push the performance envelope. We also explore the translational aspects and eventual clinical implications of these methods. We conclude by discussing future directions to address remaining challenges in MRI reconstruction. The tutorial is accompanied by a Python toolbox (https://github.com/ tutorial-MRI-recon/tutorial) to demonstrate select methods discussed in the article.

Index Terms—MRI, accelerated imaging, image reconstruction, deep learning, clinical, translation

#### I. INTRODUCTION

Magnetic Resonance Imaging (MRI) offers unparalleled diversity in contrast mechanisms to examine tissue anatomy, composition, and function [1]. Clinical MRI protocols routinely include sequences that generate anatomical images with

This study was supported in part by TUBA GEBIP 2015 and BAGEP 2017 fellowships, and a TUBITAK 1001 Grant 121E488 awarded to T. Çukur, and in part by R01 EB028797, R03 EB031175, R01 EB032378, U01 EB026996, UG3 EB034875, and P41 EB030006 grants awarded to B. Bilgic. (\*Co-corresponding authors: Tolga Çukur, cukur@ee.bilkent.edu.tr and Berkin Bilgic, berkin@nmr.mgh.harvard.edu).

T. Çukur and V.A. Nezhad are with the Dept. of Electrical-Electronics Engineering and National Magnetic Resonance Research Center (UM-RAM), Bilkent University, Ankara, Turkey, 06800. SUH. Dar is with the Institute for Artificial Intelligence in Cardiovascular Medicine, Dept. of Cardiology, Angiology, Pneumology, Heidelberg University Hospital and Heidelberg Faculty of Medicine, Heidelberg University, Germany. TH. Kim is with Dept. of Computer Engineering, Hongik University, Seoul, Korea. Y. Jun, S. Fujita, and B. Bilgic are with the Athinoula A. Martinos Center for Biomedical Imaging and Dept. of Radiology, Harvard Medical School, Massachusetts, USA.

e.g. T<sub>1</sub>, T<sub>2</sub>, T<sub>2</sub>\* and FLAIR weighting [2], [3]. These are often complemented by diffusion-weighted imaging (DWI), which plays a central role in stroke assessment [4] and cancer imaging [5]. Given access to higher-end hardware, research protocols further expand this repertoire using advanced sequences sensitive to magnetic susceptibility [6], [7], blood oxygen-level-dependent (BOLD) responses [8], [9], high-b-value diffusion [10], perfusion [11], or metabolite concentrations [12]. Yet, despite this versatility, MRI remains limited by characteristically long acquisition times, often forcing undesirable trade-offs between spatial, temporal, or angular resolution, signal-to-noise ratio (SNR), anatomical coverage, and inclusion of specific contrast mechanisms [13], [14].

Over the past two decades, significant progress has been made toward mitigating these limitations by combining developments in scanner hardware and pulse sequences with sophisticated image reconstruction algorithms. Parallel imaging (PI) [15]–[17], leveraging undersampling of k-space acquisitions and high-channel count coil arrays [18], marked a pivotal advance by enabling accelerated scans across a broad range of applications. Yet, because MRI reconstruction from undersampled acquisitions is fundamentally an ill-posed inverse problem, recovering high-fidelity images often requires integration of external knowledge in the form of priors [19]–[21]. Building upon the foundation laid by PI, a first generation of methods introduced hand-crafted priors designed by experts to constrain the solution space for plausible reconstructions [13], [20]. Prominent examples include compressed sensing (CS) methods enforcing sparsity of MR images in linear transform domains [22]–[27] and low-rank methods promoting structured representations in spatiotemporal or spectral dimensions of MRI data [28]–[31]. Although powerful, these approaches are inherently limited by the simplifying assumptions underlying their priors, which may fall short of fully capturing the rather non-trivial distribution of MR images.

More recently, a second generation of methods has emerged that leverages *learned priors*—implicit regularizers that are directly inferred from exemplar data [14], [32], [33]. These priors are typically implemented as mappings through neural network architectures, which perform hierarchical (i.e., multilayered) nonlinear transformations on the data to extract salient features and capture statistical dependencies present in MR images [34]. Whether used standalone or in hybrid formulations with classical methods [35]–[46], deep learning (DL) approaches based on neural networks have driven substantial improvements in image quality and scan acceleration. However, they also introduce potential challenges—including

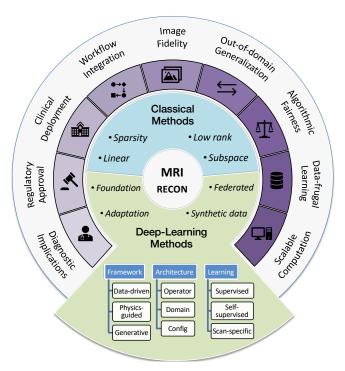


Fig. 1: A landscape of MRI reconstruction methods from classical approaches to modern deep learning are surveyed, alongside discussions on clinical impact and translation, as well as open challenges.

reliance on large-scale, high-quality training data, suboptimal generalizability across scanners or anatomies, and implicit nature of learned priors that can complicate clinical validation—all of which remain active topics of ongoing research.

Given the rapid pace of innovation, staying abreast of the evolving landscape of reconstruction methodology can be challenging. This tutorial is intended as an accessible introduction to MRI reconstruction—bridging foundational concepts, technical advances, and clinical implications (Fig. 1). To provide a conceptual road map, we begin by revisiting core principles of spatial encoding, k-space acquisition, and the inverse problem that underlies reconstruction. We then survey both classical methods based on hand-crafted priors, and DL approaches based on learned priors. We pose a tripartite categorization to group DL methods by (i) the overarching framework (e.g., data-driven vs. physics-guided vs. generative), (ii) architectural elements (e.g., convolutional vs. attentional operator, image vs. k-space domain, single- vs. multi-contrast configuration), and (iii) learning paradigms (e.g., population-level vs. scanspecific). We give special attention to the clinical implications and translation of these reconstruction tools, including their potential to reduce scan times, improve image quality, and expand access to advanced imaging protocols. We conclude the tutorial with a discussion of key open challenges and future opportunities for continued innovation in the field.

# II. FOUNDATIONS OF MR IMAGE RECONSTRUCTION

# A. Generation of Tissue Contrast

MRI relies on magnetization of spins in nuclei, most commonly hydrogen, under a large static field  $B_0$ . During imaging, the net magnetic moment in the longitudinal direction of  $B_0$ 

is tipped to the transverse plane via a radiofrequency (RF) pulse. Magnetization then recovers to its equilibrium state with tissue-specific relaxation time constants of  $T_1$  in the longitudinal axis, and  $T_2$  in the transverse plane. Hence, the transverse magnetization,  $m(\mathbf{r})$  with  $\mathbf{r}=(x,y,z)$  denoting spatial location, depends on the proton density  $\rho$ , relaxation time constants, and the amplitude and timing of RF excitations and magnetic field gradients in the pulse sequence. For common spin-echo (SE) sequences, the signal equation for a single isochromat can be expressed as:

$$m(\mathbf{r}) = K\rho(\mathbf{r})[1 - e^{-\mathrm{TR/T_1(\mathbf{r})}}]e^{-\mathrm{TE/T_2(\mathbf{r})}},$$
 (1)

where TR is the repetition time, TE is the echo time and K is a constant. Manipulating TR and TE can lead to selective weightings of  $T_1$  and  $T_2$ . For instance, moderate TR and short TE can produce  $T_1$ -weighted contrast, which is useful in e.g. imaging fatty tissues with relatively low  $T_1$  or white-, graymatter distinction in brain imaging:

$$m(\mathbf{r}) \approx K\rho(\mathbf{r})[1 - e^{-\text{TR/T}_1(\mathbf{r})}].$$
 (2)

Long TR and moderate TE can produce  $T_2$ -weighted contrast, which can assist in e.g. detecting inflammation and edema:

$$m(\mathbf{r}) \approx K \rho(\mathbf{r}) e^{-\text{TE/T}_2(\mathbf{r})}.$$
 (3)

Meanwhile, long TR and short TE can produce PD-weighted contrast, which can be useful in e.g. assessment of joints:

$$m(\mathbf{r}) \approx K \rho(\mathbf{r}).$$
 (4)

While SE sequences refocus magnetic field inhomogeneities through a refocusing pulse, gradient-echo (GRE) sequences lacking this mechanism are sensitive to additional signal decay governed by  $T_2^*$ , where  $\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'}$  with  $T_2'$  denoting the portion of signal decay that can be recovered by a refocusing pulse, attributable to field inhomogeneities across a voxel.  $T_2^*$ -weighted contrast is more sensitive to local field variations, which can be beneficial for functional and susceptibility imaging.

Beyond manipulating TE and TR to gain sensitivity to  $T_1$  and  $T_2$  differences, diffusion of water molecules can also be leveraged to provide a unique contrast. Because moving spins experience different magnetic fields under the influence of a matched pair of gradients surrounding a refocusing pulse, they cannot be fully refocused at TE, thus attenuating the spinecho signal under diffusion as [47],

$$m(\mathbf{r}) = K\rho(\mathbf{r})[1 - e^{-\text{TR/T}_1(\mathbf{r})}]e^{-\text{TE/T}_2(\mathbf{r})}e^{-bD(\mathbf{r})}, \quad (5)$$

where  $D(\mathbf{r})$  is the diffusion coefficient governing the attenuation, and  $b = (\gamma G \Delta)^2 (\Delta - \delta/3)$  depends on the amplitude (G), duration  $(\delta)$  and time separation  $(\Delta)$  of the gradients. By collecting multiple images with different b values and diffusion directions, clinicians can assess the integrity of neural pathways. Note that MRI offers a diverse arsenal of contrasts, which go well beyond these common sequences (see [1] for further examples), to investigate the structure and function of tissues, making it an indispensable diagnostic tool.

# B. Spatial Encoding in k-Space

Following an RF excitation, the pulse sequence performs spatial encoding to be able to resolve different spatial locations. Specifically, MRI data are acquired in *k-space* by

encoding spatial frequencies via magnetic field gradients that modulate the MR signal phase (Fig. 2a). For a gradient waveform  $\mathbf{G}(t) = (G_x, G_y, G_z)$  that can be independently controlled across spatial axes, the shift in the phase of the MR signal is given as:

$$\Delta \phi = -2\pi \mathbf{k}(t) \cdot \mathbf{r}$$
; such that  $\mathbf{k}(t) = \frac{\gamma}{2\pi} \int \mathbf{G}(t) dt$ . (6)

In (6),  $\mathbf{k}(t) = (k_x, k_y, k_z)$  denotes the evolution of the k-space trajectory across time, and  $\gamma$  denotes the gyromagnetic ratio. The phase modulation from gradients causes the received signal s(t) to reflect the Fourier transform of the underlying transverse magnetization  $m(\mathbf{r}, t)$ :

$$s(t)=M(\mathbf{k})=\int_{\mathbb{R}^3}m(\mathbf{r},t)e^{-j2\pi\mathbf{k}(t)\cdot\mathbf{r}}d\mathbf{r}, \tag{7}$$
 Assuming that the magnetization is quasi-stationary across the

Assuming that the magnetization is quasi-stationary across the k-space trajectory (i.e.,  $m(\mathbf{r}, t) \approx m(\mathbf{r})$ ), its spatial distribution is hence related to  $M(\mathbf{k})$  via an inverse Fourier transformation:

$$m(\mathbf{r}) = \int_{\mathbb{R}^3} M(\mathbf{k}) e^{j2\pi \mathbf{k} \cdot \mathbf{r}} d\mathbf{k}.$$
 (8)

Conventional data acquisition is performed discretely across a finite-extent Cartesian k-space grid, with sampling periods  $\Delta \mathbf{k} = \Delta(k_x, k_y, k_z)$ , and covered frequency ranges  $\mathbf{W}_{\mathbf{k}} = W_{(k_x,k_y,k_z)}$  across axes:

$$\tilde{M}(\mathbf{k}) = M(\mathbf{k}) \cdot \text{III}(\mathbf{k} \otimes \Delta \mathbf{k}) \cdot \text{rect}(\mathbf{k} \otimes \mathbf{W}_{\mathbf{k}}),$$
 (9)

where  $\oslash$  is Hadamard division, III is the Shah function, and rect is the rectangle function. Based on (8), an image  $\tilde{m}(\mathbf{r})$  can be obtained as the inverse discrete Fourier transform of  $\tilde{M}(\mathbf{k})$ , related to the actual magnetization as (Fig. 2b):

$$\tilde{m}(\mathbf{r}) = m(\mathbf{r}) \circledast \operatorname{III}(\Delta \mathbf{k} \odot \mathbf{r})$$

where  $\circledast$  is convolution, and  $\odot$  is Hadamard product. Convolution with III(·) causes  $\tilde{m}(\mathbf{r})$  to contain replicas of  $m(\mathbf{r})$ , yielding an aliasing-free field-of-view (FOV) of size  $1/\Delta \mathbf{k}$  when the Nyquist sampling criterion is met, i.e., the spatial extent of the object  $\mathbf{W}_{\mathbf{r}} = W_{(x,y,z)}$  does not exceed  $1/\Delta \mathbf{k}$ . Convolution with  $\mathrm{sinc}(\cdot)$  terms sets the spatial resolution as  $\Delta \mathbf{r} \approx 1/2\mathbf{W}_{\mathbf{k}}$ . Note that (10) describes 3D Fourier encoding in x-y-z dimensions, though many MRI sequences can also use 2D encoding in x-y with slice-selective excitation along z [1]. For notational convenience, we denote the MR image obtained through discrete k-space sampling of the magnetization as m in the remainder of the manuscript.

# C. Reconstruction: An Inverse Problem

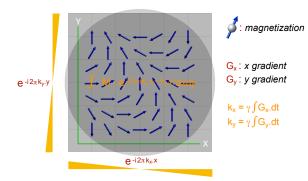
In an MRI acquisition with an acceleration factor of R, a  $1/R^{\text{th}}$  subset of samples on the k-space grid are collected, as defined via a binary sampling pattern  $\mathcal{P}(\mathbf{k})$ . Furthermore, multi-element coil arrays are used that induce spatial sensitivity profiles  $C(\mathbf{r})$  that multiply the magnetization distribution. These influences can be captured via an imaging operator A, linking the magnetization distribution to acquired data  $d(\mathbf{k})$ :

$$d(\mathbf{k}) = \mathcal{P}(\mathbf{k})\mathcal{F}\{C(\mathbf{r})m(\mathbf{r})\},\tag{11}$$

$$d(\mathbf{k}) = A \, m(\mathbf{r}). \tag{12}$$

Image reconstruction, i.e., recovering  $m(\mathbf{r})$  from  $d(\mathbf{k})$ , hence can be cast as an inverse problem aiming to minimize the

#### (a) Spatial Encoding



# (b) k-space vs. Image Domain ky y ARE IF OV Fourier Transformation

Fig. 2: (a) Illustration of spatial encoding via magnetic field gradients to localize signal origin. This encoding mechanism enables the systematic sampling of the tissue magnetization in the Fourier domain (i.e., k-space). (b) Visualization of the relationship between k-space where raw MRI data are acquired, and the image domain depicting the spatial distribution of tissue magnetization, linked through Fourier transformation.

discrepancy between acquired and reconstructed data, derived via a complex Euclidean norm since noise in MRI data is bivariate Gaussian (Fig. 3):

$$m^* = \underset{m}{\operatorname{argmin}} ||d - Am||^2, \tag{13}$$

whose analytical solution is given by  $m^* = A^\dagger d = (A^H A)^\dagger A^H d$ , with  $A^H$  denoting the Hermitian adjoint of A and  $\dagger$  denoting Moore-Penrose pseudoinverse. Since the Fourier transform is unitary (i.e.,  $\mathcal{F}^{-1} = \mathcal{F}^H$ ), A becomes relatively well-conditioned when the k-space grid is fully sampled (i.e.,  $\mathcal{P} = 1$ ). Under Cartesian k-space sampling, the solution can then be efficiently approximated by applying a coil-by-coil inverse Fast Fourier Transform (FFT) to the acquired data, followed by coil combination using sensitivity maps, yielding  $m^* \approx A^H d$  [18]. In accelerated MRI, however, undersampling may render the problem ill-posed, necessitating the use of regularization to incorporate additional prior information:

$$m^* = \underset{m}{\operatorname{argmin}} ||d - Am||^2 + \lambda \Psi(m), \tag{14}$$

where the first term enforces consistency to acquired data, and  $\Psi(m)$  weighted by  $\lambda$  enforces prior information. In classical methods,  $\Psi(m)$  has an explicit formulation based on hand-crafted regularizers that impose assumptions on the data. In DL methods,  $\Psi(m)$  has an implicit formulation based on data-driven regularizers learned from a training dataset.

#### III. CLASSICAL RECONSTRUCTION METHODS

Classical methods integrate hand-crafted physical priors (i.e., explicit signal models or assumptions on the data distribution) to regularize the solution of the inverse problem in (14) (see Fig. 4 for example reconstructions). Below, we review major categories, including linear regression/interpolation, sparsity, low rank, and subspace models, each leveraging different structural properties of MRI data.

Linear Regression/Interpolation Methods: PI reduces the FOV deliberately by omitting acquisition of k-space lines, and attempts to resolve aliasing artifacts in image (e.g., SENSE [16]) or k-space domain (e.g., GRAPPA [17]). In imagedomain techniques, coil sensitivity profiles are used explicitly to unalias voxels by solving the linear regression problem in (14). For the basic unregularized case ( $\lambda = 0$ ), this problem admits a closed-form, non-iterative least-squares solution based on the Moore-Penrose pseudoinverse (i.e.,  $m^* = A^{\dagger}d$ ) [16]. While this non-iterative solution is commonly used for the basic case, iterative solvers can provide improved numerical stability and efficiency in higher-dimensional or more complex reconstruction settings. For instance, a gradient descent algorithm<sup>1</sup> can compute  $m^* = \operatorname{argmin} f_{pi}(m) =$  $\operatorname{argmin} ||d - Am||^2$ :

$$m_{t+1} = m_t - \alpha \cdot \partial f_{pi}(m_t), \tag{15}$$

where  $\alpha$  is a step size that needs to be chosen appropriately, and  $m_t$  and  $m_{t+1}$  denote the image estimates at steps t and t+1. Evaluating the gradient above leads to the update rule,

$$m_{t+1} = m_t - \alpha \cdot A^H (Am_t - d), \tag{16}$$

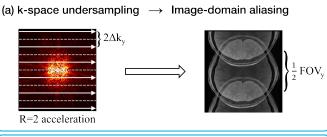
which is run until a convergence criterion is satisfied.

Unlike SENSE, GRAPPA uses inter-coil relationships to interpolate missing k-space data. This interpolation is performed across coils and within a small k-space neighborhood:

$$M_j(k_x, k_y) = \sum_{i} \sum_{u \in U} \sum_{v \in V} g_j(i, \mu, v) M_i(k_x + u\Delta k_x, k_y + v\Delta k_y)$$

with  $i = 1, ..., N_c$  where  $N_c$  is the number of coils. Here,  $M_j(k_x, k_y)$  is the target missing k-space sample in the  $j^{th}$  coil and  $g_i$  is the k-space "kernel" used for interpolating the data for the  $j^{th}$  coil. This kernel computes a linear combination of the acquired points  $M_i(k_x + u\Delta k_x, k_y + v\Delta k_y)$  inside a small neighborhood in k-space, denoted with U and V, across all coils. The kernel weights  $q_i$  are estimated from low-resolution, fully-sampled calibration data, which is collected via a prescan or formed from the central portion of undersampled data.

Note that the estimation of coil sensitivity profiles C from calibration data is critical to construct an accurate estimate of the imaging operator in image-domain or analogously the interpolation kernel in k-space techniques. While heuristic approaches such as polynomial fitting can be used to estimate smoothly-varying profiles, structured methods have become prominent that estimate them as eigenvectors in the null space of a calibration matrix formed using calibration data (e.g.,



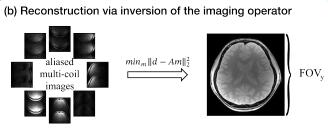


Fig. 3: (a) Undersampling in k-space leads to aliasing artifacts in the image domain following inverse Fourier transformation. (b) Multi-coil aliased images are processed to invert the imaging operator A, which encapsulates the effects of the sampling pattern, Fourier encoding, and coil sensitivities. A reliable inversion can produce an artifact-reduced, coil-combined image that approximates the fully-sampled reference.

ESPIRiT [48], PISCO [49]). Beyond such a priori estimation strategies, it is also possible to formulate nonlinear optimization problems that jointly estimate coil sensitivity profiles and the underlying image (e.g., JSENSE [50], IRGNM [51]). These nonlinear approaches can yield more accurate reconstructions, albeit at the cost of increased computational complexity.

Sparsity Methods:  $\Psi(m)$  in (14) can make use of a sparsity prior as popularized by CS methods that recover images from compressible measurements [22]. These priors are often formulated as  $\Psi(m) = ||\Phi m||_1$ , where  $\Phi$  performs linear transformation to a domain (e.g. wavelet or gradient domain), whose coefficients are assumed to be compressible. Such  $M_j(k_x,k_y) = \sum_i \sum_{u \in U} \sum_{v \in V} g_j(i,\mu,v) \\ M_i(k_x + u\Delta k_x, k_y + v\Delta k_y), \text{sparsity priors can be incorporated into PI to perform a PICS reconstruction}^2 \ [23], \ [52], \ [53]:$ 

$$m^* = \underset{m}{\operatorname{argmin}} f_{pics}(m) = \underset{m}{\operatorname{argmin}} ||d - Am||^2 + \lambda ||\Phi m||_1.$$
(18)

When  $\Phi$  is orthogonal or a tight frame, (18) can be solved via a proximal gradient descent algorithm [54],

$$v_t = m_t - \alpha A^H (Am_t - d), \tag{19}$$

$$m_{t+1} = \Phi^H \mathcal{S}_{\alpha\lambda}(\Phi v_t), \tag{20}$$

where  $\alpha$  denotes step size,  $v_t$  denotes an intermediate iterate, and  $S_{\alpha\lambda}(\cdot)$  denotes element-wise soft-thresholding with a threshold of  $\alpha\lambda$ . Note that, unlike unregularized PI methods, PICS algorithms can benefit from the noise-like aliasing interference elicited by randomized undersampling patterns [25].

Low-Rank Methods: Low-rank methods typically extract local patches from k-space data and stack them as rows of a structured matrix [28]-[31], [55]. Linear dependencies across coils, limited spatial support, phase smoothness, and transform-domain sparsity properties confer this matrix with

<sup>&</sup>lt;sup>1</sup>An implementation of SENSE is available at https://github.com/ tutorial-MRI-recon/tutorial/tree/main/basic\_sense\_

<sup>&</sup>lt;sup>2</sup>An implementation of PICS is available at https://github.com/ tutorial-MRI-recon/tutorial/tree/main/basic\_sense\_

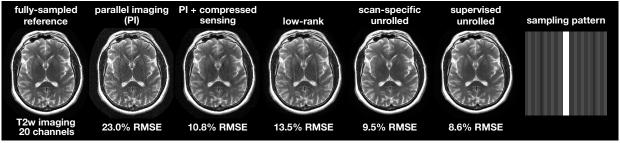


Fig. 4: Reconstruction of an axial T<sub>2</sub>-weighted brain image from undersampled k-space data using various methods. From left to right: fully-sampled reference, parallel imaging (PI), parallel imaging with compressed sensing (PICS), low-rank reconstruction, scan-specific unrolled network, supervised unrolled network, and the sampling pattern. Root-mean-squared error (RMSE) between reconstructed and reference images is reported below each method. PI and low-rank reconstructions show higher noise and residual aliasing, while PICS reduces noise but introduces block-like artifacts. DL methods achieve lower RMSE values, with visually reduced artifacts and improved sharpness relative to classical methods. The supervised model provides the strongest artifact suppression, albeit with a few subtle synthetic features, whereas the scan-specific model offers higher fidelity at the expense of residual artifacts.

low-rankness, which can be exploited as prior information [28], [30], [56]–[59]. Specifically, given d, local neighborhood of k-space samples are arranged into  $\mathcal{H}(k)$ , forming a Hankel or Toeplitz matrix that exhibits low-rank characteristics,

$$\operatorname{rank}(\mathcal{H}(k)) \le p$$
, such that  $d = Am$ , (21)

However, directly enforcing this rank constraint results in a non-convex, NP-hard problem, which is challenging to solve. To make the problem tractable, a typical approach is to convert the rank constraint into a minimization problem<sup>3</sup>,

$$m^* = \underset{m}{\operatorname{argmin}} \|d - Am\|^2 + \lambda \Psi(\mathcal{H}(k)), \tag{22}$$

where  $\Psi(\cdot)$  is a surrogate rank penalty. Commonly used formulations for  $\Psi(\cdot)$  include [30], [56]–[58],

$$\Psi(\mathbf{X}) = \min_{\mathbf{T}, \, \text{rank}(\mathbf{T}) \le p} \|\mathbf{X} - \mathbf{T}\|_F^2 = \sum_{i > p} \sigma_i^2$$
 (23)

or [31], [60],

$$\Psi(\mathbf{X}) = \|\mathbf{X}\|_* = \sum_i \sigma_i \tag{24}$$

where  $\sigma_i$  denotes the singular values of matrix **X**, and p is a user-defined rank parameter.

Subspace Methods: Quantitative MRI (qMRI) reconstructions can leverage the knowledge of spin evolution by creating a low-rank subspace  $\Gamma$  for MR signals, where  $\Gamma$  is formed with singular value decomposition (SVD) of a dictionary generated through Bloch signal simulations or learned from sample image patches [61]. Subspace coefficients s in  $\Gamma$  can be estimated from undersampled k-space data, such that  $\Gamma s$  is an accurate approximation to the measured signal evolution [62]. This estimation can be performed as,

$$s^* = \operatorname*{argmin}_s ||d - A\Gamma s||^2 + \lambda \cdot \Psi(s) \tag{25}$$
 where the regularizer  $\Psi(s)$  could impose sparsity or low-

where the regularizer  $\Psi(s)$  could impose sparsity or low-rank constraints on the subspace coefficients. Having obtained the coefficients, images with various different contrasts can then be synthesized. A similar approach is adopted in MR Fingerprinting (MRF) [63], where Bloch simulations lend themselves to creation of a dictionary that explains the behavior of transverse signal as the flip angle and TRs are evolving throughout the acquisition. Rather than using such a dictionary only at the signal matching step after image reconstruction, it can be incorporated into the reconstruction process itself,

so that dictionary coefficients that approximate the acquired data in a parsimonious manner can be directly estimated [64]. This subspace approach can help mitigate relaxation-related blurring during the qMRI echo train, and boost SNR and mitigate artifacts in MRF [65].

Note that the subspace formulation can be viewed as a linear approximation of the underlying MR signal evolution. Rather than projecting signals into a fixed low-rank basis, an alternative is to explicitly enforce the nonlinear signal equations during reconstruction, enabling direct estimation of quantitative parameters (e.g., relaxation times) from undersampled data [66]. Such nonlinear formulations can improve parameter accuracy by adhering more closely to MR physics, but they lead to nonlinear inverse problems that are often computationally more demanding to solve. These nonlinear formulations are applicable to steady-state sequences that admit a closed-form expression for their signal evolution, such as those in (1-3).

#### IV. DEEP-LEARNING RECONSTRUCTION METHODS

DL methods learn data-driven priors from exemplar MRI datasets to help solve (14). To structure discussions on DL, we dissect methods in the literature along three fundamental axes: modeling framework, network architecture, and learning paradigm. This tripartite view reflects the core design choices that govern how DL reconstruction methods incorporate domain knowledge, compose their networks, and learn from data.

# A. Modeling Frameworks

To avoid overlapping terminology and blurred boundaries that may obscure complementary strengths and limitations, here we categorize current DL methods under three distinct frameworks: data-driven, physics-guided, and generative (Fig. 5; see Fig. 4 for example reconstructions). This taxonomy is grounded in fundamental distinctions regarding how DL methods leverage domain knowledge, i.e., how they incorporate data-driven and physical priors, to solve the inverse problem. Accordingly, data-driven and physics-guided models refer to conditional networks<sup>4</sup> trained to predict MR images

<sup>&</sup>lt;sup>3</sup>An implementation of LORAKS is available at https://github.com/tutorial-MRI-recon/tutorial/tree/main/low\_rank.

 $<sup>^4</sup>$ Conditional networks generate outputs based on a measured input and a learned mapping, effectively modeling the conditional distribution  $P(\text{output} \mid \text{input})$ .

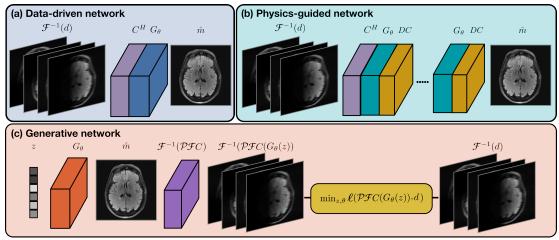


Fig. 5: (a) A data-driven network  $G_{\theta}$  reconstructs an MR image  $\hat{m}$  from a linear transformation of undersampled multi-coil data d. In this example, the linear transformation is a zero-filled Fourier reconstruction  $m_u = A^H d$ , obtained by inverse Fourier transformation of d followed by coil combination based on estimated profiles C. The reconstruction  $\hat{m}$  is obtained via a forward pass of  $m_u$  through network layers. (b) A physics-guided network reconstructs an MR image from a linear transformation of undersampled data. In this example, the zero-filled Fourier reconstruction  $m_u$  is forward passed iteratively through interleaved network layers  $G_{\theta}$  and data-consistency (DC) layers. In DC layers, the input coil-combined image is first projected onto a multi-coil image based on C, data consistency is enforced in k-space, and the resultant multi-coil image is coil-combined. (c) A generative network  $G_{\theta}$  trained to capture the distribution of high-quality MR images is combined with the imaging operator to reconstruct an MR image at inference. In this example,  $G_{\theta}$  synthesizes a random MR image from a latent representation z—this may be a random code (GANs), Gaussian noise (diffusion), or a latent code produced by an encoder.  $G_{\theta}$  and z can be jointly optimized to ensure that reconstructed k-space data are consistent with acquired data d based on the imaging operator A.

that would resemble reconstructions of fully-sampled data, using as input a linear transformation of undersampled data, and differing in whether they exploit physical priors. By contrast, generative models refer to unconditional networks<sup>5</sup> trained to synthesize random MR images that are not tied to undersampled data, integrating physical priors only during inference. This classification partly diverges from the broader machine learning convention, where 'generative models' can encompass any architecture capable of producing stochastic images. In MRI reconstruction, however, we argue that the critical distinction lies not in output stochasticity, but in how models fundamentally approach the inverse problem. Conditional networks—even those with stochastic components like adversarial or variational architectures—receive undersampled data as explicit input and are thus more appropriately classified within our data-driven or physics-guided categories based on their treatment of physical priors.

A.1 Data-Driven Models: Instead of relying on predefined priors, these methods learn data-driven priors directly from available datasets [38], [46], [67], [68]. In a common formulation, a network is trained to map from undersampled to fully-sampled data, using a training dataset  $\left\{\left(d^{\operatorname{tr}(j)}, m^{\operatorname{tr}(j)}\right)\right\}_{j=1}^{N_{\operatorname{tr}}}$  where j is the sample index,  $N_{\operatorname{tr}}$  is the number of training samples,  $d^{\operatorname{tr}(j)}$  denotes acquired data for the jth sample, and  $m^{\operatorname{tr}(j)}$  is the ground-truth image. Note that the network input/output can be expressed in image domain, k-space or both, with image-domain networks often using the least-squares solution  $m_u = A^H d$  (i.e., zero-filled Fourier reconstruction) as input, and trying to recover the ground-truth image m. The training objective for a data-driven (dd) network can then be

expressed as:

$$\mathcal{L}_{dd}(\theta) = \mathbb{E}_{(d^{\text{tr}}, m^{\text{tr}})} \left[ \ell \left( G_{\theta}(m_u^{\text{tr}}), m^{\text{tr}} \right) \right]$$
 (26)

where  $\mathbb E$  denotes expectation,  $G_{\theta}$  is the network parameterized by  $\theta$ , and  $\ell(\cdot,\cdot)$  denotes a loss function that reflects the discrepancy between network-predicted and ground-truth reconstructions. Note that this discrepancy can be measured in image domain as depicted in (26), in k-space, or both. Furthermore, while the  $\ell_2$ -norm is a common choice for comparing network and ground-truth reconstructions—mirroring its role in quantifying data consistency—alternative losses are often employed to promote desired image attributes. Prominent examples include  $\ell_1$ , hybrid  $\ell_1$ - $\ell_2$ , perceptual, adversarial, or diffusion losses [69].

During inference, the trained network with optimized parameters  $\theta^*$  processes undersampled data d to predict  $\hat{m}$  through a forward pass across the network architecture:

$$\hat{m} = G_{\theta^*}(m_u). \tag{27}$$

Learning an end-to-end mapping from undersampled to fully-sampled data, purely data-driven models are computationally efficient and can outperform traditional methods. Yet, they may suffer performance losses due to the lack of an explicit mechanism to enforce fidelity to the imaging operator or acquired data [70]. As remedy, some approaches embed data-driven networks—typically pretrained for denoising or artifact suppression—within iterative optimization algorithms [71], [72]. This keeps the network agnostic to the imaging operator during training, while imposing physical priors externally at inference—a middle ground between purely data-driven and tightly integrated physics-guided designs.

A.2 Physics-Guided Models: To benefit from knowledge of acquired k-space locations and coil sensitivities, physics-guided models typically inject data-consistency blocks based on the imaging operator within the network architecture [14], [37], [40], [73]–[80]. A common motivation for this design

 $<sup>^5</sup>$ Unconditional networks generate outputs without an explicit measured input, often leveraging random latent variables, and model the marginal distribution  $P({\rm output})$ .

is a composite objective that balances physics-driven data consistency with data-driven regularization:

$$m^* = \operatorname*{argmin}_m \|d-Am\|^2 + \lambda \, \ell(G_\theta(m), m), \qquad (28)$$
 where  $G_\theta$  denotes the network module. In practice, unrolled

where  $G_{\theta}$  denotes the network module. In practice, unrolled models do not directly solve (28), but instead alternate a learned prior block with a quadratic data-consistency update as described in (29-30). This iterative structure resembles fixed-point iterations, and has close connections to proximal splitting methods [81]–[84]:

$$z^{(i+1)} = G_{\theta}(m^{(i)}), \tag{29}$$

$$m^{(i+1)} = \underset{m}{\operatorname{argmin}} \|d - Am\|^2 + \lambda \ell(z^{(i+1)}, m), \qquad (30)$$

where I is the total number of iterations, and  $0 \le i < I$ . The training objective for a physics-guided (pg) network can be formulated as:

$$\mathcal{L}_{pg}(\theta) = \mathbb{E}_{(d^{\text{tr}}, m^{\text{tr}})} \left\{ \ell(m^{(I)}, m^{\text{tr}}) \right\}$$
(31)

Image-domain unrolled networks typically set the input as  $m^{(0)} = m_u$ , and given a trained network, the reconstructed image during inference is taken as  $\hat{m} = m^{(I)6}$ .

Traditional methods for quantitative MRI often rely on analytical models that describe how tissue parameters affect the measured signals based on the physics of spin dynamics (e.g., mono-exponential decay in multi-echo spin-echo sequences for  $T_2$  mapping). Recent physics-guided DL methods also integrate such analytical priors to estimate quantitative maps of tissue parameters q from undersampled data [85], [86]:

$$q^* = \underset{q}{\operatorname{argmin}} ||d - Ah(q)||^2 + \lambda \ell(G_{\theta}(q), q),$$
 (32)

where  $h(\cdot)$  is the analytical signal model. By jointly optimizing over parameter and data fidelity, such unrolled networks can mitigate errors that arise from treating image reconstruction and parameter mapping disjointly.

Physics-guided models incorporate the imaging operator to improve reconstruction performance, but they may face generalization issues when tested out-of-domain—e.g. with different undersampling patterns or rates—and become computationally expensive when using a large number of cascades.

A.3 Generative Models: Unlike models subject to task-specific training involving undersampled and fully-sampled data, generative priors agnostic to the imaging operator (i.e., the k-space sampling patterns and coil sensitivity maps) have come forth as a promising alternative to boost generalization by learning the distribution of fully-sampled data alone [41], [87]. For inference, a generative prior that synthesizes MR images can be combined with the imaging operator, typically through a conventional optimization problem. A variety of machine learning approaches from autoencoding to flow-matching models are available for building generative priors [88], yet arguably adversarial and diffusion models have gained particularly widespread adoption.

Earlier works utilized Generative Adversarial Networks (GANs) [89], [90] comprising a generator network  $G_{\theta}$  that synthesizes images from latent variables, and a discriminator

 $^6 An$  implementation of an unrolled model based on hybrid  $\ell_1\text{-}\ell_2$  loss is available at <code>https://github.com/tutorial-MRI-recon/tutorial/tree/main/supervised\_dl.</code>

network  $D_{\theta}$  that distinguishes between actual and synthetic images, trained via the following objective:

$$\mathcal{L}_{GAN}(\theta) = \mathbb{E}_{m \sim m^{tr}} \left[ \log D_{\theta}(m) \right] + \mathbb{E}_{z \sim p_{z}(z)} \left[ \log \left( 1 - D_{\theta}(G_{\theta}(z)) \right) \right], \tag{33}$$

where  $p_z(z)$  denotes the distribution of latent variables. During inference, the latent code and, sometimes, the generator itself are optimized so as to satisfy fidelity with the imaging operator and acquired data:

$$(z,\theta)^* = \operatorname*{argmin}_{z,\theta} \ell(AG_{\theta}(z),d)$$
 (34)

where z denotes a vector of random latent variables, and  $\ell$  is a loss function that measures discrepancy between reconstructed and acquired k-space data (e.g.,  $\ell_2$ -norm,  $\ell_1$ -norm, or hybrid  $\ell_1$ - $\ell_2$  norm). The final reconstruction can be computed via the generator network as:

$$\hat{m} = G_{\theta^*}(z^*) \tag{35}$$

More recent approaches have leveraged diffusion-based priors that operate through two complementary processes: a forward process that progressively adds noise to training images, and a reverse process that learns to denoise images step-by-step [91]–[93]. During training, the forward process systematically corrupts clean images by adding Gaussian noise over T timesteps, while the reverse process trains a neural network  $G_{\theta}$  to predict and remove this added noise. The training objective can be expressed as:

$$\mathcal{L}_{\text{diff}}(\theta) = \mathbb{E}_{t \sim [1,T], m_0^{\text{tr}}, \epsilon} \left[ \| \epsilon - G_{\theta} \left( \sqrt{\bar{\alpha}_t} m_0^{\text{tr}} + \sqrt{1 - \bar{\alpha}_t} \epsilon, t \right) \|^2 \right],$$
(36)

where  $m_0^{\rm tr}$  refers to ground-truth images from the training set,  $\epsilon \sim \mathcal{N}(0,I)$  is Gaussian noise,  $G_{\theta}(\cdot,t)$  is the neural network that predicts the added noise at timestep t, and  $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$  defines the cumulative noise schedule that controls the amount of noise added during the forward process. During inference, only the reverse process is used to reconstruct the MR image corresponding to an undersampled acquisition, starting from a pure Gaussian noise image. To do this, the denoising process is guided between successive steps to ensure that the synthesized images remain faithful to acquired data [94], analogously to an unrolled network operating with T iterations. At each step t, the sampling procedure for a diffusion model can be expressed as:

$$\tilde{m}_{t-1} = G_{\theta}(m_t, t), t \in [T, 0),$$
(37)

$$m_{t-1} = \tilde{m}_{t-1} + A^H(d - A(\tilde{m}_{t-1})),$$
 (38)

where  $\tilde{m}_{t-1}$  is the image at step (t-1) that can be predicted using  $G_{\theta}$ ,  $m_{t-1}$  is the corrected data-consistent version, and  $m_T = \epsilon$  represents the starting point of pure Gaussian noise. The final reconstructed image is taken as  $\hat{m} = m_0$ . Image quality can be further enhanced by adapting the generative prior itself to acquired data, as in (34) [95].

Generative methods decouple the imaging operator from the training of the network to arrive at a formulation that improves generalization. However, they often suffer from prolonged inference times, and determining appropriate stopping criteria for test-time adaptation can be challenging.

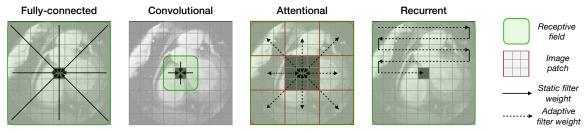


Fig. 6: Receptive fields for fully-connected, convolutional, attentional, and recurrent (e.g., state-space) operators, illustrating differences in filtering primitives. Fully-connected operators perform static global filtering across the image. Convolutional operators perform static local filtering constrained to a compact neighborhood. Attentional operators perform input-adaptive global filtering by computing pair-wise similarities of image tokens (e.g., patches). Recurrent operators perform input-adaptive global filtering by aggregating high-order dependencies across image pixels through sequence modeling.

# B. Network Architectures

Although the modeling framework dictates the overarching network structure for reconstruction models, a range of architectural decisions from operator type to contrast configuration must be made when composing individual layers.

B.1 Operator Type: Each layer applies mathematical operators that capture spatial dependencies in the input-in conjunction with normalization and non-linear activation functionsenabling the extraction of task-relevant features that are propagated through the network (Fig. 6). These operators define distinct filtering primitives, imparting varying degrees of representational power and computational efficiency.

Fully-connected operators establish an explicit relationship between each pair of elements in the input feature map versus the output feature map:

$$m_{\text{out}}[j] = \sum_{i} w[i, j], m_{\text{in}}[i],$$
 (39)

where  $m_{\text{in}} \in \mathbb{C}^M$  denotes the vectorized input map,  $m_{\text{out}} \in \mathbb{C}^M$  is the output map, and  $w \in \mathbb{C}^{M \times M}$  is the static filter weights of the operator. As complex-valued operators are relatively uncommon in DL programming libraries, practical implementations often map real/imaginary or magnitude/phase components across the channel dimension of feature maps and prescribe real-valued operators [44]. Commonly found in implicit generative methods that synthesize images as a function of spatial coordinates [96], fully-connected operators enable modeling of all pairwise interactions between input and output elements [97], yet they are parameter-intensive and often computationally prohibitive for many network designs.

Convolutional operators use a translation-invariant kernel to filter the input feature map:

$$m_{\text{out}}[u, v] = \sum_{i} \sum_{j} w[u - i, v - j] m_{\text{in}}[i, j],$$
 (40)

where  $m_{\text{in, out}} \in \mathbb{C}^{U \times V}$ , and  $w \in \mathbb{C}^{r \times r}$  denotes the static filter weights with  $r \ll U, V$ . Convolutional operators are parameter-efficient due to their localized receptive fields and hence easier to train on compact datasets [40], [98], which have contributed significantly to their widespread use in DL models for MRI reconstruction. However, these benefits come at the cost of a bias toward localized dependencies between neighboring elements, which limits sensitivity to long-range interactions across the input map.

Attentional operators compute interactions across the input feature map  $m_{\rm in} \in \mathbb{C}^{U \times V}$  by partitioning it into  $N = \frac{UV}{P^2}$ 

patches  $\{m_{in}^k\}_{k=1}^N$ , each of size  $P \times P$ . Each patch is flattened and linearly projected into a vector  $z_k \in \mathbb{C}^D$ . The output patch representations  $\{z_j'\}_{j=1}^N$  are derived as [90], [99]:

$$z'_{j} = \sum_{k=1}^{N} \alpha_{j,k}(z_{k}W_{V}), \tag{41}$$

$$\alpha_{j,k} = \frac{\exp\left((z_j W_Q)(z_k W_K)^\top / \sqrt{D}\right)}{\sum_{l=1}^N \exp\left((z_j W_Q)(z_l W_K)^\top / \sqrt{D}\right)}, \quad (42)$$
 where  $W_Q, W_K, W_V \in \mathbb{C}^{D \times D}$  are learnable query, key, and

value projection matrices, and the attentional filter weights  $\alpha$ are input adaptive. The output feature map  $m_{\text{out}} \in \mathbb{C}^{U \times V}$  is reconstructed by reshaping and aggregating the output patches  $\{z_i'\}$ . Attention enables modeling of long-range interactions across the input map, but it incurs higher computational and memory costs compared to convolutional operators, particularly for high-resolution images with many patches [100]-[102].

**Recurrent operators** propagate information sequentially across the spatial dimensions, providing an alternative to attention primitives for capturing global context [103]-[105]. A representative example is the state-space model (SSM) operator that describes the dynamics of hidden states  $h[i] \in$  $\mathbb{C}^N$  in response to an input sequence of pixels  $m_{\rm in}[i] \in \mathbb{C}^M$ indexed by i through a linear system [106]:

$$h[i+1] = A_{\text{ssm}} h[i] + B_{\text{ssm}} m_{\text{in}}[i],$$
 (43)

$$m_{\text{out}}[i] = C_{\text{ssm}} h[i] + D_{\text{ssm}} m_{\text{in}}[i], \tag{44}$$

 $m_{\rm out}[i] = C_{\rm ssm} \, h[i] + D_{\rm ssm} \, m_{\rm in}[i], \tag{44}$  where  $A_{\rm ssm} \in \mathbb{C}^{N \times N}, \; B_{\rm ssm} \in \mathbb{C}^{N \times M}, \; C_{\rm ssm} \in \mathbb{C}^{M \times N},$   $D_{\rm ssm} \in \mathbb{C}^{M \times M}$  are learnable matrices for adaptive filtering. Here, flattening the input feature map  $m_{\rm in}$  into a 1D sequence via a specific pixel ordering across the spatial grid constrains the inherently multidimensional spatial relationships to a single traversal path. Yet, this sequentialization enables SSMs to summarize long-range spatial dependencies at lower computational cost compared to attentional operators [107]-[109]. Ongoing works explore alternative sequentialization approaches (e.g., raster scan, spiral, space-filling curves) to balance fidelity to image structure with efficiency [106], [110].

B.2 Input Domain: Architectures vary in the domain in which they operate on MRI data, with each domain offering distinct advantages for capturing structural, frequency-specific, or semantic features essential to the reconstruction task.

Image-domain networks typically operate on aliased images, obtained via zero-filled inverse Fourier transformation of undersampled data [37]. Within physics-guided or generative frameworks, image-domain feature maps are often propagated between the image and Fourier domains to enforce data consistency based on the imaging operator [40].

**Fourier-domain networks** directly process undersampled k-space data, offering an intrinsic advantage for enforcing data-consistency, as the data remain in the Fourier domain [45], [111], [112]. However, special care is often required when handling k-space inputs, as the energy distribution varies significantly across spatial frequencies.

**Hybrid-domain networks** employ separate branches or interleaved cascades to process both image- and Fourier-domain representations [73], [113], [114]. By combining spatially localized image features with globally distributed spectral features, these architectures can achieve enhanced reconstruction performance compared to single-domain counterparts.

Latent-domain networks abstract away from image or k-space representations by operating in a learned latent space. Such networks encode MRI data onto a lower-dimensional semantic representation, and decode the processed semantic representation back onto MRI data [43], [87], [115], [116]. This strategy enables the more efficient processing of high-dimensional and multi-contrast MRI datasets.

*B.3 Contrast Configuration:* Architectures can also differ in whether they process different tissue contrasts in an MRI protocol independently or jointly. This typically shapes the branching strategies within the network and directly impacts the network's performance.

**Independent reconstruction** refers to the most common formulation adopted in MRI, where each contrast (e.g.,  $T_1$ ,  $T_2$ ) is reconstructed separately using contrast-specific inputs [117]. This approach simplifies modeling, and generally lowers model complexity. However, it also ignores potential redundancies and synergies across contrasts, which may limit reconstruction quality towards higher acceleration factors.

Joint reconstruction aims to reconstruct multiple contrasts acquired from the same subject simultaneously by leveraging shared structural features arising from the underlying common anatomy [118], [119]. Such networks typically receive multicontrast inputs, either as multi-channel tensors or separate inputs fused through cross-modal attention or shared encoding [75], [77], [120], [121]. Joint reconstruction can significantly improve performance, particularly under high accelerations, albeit requires careful registration of contrasts to avoid artifacts due to spatial misalignment and feature leakage [122].

**Prior-conditioned reconstruction** refers to setups where the reconstruction of a target contrast is guided by one or more auxiliary contrasts [123], [124]. This formulation is particularly relevant in settings where one contrast carries a fully-sampled acquisition, and another has a heavily undersampled acquisition. Prior-conditioning mechanisms may include concatenation, attention, or learned cross-contrast priors [125]. Prior-conditioned reconstruction offers similar benefits to joint reconstruction, albeit proper alignment between contrasts is often key to model performance [126]. Note that prior-conditioning does not have to be restricted to distinct contrasts;

it can also be employed to process repeated scans of a given contrast. For instance, patient-specific conditioning can be realized to improve reconstruction performance using prior scans of the same contrast for a subject, even when acquired on different MRI systems (e.g., high-field versus low-field) [127].

# C. Learning Paradigms

The choice of training strategy for DL methods is closely tied to the availability and structure of training data. Depending on whether fully-sampled acquisitions are available to serve as ground truth, or whether a training set exists at all, different paradigms become suitable. Below, we review the three main paradigms commonly adopted in the literature that aim for population-level, scan-specific or hybrid training.

C.1 Population-Level Training: Population-level training aims to learn a reconstruction network from a cohort of training subjects, assuming access to multiple independent acquisitions. Depending on whether paired (fully-sampled and undersampled) data are available, supervised or self-supervised learning might be adopted.

**Supervised learning** assumes access to paired datasets  $\{(m_u^{(i)}, m^{(i)})\}_{i=1}^N$ , where i is sample index, N is the number of samples in the training set,  $m_u$  is the zero-filled Fourier reconstruction of an undersampled acquisition (or its corresponding k-space data), m is the ground-truth image Fourier reconstructed from a fully-sampled acquisition. The reconstruction network is trained to minimize the discrepancy between its prediction  $\hat{m}^{(i)} = G_{\theta}(m_u^{(i)})$  and the ground truth (Fig. 7a), general objective can be expressed as:

$$\mathcal{L}_{\sup}(\theta) = \mathbb{E}_{(m_u, m)} \left[ \ell \left( G_{\theta}(m_u), m \right) \right], \tag{45}$$

where  $\mathbb{E}$  denotes expectation typically estimated via Monte Carlo sampling over training samples. Supervised methods yield high reconstruction accuracy when the distributions of training versus test samples are reasonably well aligned. However, they require large datasets of fully-sampled acquisitions, which are costly to collect in clinical settings [34], [128].

Note that generative networks are typically trained to synthesize MR images derived from fully-sampled data, without considering the imaging operator. At inference time, they are combined with an inverse sampling technique (see Section IV-A.3) to reconstruct undersampled data from test subjects. Formally, this setup aligns with the definition of unsupervised learning, as networks are trained without paired input-output examples, learning solely from the target distribution [90].

Self-supervised learning. With MRI protocols that require significantly high spatial-temporal resolution and coverage, it may be infeasible to perform fully-sampled acquisitions to curate training sets in clinically acceptable scan times [34], [117]. In such scenarios, reconstruction models can be trained solely on undersampled acquisitions by adopting self-supervised learning [129], [130]. A common strategy partitions available measurements into training and supervision sets [72], [76], [131], where acquired k-space indices  $\Omega$  are split into two disjoint sets  $\Theta$  and  $\Lambda$  such that  $\Theta \cup \Lambda = \Omega$ . The network is trained to reconstruct the signal over  $\Lambda$  using predictions based

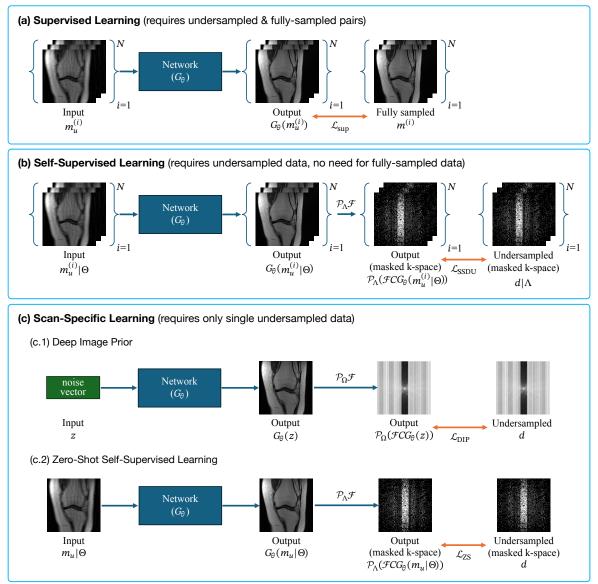


Fig. 7: Different learning paradigms depending on availability and structure of training data. (a) Supervised learning performs population-level training of a reconstruction network using paired sets of undersampled and fully-sampled data, where training loss is expressed using ground-truth images derived from fully-sampled data. (b) Self-supervised learning performs population-level training using only undersampled data, in cases where it is not feasible to collect fully-sampled acquisitions. Training loss is expressed using a masked portion of undersampled data. (c) Scan-specific learning performs test-time training using undersampled data from an individual test subject. Reconstruction can be achieved via either an unconditional network that maps a noise vector z onto an image as in deep image prior (DIP) methods, or a conditional network that maps a zero-filled Fourier reconstruction of undersampled data  $m_u$  onto an image as in zero-shot (ZS) methods. In both cases, training loss is expressed using single-scan undersampled data.

on data from  $\Theta$ , as inspired by masked modeling techniques in computer vision [99] (Fig. 7b):

 $\mathcal{L}_{\text{SSDU}}(\theta) = \mathbb{E}_{m_u} \left[ \ell \left( \mathcal{P}_{\Lambda} \left( \mathcal{F} C G_{\theta}(m_u|_{\Theta}) \right), d|_{\Lambda} \right) \right],$  (46) where  $\mathcal{P}_{\Lambda}$  is a projection onto the supervision set in k-space. This technique enables network training without ever accessing fully-sampled ground truth, making it particularly appealing in real-world clinical scenarios [132], [133].

C.2 Scan-Specific Training: When training data are entirely absent due to limited scanning resources, network models must be learned exclusively on test data during inference [98], [134], [135]. Note that such scan-specific learning can be categorized as a subclass of self-supervised techniques that learn from a single test sample (i.e., a single test subject). A prominent approach uses a randomly-initialized unconditional network  $G_{\theta}$  to map a fixed noise input z onto an image-

domain reconstruction [136]–[139]. The network is optimized to satisfy data consistency in k-space based on the available undersampled acquisition, as inspired by deep image prior (DIP) techniques [140] (Fig. 7c.1):

$$\mathcal{L}_{\text{DIP}}(\theta) = \ell \left( \mathcal{P}_{\Omega} \left( \mathcal{F} C G_{\theta}(z) \right), d \right), \tag{47}$$

where  $\mathcal{P}_{\Omega}$  denotes the sampling mask.

Another approach uses a randomly-initialized conditional network  $G_{\theta}$  to map a linear transformation of a subset  $\Theta$  of undersampled data (e.g., the zero-filled Fourier reconstruction  $m_u|_{\Theta}$  of  $d|_{\Theta}$  onto a reconstruction [135], [141]. In this case, the network is optimized to satisfy consistency to another

subset  $\Lambda$  of undersampled data (Fig. 7c.2):<sup>7</sup>

$$\mathcal{L}_{ZS}(\theta) = \ell\left(\mathcal{P}_{\Lambda}\left(\mathcal{F}CG_{\theta}(m_u|_{\Theta})\right), d|_{\Lambda}\right). \tag{48}$$

Test-time trained networks can adapt to the unique characteristics of each scan, albeit running optimization during inference is computationally expensive. Moreover, when the entire set of acquired data is used during training, the network might suffer from overfitting to noise if not early stopped [142]. A promising strategy to alleviate this issue is to hold out an independent subset of k-space samples as validation data to guide early stopping [141]. For instance, a third disjoint set of k-space data  $\mathcal V$  (i.e.,  $\Theta \cup \Lambda \cup \mathcal V = \Omega$ ) can be reserved to assess the validation performance of the conditional network described in (48).

C.3 Hybrid Training: A network learned on a training dataset collected at a particular imaging site or with a specific protocol might later be deployed in different settings, yielding shifts in the data distribution. These shifts can in turn cause poor generalization in learned networks, necessitating domain alignment. Hybrid training combines the strengths of population-level and scan-specific paradigms to enhance generalization on out-of-distribution test data [43], [87], [90], [143]. In this framework, a network  $G_{\theta_0}$  initially trained on a cohort of training subjects is later adapted on the undersampled test data from a specific subject (Fig. 7c.2):

 $\mathcal{L}_{\text{hybrid}}(\theta) = \ell\left(\mathcal{P}_{\Omega}\left(\mathcal{F}CG_{\theta}(z)\right), d\right)$  with  $\theta \leftarrow \theta_{0}$ , (49) where z is a noise variable or an aliased image,  $\theta$  is initialized from  $\theta_{0}$ . Hybrid training leverages population priors while retaining the adaptability of scan-specific optimization.

#### V. CLINICAL IMPLICATIONS AND TRANSLATION

# A. General Implications

A.1 Improved Image Quality for Visual Assessment: Technical advances in MRI encoding and reconstruction have enabled recovery of high-fidelity images from heavily undersampled data, which is particularly impactful in clinical settings where scan time is limited but the need for diagnostic accuracy is paramount [144]. In essence, improvements in image quality—such as enhanced contrast, reduced noise, and clearer delineation of anatomical structures—enable radiologists to make more accurate and confident interpretations, particularly in complex or borderline cases. For instance, high tissue contrast and structure preservation contribute to improved detection and delineation of pathological findings such as ischemic lesions, demyelinating plaques, or small tumors [145]. The extent of the lesion delineation is also important [146]. In addition, improved consistency in image texture and sharpness across slices and contrasts contributes to a more coherent visual impression [147]. This uniformity reduces cognitive load on the reader, minimizes interpretation time, and decreases intrareader and inter-reader variability [146].

A.2 Enhanced Patient Compliance and Comfort: Accelerated MRI reconstruction techniques offer benefits not only

for image quality but also for patient compliance and comfort [22], [148]. One of the most direct advantages is the reduction of scan duration. Shorter scans can minimize the time patients are required to remain still inside the scanner, thereby reducing physical discomfort, especially for populations such as children, elderly individuals, or critically ill patients [149], [150].

Increased efficiency in MRI scanning schedules improves throughput and access to imaging services. Shortened scan durations also decrease the likelihood of motion-related artifacts, which are one of the most common sources of non-diagnostic or inconclusive scans [151]. In high-demand clinical settings, minimizing the need for repeat acquisitions and completing exam protocols within limited time windows help alleviate scheduling bottlenecks, reduce patient wait times, lower operational costs, and enhance the overall patient experience [152].

These safety and comfort gains, while sometimes considered secondary to diagnostic efficacy, are in fact essential to patient-centered care. They contribute to higher compliance rates, fewer aborted scans, and an overall improvement in the quality of care delivered in the imaging suite.

A.3 Expanded Access to Advanced Imaging Protocols: The integration of advanced MRI reconstruction techniques, particularly those enabling high acceleration factors, has the potential to allow for the incorporation of additional or more complex pulse sequences that would otherwise be excluded due to time constraints. This is particularly significant for multi-parametric, multi-contrast, or dynamic imaging protocols that can provide more comprehensive diagnostic information. For instance, advanced reconstructions can facilitate the routine inclusion of diffusion, perfusion, spectroscopy, or relaxometry-all of which are often omitted in time-limited clinical workflows despite their diagnostic value [153]. In neuroimaging, this may enable broader application of multishell diffusion protocols or resting-state functional MRI, while in body imaging, time-resolved sequences (e.g., GRASP) may become more feasible [154].

From a practical perspective, this expansion can contribute to earlier detection of disease, improved lesion characterization, and more precise treatment planning. Furthermore, it can help bring research-grade imaging techniques into standard clinical practice, thereby narrowing the translational gap between academic innovation and real-world diagnostics [155]. Additionally, advanced protocols supported by fast acquisitions are particularly valuable in longitudinal studies or follow-up examinations, where consistent but fast imaging is critical to monitor disease progression or treatment response. Ultimately, the capacity to deliver richer imaging datasets within standard clinical time slots represents a major step toward personalized and precision medicine [156].

A.4 Facilitation of Data-Demanding Analysis Pipelines: As modern radiology increasingly relies on computational tools for image processing and decision support, the role of image reconstruction extends beyond generating visually interpretable images. It now serves as the foundational step in a data pipeline that feeds into downstream applications such as automated segmentation, radiomics, disease classification, and predictive modeling [157], [158].

<sup>&</sup>lt;sup>7</sup>An implementation of a scan-specific model is available at https://github.com/tutorial-MRI-recon/tutorial/tree/main/scan\_specific\_dl.

	Concept and Basic Research	Clinical Validation	Regulatory Approval	Commercialization and Deployment	Post-market Monitoring and Improvement
Milestones	Identification of unmet clinical needs, initial feasibility testing, preliminary algorithm prototype. Validate the technique in phantoms.	Validate the technique in controlled settings in retrospective patient datasets. Pilot and expanded clinical trials, demonstration of diagnostic efficacy and stability	Submission of FDA/CE regulatory applications, obtaining certification or clearance	Product launch, clinical integration, user training, scalability of deployment	Real-world performance tracking, ongoing quality assurance, iterative improvements and updates
Stakeholders	Imaging scientists, physicists, clinicians, engineers	Radiologists, clinical investigators, technologists, physicists, clinicians, engineers, patients	Regulatory specialists, quality assurance teams, legal teams, regulatory bodies	Industry representatives, hospital administrators, clinical users (radiologists, technologists)	Clinical users, regulatory bodies, R&D teams, professional societies

Fig. 8: Stages of translational development for MRI reconstruction techniques, alongside the critical milestones and involved stakeholders.

In addition, accelerated reconstruction techniques facilitate the acquisition of large volumes of imaging data without exceeding clinically acceptable scan durations. This is especially relevant for 3D or 4D protocols, multi-contrast studies, or high-resolution imaging where large data throughput is needed for training and inference with DL models. For example, in radiomic workflows, the variability in input quality due to undersampled acquisitions can be mitigated by advanced reconstruction algorithms, thereby enhancing the reproducibility and robustness of the extracted features [159], [160]. In this context, the synergy between advanced reconstruction and data-centric analysis tools forms a critical axis for the future of precision radiology [161].

# B. Application-Specific Implications

Although many current clinical MRI protocols employ modest acceleration supported by classical reconstruction methods, e.g. R=2-4 using PI, substantial opportunities remain to further improve scan efficiency via adoption of advanced reconstruction techniques. Benefits of additional acceleration would be particularly evident in longer protocols, where reduced scan times can significantly improve workflow, enhance patient comfort, and lower the risk of motion-related artifacts [162]. The clinical value of MRI reconstruction should therefore be evaluated in relation to exam durations and diagnostic requirements across anatomical regions and applications.

**Neuro MRI.** In neuroimaging, standard clinical protocols are designed to acquire multiple image contrasts, such as T<sub>1</sub>-weighted, T<sub>2</sub>-weighted, FLAIR, diffusion, and susceptibility-weighted imaging, typically taking 20 to 30 minutes [128]. Advanced reconstruction techniques can enable higher acceleration rates to be used, reducing the scan time to as little as 5 to 10 minutes. Protocols that involve high-resolution volumetric imaging across multiple contrasts, such as those used in epilepsy evaluation (usually 30 to 40 minutes [163]), stand to benefit significantly. In addition to improving patient compliance, especially among pediatric or neurologically impaired populations, the time saved can be reinvested to increase spatial resolution or to include additional contrasts, thereby enhancing lesion conspicuity and anatomical detail without prolonging the overall exam duration.

MSK MRI. In musculoskeletal imaging, fine anatomical delineation is required for assessing small structures such as ligaments, tendons, cartilage, or subchondral bone [164]. Due to the requirement for high spatial resolution and multiple oblique imaging planes aligned with complex joint anatomy, musculoskeletal imaging typically takes 15 to 30 minutes [165]. The time saved with higher acceleration using advanced reconstruction techniques can be reinvested to improve SNR or further enhance spatial resolution, thereby facilitating detection of subtle findings such as partial tendon tears, cartilage defects, and marrow lesions.

Abdominal MRI. In abdominal imaging, challenges related to respiratory motion and contrast timing often necessitate extended protocols over multiple breath-holds, lasting 30–45 minutes [27]. Free-breathing acquisitions combined with advanced reconstruction methods can reduce scan times to 10–15 minutes. This can enable more consistent image quality across patients, mitigate motion artifacts, and preserve key contrasts such as lesion enhancement. For prostate exams, high quality T<sub>2</sub>-weighted images in multiple planes and DWI are often required, while contrast enhanced scans are also used in subpopulations [166]. Reconstruction techniques can enable shorter and/or higher-resolution scans, improving lesion detection while streamlining screening and diagnostic workflows.

Cardiac MRI. Used pervasively in evaluation of cardiovascular diseases, cardiac MRI involves detailed assessments across multiple domains, such as cardiac morphology and function, myocardial tissue characterization, perfusion imaging, flow and hemodynamics, or coronary artery evaluation [167]. These assessments are encoding-intensive due to the need for dynamic, multi-phase, and multi-contrast data acquisitions. Additionally, respiratory and cardiac motion often necessitate multiple breath-holds during scanning. The average scan duration ranges from 30 to 60 minutes, depending on protocol complexity (e.g., inclusion of perfusion, late gadolinium enhancement, or mapping sequences) and patient-related factors such as breath-holding capacity and the presence of arrhythmias [168], [169]. Reconstruction techniques can facilitate the use of higher acceleration and reduce scan time to 10-20 minutes, enabling higher temporal resolution or more effective motion correction. These improvements can enhance the visualization of wall motion, myocardial strain, and perfusion dynamics.

Whole-body MRI. In whole-body imaging, increasingly adopted in oncology for metastatic screening and treatment monitoring, protocols often extend to 45–90 minutes [170]. Advanced acquisition/reconstruction, coupled with motion-compensation techniques, can reduce exam times to 15–30 minutes without compromising diagnostic accuracy.

MR-guided Radiotherapy. MR-guided radiotherapy leverages MRI's superior soft-tissue contrast and real-time imaging for precise tumor targeting and on-table adaptive planning without ionizing radiation [171], [172]. Yet, cine MR sequences for gating and intrafraction tracking require high temporal resolution, minimal latency, and robustness to motion [173]. Advanced reconstruction methods can provide low-latency image formation with improved artifact suppression, supporting reliable margin definition in abdominal and tho-racic sites and enhancing treatment precision.

Interventional MRI. Guiding procedures such as biopsies, ablations, and catheter-based therapies [174], interventional MRI requires rapid, high-temporal-resolution visualization of a limited field of view to track instruments and target tissues in real time. Frequent presence of needles, wires, and ablation devices further necessitates robust artifact mitigation [175]. Advanced reconstructions that accelerate imaging with minimal latency while preserving spatial fidelity can support continuous procedural monitoring, faster intra-procedural decision-making, and safer, more effective interventions overall.

**Low-field MRI.** Low-field MRI ( $\lesssim$ 0.5T) offers reduced equipment cost, portability, and improved safety for patients with implants, pediatric populations, or those in resource-limited or bedside settings [176]. However, lower SNR and reduced spatial resolution are common limitations at these field strengths that can compromise diagnostic utility. Advanced reconstruction methods can mitigate these challenges by improving scan efficiency, lowering image noise, and improving spatial resolution.

Others. Finally, emerging applications such as lung, breast, and fetal MRI stand to benefit substantially from advanced reconstruction. These techniques can help address longstanding barriers including low signal, motion artifacts, and time constraints. As such methods continue to mature, they may drive broader clinical adoption and enable more accessible, efficient MRI workflows for screening and point-of-care evaluations.

# C. Stages of Translational Development

The translation of MRI reconstruction technologies into clinical practice generally follows a multi-stage pipeline. Although several definitions of translational research phases have been proposed in the biomedical sciences, especially in genomics, they may not fully capture the specific challenges associated with imaging technologies [177], [178]. As summarized in Fig. 8, the key stages include conceptual development, clinical validation, regulatory approval, commercialization, and post-market monitoring.

In the early conceptual stage, reconstruction methods are

developed and tested on simulated or phantom data to establish feasibility. Subsequent clinical validation involves retrospective evaluation on patient datasets and prospective pilot studies, aiming to demonstrate diagnostic efficacy, reproducibility, and robustness in real-world settings. Most studies leverage image fidelity metrics such as MSE, PSNR, and SSIM, which are widely used in computer vision, to provide a quantitative benchmark of reconstruction performance. While valuable for assessing basic image similarity, these conventional metrics can have limited correspondence to clinical utility and should be interpreted with caution when used as indicators of diagnostic value. Ideally, metrics that better align with human perception of image quality can be incorporated to improve fidelity of metric-based evaluations [179], [180]. More importantly, clinical validation should not rely solely on quality metrics but rather on demonstrated clinical value, that is, whether the reconstruction supports accurate diagnosis and appropriate patient management [181].

A particularly challenging step in this pipeline is regulatory approval. MRI reconstruction techniques, especially those involving machine learning, may be subject to rigorous review by agencies such as the U.S. Food and Drug Administration (FDA) or European CE marking authorities. Demonstrating safety, interpretability, and generalizability is essential.

Following regulatory clearance, commercialization and deployment require integration into clinical workflows and imaging platforms. This involves not only technical compatibility with PACS (picture archiving and communication systems) and scanner software, but also attention to computational demands, user training, and support infrastructure as scalable deployment demands efficient algorithms.

Finally, post-market surveillance is critical to ensure continued safety and efficacy, identify performance issues across diverse populations or institutions, and guide iterative updates [182]. Collaborative efforts between industry, regulatory bodies, and clinical users are essential to support this continuous improvement process. Overall, the successful translation of advanced MRI reconstruction techniques depends on coordinated progress across technical innovation, regulatory navigation, and system-level implementation.

#### D. Clinical Workflow Integration

Despite their diagnostic potential, advanced reconstruction techniques face several challenges in clinical implementation. Each technique has its own strengths and limitations as briefly summarized in Table I, and it is important to understand these characteristics to select the most appropriate method for the task at hand. Beyond implications directly related to the level of acceleration, artifact suppression, and image quality, an important challenge in clinical implementation is the computational demand of advanced reconstruction methods. Integration of these methods into routine workflows inevitably depends on infrastructure readiness. High-performance computing resources, often including GPUs (graphical processing units), are typically required—posing constraints for on-scanner deployment. Off-scanner solutions, such as PACS-integrated servers

TABLE I: Categorization of MRI reconstruction methods and considerations for their clinical translation.

Category	Methods	Key Characteristics	Pros	Cons	Current Trans- lational Stage
Classical Methods	Linear Regression / Interpo- lation	Exploits coil-sensitivity encoding; straightforward implementation; clinically established.	Computationally efficient implementation; clinically accepted; minimal regulatory hurdles.	Limited acceleration (R), residual artifacts/noise towards higher R.	Post-market Monitoring and Improvement
	Sparsity	Exploits transform-domain spar- sity; iterative optimization-based implementation; broad applica- bility across MRI sequences.	Moderate R, strong artifact reduction, clinically validated.	High computational cost due to iterative optimization; requires careful parameter tuning.	Commercialization and Deployment Stage
	Low-Rank	Exploits structured low-rankness; especially beneficial on dynamic MRI data.	Decent artifact suppression; well-suited for data with significant spatial/temporal correlations.	Computationally intensive; relatively complex, non-standard implementation.	Mainly Concept and Basic Re- search
	Subspace	Utilizes Bloch-simulated signal models; beneficial in parameter mapping applications.	Quantitative imaging capability; high flexibility in post-hoc con- trast generation.	Complexity of dictionary calibration; longer development cycles.	Some are FDA/CE approved
Deep- Learning Methods	Data- Driven	Leverages priors learned from exemplar data; reconstruction cast as a single forward pass through the network.	High reconstruction quality; rapid inference; minimal parameter tuning.	Limited out-of-domain general- ization; regulatory complexities due to black-box nature.	Commercialization and Deployment Stage
	Physics- Guided	Combines data- and physics- driven priors; reconstruction cast as iterations through the two types of priors.	Better physical grounding; higher user trust; improved generalization under changes to the imaging operator.	Higher computational burden than data-driven methods; lim- ited structural details in recon- structed images.	Mainly Clinical Validation, some are FDA/CE approved
	Generative	Leverages task-agnostic generative priors to synthesize MR images; enhances structural details.	High visual quality; excellent detail preservation.	Susceptibility to hallucinations; reduced user trust; typically heavy computational burden.	Mainly Concept and Basic Re- search

or secure local clusters, provide a feasible alternative, but may introduce latency and complicate data routing.

Seamless adoption also demands interoperability with hospital IT (information technology) systems and compliance with DICOM (digital imaging and communications in medicine) standards. Deployment must be planned with attention to institutional regulations, particularly in regions with strict data governance policies. These operational factors underscore the importance of aligning algorithm development with practical implementation pathways. Effective integration requires not only technical compatibility, but also organizational adaptation and clinician engagement. Technologists and radiologists must be trained to understand both the capabilities and limitations of novel reconstruction methods. Trust in these systems is reinforced through transparent validation, consistent performance, and the ability to recognize artifacts or failure modes.

# VI. FUTURE DIRECTIONS

Learning-based MRI reconstruction has achieved impressive gains in scan acceleration and image quality in recent years. However, as modern methods transition from research to widespread clinical use, new challenges emerge that go beyond reconstruction performance. Future progress must also prioritize reliability and adaptability across diverse patient populations and clinical settings. Key areas of focus include mitigating the risk of hallucinations and bias, designing learning paradigms that scale with limited or heterogeneous data, and ensuring that methods are computationally efficient and

infrastructure-compatible for real-world deployment.

# A. Enhancing Reliability Against Hallucinations

Advances in MRI reconstruction have enabled the recovery of high-fidelity images from heavily undersampled acquisitions, significantly improving scan efficiency. To suppress aliasing artifacts in such acquisitions, reconstruction methods balance the support of acquired data with that of the imposed prior. As acceleration rates increase and fewer data are collected, reconstructions become more susceptible to hallucinations—manifesting as either loss or leakage of tissue features [183]. Loss occurs when true features cannot be recovered due to aliasing, while leakage refers to the appearance of features not present in the actual anatomy. Since hallucinatory features can closely resemble realistic anatomy, they pose serious risks for diagnostic reliability [184].

Classical methods and some scan-specific DL methods rely on hand-crafted priors, which may lead to visible losses such as over-smoothing or signal attenuation in fine structures—potentially alleviated by tuning down the regularization strength at the expense of weaker artifact suppression. In contrast, non-scan-specific DL methods often learn from population data, risking leakage of features commonly seen in the pool of training subjects but absent in the test subject. A promising direction to mitigate hallucinations is the integration of priors grounded in physical signal models, resulting in hybrid approaches that combine classical and DL methods [185]–[187]. Ensemble approaches that aggregate predictions

from diverse DL methods (e.g., physics-guided and generative networks) might also offer a level of safeguard by blending complementary strengths [94], [188]. Additionally, joint reconstruction strategies for multiple MRI contrasts can help enhance anatomical consistency across sequences, reducing the likelihood of hallucinated features appearing in any single contrast when proper measures are taken to avoid cross-contrast leakages [74], [189]. Incorporating explicit anatomical constraints such as shape-aware priors might further regularize the solution space to favor plausible reconstructions [190].

While hallucinations may not be entirely avoidable, providing uncertainty or hallucination maps alongside reconstructed images [191], [192] or using quantitative metrics that may be sensitive to presence of hallucinatory features [193] can guide interpretation. Moreover, explainability tools—such as layerwise relevance propagation [194], counterfactual generation [195], or network gradients [196]—can help identify which components of the network are responsible for hallucinations. Such retrospective analyses may support decisions about whether clinically relevant features are real or hallucinated.

# B. Mitigating Biases from Data and Methodology

As DL methods become increasingly prevalent in medical imaging, concerns about algorithmic bias have gained growing attention [197]. One major source of bias arises when reconstruction models, often trained on datasets that overrepresent certain populations or anatomical characteristics, systematically underperform on underrepresented groups [198]. DL-based reconstruction models may inadvertently encode population-level priors that favor the majority distributions in the training data—at the expense of generalizability. Unlike hallucinations, which typically manifest as localized errors such as feature leakage or loss, bias often presents more subtly as systematic degradations in SNR, resolution, or contrast. It should be recognized that biases are not exclusive to DL methods—classical reconstructions can also produce characteristic artifacts due to over-emphasis of image features compatible with selected hand-crafted priors, such as block artifacts in total variation regularization. These biases can lead to inequities in image quality, diagnostic accuracy, and clinical decision-making.

Another important source of bias affecting the fidelity of model evaluation arises from shortcomings in methodological practices [184]. Common issues include assessing reconstruction models on MRI data that have been preprocessed in ways that simplify the underlying inverse problem—for example, using magnitude-only, coil-combined, or compressed images during modeling. Such simplifications can lead to overly optimistic performance metrics and poor generalization when models are eventually deployed on real-world MRI data, which are inherently complex and multi-coil.

Addressing bias in MRI reconstruction therefore requires a multifaceted strategy that combines data-centric and model-centric solutions. On the data side, curating diverse and representative datasets across demographics that contain raw, multi-coil, complex MRI data with minimal processing is fundamental. Collaborative learning frameworks offer a promising

avenue to enhance training diversity while preserving patient privacy [116], [199]. On the model side, advances in domain adaptation and model personalization can help mitigate performance degradations under domain shifts [200], [201]. Evaluating reconstruction performance across predefined subgroups—by demographic attributes or clinical conditions—is equally critical to reveal hidden disparities. Fairness-aware evaluation frameworks may provide systematic tools to quantify and address potential biases [202]. Ultimately, transparent reporting of dataset composition, preprocessing steps, and model development/evaluation settings will be essential for equitable adoption of DL-based MRI reconstruction.

# C. Learning Under Data Constraints

By learning task-specific data representations, DL architectures can surpass classical methods in capturing complex signal structure during MRI reconstruction. However, many supervised DL methods require large training datasets consisting of paired undersampled and fully-sampled acquisitions, and are typically trained for specific anatomies, contrasts, or undersampling patterns [37]–[40]. Unfortunately, even for a given anatomy, MRI scanners operate with a number of distinct protocols that vary widely in contrast, resolution, and other acquisition parameters. As a result, these models can suffer from degraded generalization performance when inevitably applied to new settings that differ from their training distribution [203], [204]. Curating diverse and representative training datasets is particularly challenging in clinical environments, and even more so for multi-contrast protocols [205]. This data dependency raises a significant barrier for deployment of supervised DL in general-purpose MRI reconstruction.

To address these challenges, several alternative learning strategies can be considered. Common approaches include unsupervised paradigms where generative models learn the distribution of fully-sampled data without requiring paired examples [206]–[208], and self-supervised methods that exploit inherent redundancies in the undersampled data to guide training [76]. More broadly, federated learning provides a framework for decentralized training across institutions, enabling collaborative model development without the need to share any MRI data [209], [210]. In scenarios where collaboration is impractical, synthetic datasets generated using physically-informed signal models can help alleviate data scarcity [211].

In resource-constrained settings, adaptation of foundation models on limited fine-tuning sets, or the development of scan-specific models without separate training data might also prove effective. Foundation models, pre-trained on large-scale, diverse datasets, can offer improved generalization through fine-tuning with smaller domain-specific datasets or few-shot learning techniques [212], [213]. In contrast, scan-specific approaches typically perform reconstruction modeling on a single undersampled scan using the acquired data and the imaging operator [135]. While this approach enhances generalization to the individual scan's characteristics and avoids distribution shifts, it often requires longer reconstruction times compared to pre-trained models, which can be a trade-off in

time-sensitive clinical workflows.

# D. Resource-Aware Computational Design

A critical consideration for deploying MRI reconstruction methods in practice is their computational cost. MRI technicians often need to assess scan quality in real-time to determine whether a scan is diagnostically adequate or needs to be repeated due to motion artifacts, improper slice placement, or other issues [25]. Providing timely feedback is especially important in busy clinical environments or for motion-sensitive populations such as pediatric or elderly patients. Furthermore, MRI protocols frequently involve multi-contrast imaging with high-channel receiver arrays, resulting in large data volumes that pose significant computational burden—particularly when real-time or near-real-time reconstruction is desired [214].

Several strategies can help mitigate these challenges. Coil compression techniques can reduce channel dimensionality, for instance by projecting 32-channel data to 16 virtual coils, thereby accelerating subsequent processing [214]. Similarly, compression along contrast dimensions can reduce redundancy in multi-contrast acquisitions [63]. Parallelized implementations leveraging multiple GPUs or distributed compute nodes can further accelerate workflows [215]. Alternatively, sequential reconstruction pipelines may be employed, wherein an initial reference contrast is reconstructed with light undersampling and then used as a prior to guide joint reconstruction of the remaining contrasts [118]. While this approach may not fully exploit inter-contrast correlations, it offers a practical trade-off between computational efficiency and image quality.

#### VII. CONCLUSION

We reviewed a broad range of MRI reconstruction methods that incorporate prior knowledge to improve quality of images recovered from accelerated acquisitions. Modern DL methods show strong potential to shorten scan times, expand access to advanced imaging protocols previously limited to research settings, and provide higher-quality inputs for downstream analysis. However, their clinical translation depends on rigorous validation, clear regulatory frameworks, scalable computational designs, and ongoing collaboration with technologists and radiologists to identify failure modes and sustain user trust. Ultimately, realizing the full potential of these methods in routine patient care will require interdisciplinary efforts among researchers, clinicians, vendors, and regulators.

# REFERENCES

- [1] M. A. Bernstein, K. F. King, and X. J. Zhou, *Handbook of MRI pulse sequences*. Academic Press, 2004.
- [2] B. De Coene et al., "MR of the brain using fluid-attenuated inversion recovery (FLAIR) pulse sequences." AJNR Am. J. Neuroradiol., vol. 13, no. 6, pp. 1555–1564, 1992.
- [3] J. V. Hajnal et al., "Use of fluid attenuated inversion recovery (FLAIR) pulse sequences in MRI of the brain," J. Comput. Assist. Tomogr., vol. 16, no. 6, pp. 841–844, 1992.
- [4] M. E. Moseley et al., "Early detection of regional cerebral ischemia in cats: comparison of diffusion- and T2-weighted MRI and spectroscopy," Magn. Reson. Med., vol. 14, no. 2, pp. 330–346, 1990.

- [5] D.-M. Koh and D. J. Collins, "Diffusion-weighted MRI in the body: applications and challenges in oncology," AJR Am. J. Roentgenol., vol. 188, no. 6, pp. 1622–1635, 2007.
- [6] L. De Rochefort et al., "Quantitative susceptibility map reconstruction from MR phase data using Bayesian regularization: validation and application to brain imaging," Magn. Reson. Med., vol. 63, no. 1, pp. 194–206, 2010.
- [7] E. M. Haacke et al., "Susceptibility weighted imaging (SWI)," Magn. Reson. Med., vol. 52, no. 3, pp. 612–618, 2004.
- [8] J. W. Belliveau *et al.*, "Functional mapping of the human visual cortex by magnetic resonance imaging," *Science*, vol. 254, no. 5032, pp. 716– 719, 1991.
- [9] S. Ogawa et al., "Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields," Magn. Reson. Med., vol. 14, no. 1, pp. 68–78, 1990.
- [10] K. Setsompop et al., "Pushing the limits of in vivo diffusion MRI for the Human Connectome Project," *Neuroimage*, vol. 80, pp. 220–233, 2013.
- [11] J. A. Detre et al., "Perfusion imaging," Magn. Reson. Med., vol. 23, no. 1, pp. 37–45, 1992.
- [12] F.-H. Lin et al., "Sensitivity-encoded (SENSE) proton echo-planar spectroscopic imaging (PEPSI) in the human brain," Magn. Reson. Med., vol. 57, no. 2, pp. 249–257, 2007.
- [13] J. A. Fessler, "Model-based image reconstruction for MRI," *IEEE Signal Process. Mag.*, vol. 27, no. 4, pp. 81–89, 2010.
- [14] K. Hammernik et al., "Physics-driven deep learning for computational magnetic resonance imaging: combining physics and machine learning for improved medical imaging," *IEEE Signal Process. Mag.*, vol. 40, no. 1, pp. 98–114, 2023.
- [15] D. K. Sodickson and W. J. Manning, "Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays," *Magn. Reson. Med.*, vol. 38, no. 4, pp. 591–603, 1997.
- [16] K. P. Pruessmann et al., "SENSE: sensitivity encoding for fast MRI," Magn. Reson. Med., vol. 42, no. 5, pp. 952–962, 1999.
- [17] M. A. Griswold et al., "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," Magn. Reson. Med., vol. 47, no. 6, pp. 1202– 1210, 2002.
- [18] P. B. Roemer et al., "The NMR phased array," Magn. Reson. Med., vol. 16, no. 2, pp. 192–225, 1990.
- [19] S. Ravishankar, J. C. Ye, and J. A. Fessler, "Image reconstruction: from sparsity to data-adaptive methods and machine learning," *Proc. IEEE*, vol. 108, no. 1, pp. 86–109, 2020.
- [20] X. Wang et al., "Physics-based reconstruction methods for magnetic resonance imaging," Philos. Trans. R. Soc. A, Math. Phys. Eng. Sci., vol. 379, no. 2200, p. 20200196, 2021.
- [21] D. Liang et al., "Deep magnetic resonance image reconstruction: Inverse problems meet neural networks," IEEE Signal Process. Mag., vol. 37, no. 1, pp. 141–151, 2020.
- [22] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: the application of compressed sensing for rapid MR imaging," *Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [23] K. T. Block, M. Uecker, and J. Frahm, "Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint," *Magn. Reson. Med.*, vol. 57, no. 6, pp. 1086–1098, 2007.
- [24] M. Akçakaya *et al.*, "Compressed sensing with wavelet domain dependencies for coronary MRI: a retrospective study," *IEEE Trans. Med. Imaging*, vol. 30, no. 5, pp. 1090–1099, 2011.
  [25] M. Murphy *et al.*, "Fast \(\ell\_1\)-SPIRiT compressed sensing parallel
- [25] M. Murphy et al., "Fast ℓ<sub>1</sub>-SPIRiT compressed sensing parallel imaging MRI: scalable parallel implementation and clinically feasible runtime," *IEEE Trans. Med. Imaging*, vol. 31, no. 6, pp. 1250–1262, 2012.
- [26] D. Liang et al., "Accelerating SENSE using compressed sensing," Magn. Reson. Med., vol. 62, no. 6, pp. 1574–1584, 2009.
- [27] L. Feng et al., "Compressed sensing for body MRI," J. Magn. Reson. Imaging, vol. 45, no. 4, pp. 966–987, 2017.
- [28] P. J. Shin et al., "Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion," Magn. Reson. Med., vol. 72, no. 4, pp. 959–970, 2014.
- [29] M. Mani et al., "Multi-shot sensitivity-encoded diffusion data recovery using structured low-rank matrix completion (MUSSELS)," Magn. Reson. Med., vol. 78, no. 2, pp. 494–507, 2017.
- [30] J. P. Haldar, "Low-rank modeling of local k-space neighborhoods (LORAKS) for constrained MRI," *IEEE Trans. Med. Imaging*, vol. 33, no. 3, pp. 668–681, 2014.
- [31] D. Lee et al., "Acceleration of MR parameter mapping using annihilating filter-based low rank Hankel matrix (ALOHA)," Magn. Reson. Med., vol. 76, no. 6, pp. 1848–1864, 2016.

- [32] C. M. Sandino et al., "Compressed sensing: From research to clinical practice with deep neural networks: Shortening scan times for magnetic resonance imaging," *IEEE Signal Process. Mag.*, vol. 37, no. 1, pp. 117–127, 2020.
- [33] Y. Chen et al., "AI-based reconstruction for fast MRI—A systematic review and meta-analysis," Proc. IEEE, vol. 110, no. 2, pp. 224–245, 2022.
- [34] R. Heckel et al., "Deep learning for accelerated and robust MRI reconstruction," MAGMA, vol. 37, no. 3, pp. 335–368, 2024.
- [35] S. Wang et al., "Accelerating magnetic resonance imaging via deep learning," in *IEEE ISBI*, 2016, pp. 514–517.
- [36] Y. Yang et al., "ADMM-CSNet: A deep learning approach for image compressive sensing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 3, pp. 521–538, 2020.
- [37] K. Hammernik et al., "Learning a variational network for reconstruction of accelerated MRI data," Magn. Reson. Med., vol. 79, no. 6, pp. 3055–3071, 2018.
- [38] M. Mardani et al., "Deep generative adversarial neural networks for compressive sensing MRI," *IEEE Trans. Med. Imaging*, vol. 38, no. 1, pp. 167–179, 2019.
- [39] J. Schlemper et al., "A deep cascade of convolutional neural networks for MR image reconstruction," in IPMI. Springer, 2017, pp. 647–658.
- [40] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL: Model-Based deep learning architecture for inverse problems," *IEEE Trans. Med. Imaging*, vol. 38, no. 2, pp. 394–405, 2019.
- [41] T. M. Quan, T. Nguyen-Duc, and W.-K. Jeong, "Compressed sensing MRI reconstruction using a generative adversarial network with a cyclic loss," *IEEE Trans. Med. Imaging*, vol. 37, no. 6, pp. 1488–1497, 2018.
- [42] A. Sriram et al., "End-to-end variational networks for accelerated MRI reconstruction," in MICCAI, 2020, pp. 64–73.
- [43] S. Biswas, H. K. Aggarwal, and M. Jacob, "Dynamic MRI using model-based deep learning and SToRM priors: MoDL-SToRM," *Magn. Reson. Med.*, vol. 82, no. 1, pp. 485–494, 2019.
- [44] S. U. H. Dar et al., "A transfer-learning approach for accelerated MRI using deep neural networks," Magn. Reson. Med., vol. 84, no. 2, pp. 663–685, 2020.
- [45] B. Zhu et al., "Image reconstruction by domain-transform manifold learning," *Nature*, vol. 555, no. 7697, pp. 487–492, 2018.
- [46] R. Nath et al., "Accelerated phase contrast magnetic resonance imaging via deep learning," in IEEE ISBI, 2020, pp. 834–838.
- [47] D. K. Jones, Ed., Diffusion MRI: Theory, methods, and applications. Oxford Univ. Press, 2010.
- [48] M. Uecker et al., "ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA," Magn. Reson. Med., vol. 71, no. 3, pp. 990–1001, 2014.
- [49] R. A. Lobos, C.-C. Chan, and J. P. Haldar, "New theory and faster computations for subspace-based sensitivity map estimation in multichannel MRI," *IEEE Trans. Med. Imaging*, vol. 43, no. 1, pp. 286–296, 2024.
- [50] L. Ying and J. Sheng, "Joint image reconstruction and sensitivity estimation in SENSE (JSENSE)," Magn. Reson. Med., vol. 57, no. 6, pp. 1196–1202, 2007.
- [51] M. Uecker et al., "Image reconstruction by regularized nonlinear inversion—Joint estimation of coil sensitivities and image content," Magn. Reson. Med., vol. 60, no. 3, pp. 674–682, 2008.
- [52] B. Liu et al., "Regularized sensitivity encoding (SENSE) reconstruction using Bregman iterations," Magn. Reson. Med., vol. 61, no. 1, pp. 145– 152, 2009.
- [53] R. Otazo et al., "Combination of compressed sensing and parallel imaging for highly accelerated first-pass cardiac perfusion MRI," Magn. Reson. Med., vol. 64, no. 3, pp. 767–776, 2010.
- [54] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," SIAM J. Imaging Sci., vol. 2, no. 1, pp. 183–202, 2009.
- [55] R. Otazo, E. Candes, and D. K. Sodickson, "Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components," *Magn. Reson. Med.*, vol. 73, no. 3, pp. 1125–1136, 2015.
- [56] J. P. Haldar, "Autocalibrated LORAKS for fast constrained MRI reconstruction," in *IEEE ISBI*. IEEE, 2015, pp. 910–913.
- [57] J. P. Haldar and J. Zhuo, "P-LORAKS: low-rank modeling of local k-space neighborhoods with parallel imaging data," *Magn. Reson. Med.*, vol. 75, no. 4, pp. 1499–1514, 2016.
- [58] T. H. Kim, K. Setsompop, and J. P. Haldar, "LORAKS makes better SENSE: phase-constrained partial Fourier SENSE reconstruction without phase calibration," *Magn. Reson. Med.*, vol. 77, no. 3, pp. 1021–1035, 2017.

- [59] B. Bilgic et al., "Improving parallel imaging by jointly reconstructing multi-contrast data," Magn. Reson. Med., vol. 80, no. 2, pp. 619–632, 2018
- [60] G. Ongie and M. Jacob, "A fast algorithm for convolutional structured low-rank matrix recovery," *IEEE Trans. Comput. Imaging*, vol. 3, no. 4, pp. 535–550, 2017.
- [61] S. Ravishankar and Y. Bresler, "MR image reconstruction from highly undersampled k-space data by dictionary learning," *IEEE Trans. Med. Imaging*, vol. 30, no. 5, pp. 1028–1041, 2011.
- [62] J. I. Tamir et al., "T2 shuffling: sharp, multicontrast, volumetric fast spin-echo imaging," Magn. Reson. Med., vol. 77, no. 1, pp. 180–195, 2017.
- [63] D. Ma et al., "Magnetic resonance fingerprinting," Nature, vol. 495, no. 7440, pp. 187–192, 2013.
- [64] B. Zhao et al., "Improved magnetic resonance fingerprinting reconstruction with low-rank and subspace modeling," Magn. Reson. Med., vol. 79, no. 2, pp. 933–942, 2018.
- [65] X. Cao et al., "Optimized multi-axis spiral projection MR fingerprinting with subspace reconstruction for rapid whole-brain high-isotropicresolution quantitative imaging," Magn. Reson. Med., vol. 88, no. 1, pp. 133–150, 2022.
- [66] T. J. Sumpf et al., "Model-based nonlinear inverse reconstruction for T2 mapping using highly undersampled spin-echo MRI," J Magn. Reson. Imaging, vol. 34, no. 2, pp. 420–428, 2011.
- [67] G. Yang et al., "DAGAN: Deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction," IEEE Trans. Med. Imaging, vol. 37, no. 6, pp. 1310–1321, 2018.
- [68] G. Wang et al., "B-spline parameterized joint optimization of reconstruction and k-space trajectories (BJORK) for accelerated 2D MRI," IEEE Trans. Med. Imaging, vol. 41, no. 9, pp. 2318–2330, 2022.
- [69] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT Press, 2016.
- [70] C. G. Radhakrishna and P. Ciuciu, "Jointly learning non-Cartesian k-space trajectories and reconstruction networks for 2D and 3D MR imaging through projection," *Bioengineering*, vol. 10, no. 2, p. 158, 2023.
- [71] R. Ahmad et al., "Plug-and-play methods for magnetic resonance imaging: Using denoisers for image recovery," *IEEE Signal Process.* Mag., vol. 37, no. 1, pp. 105–116, 2020.
- [72] J. Liu et al., "RARE: Image reconstruction using deep priors learned without groundtruth," *IEEE J. Sel. Top. Signal Process.*, vol. 14, no. 6, pp. 1088–1099, 2020.
- [73] B. Zhou and S. K. Zhou, "DuDoRNet: Learning a dual-domain recurrent network for fast MRI reconstruction with deep T1 prior," in *IEEE CVPR*, 2020.
- [74] D. Polak et al., "Joint multi-contrast variational network reconstruction (jVN) with application to rapid 2D and 3D imaging," Magn. Reson. Med., vol. 84, no. 3, pp. 1456–1469, 2020.
- [75] L. Sun *et al.*, "A deep information sharing network for multi-contrast compressed sensing MRI reconstruction," *IEEE Trans. Image Process.*, vol. 28, no. 12, pp. 6141–6153, 2019.
  [76] B. Yaman *et al.*, "Self-supervised learning of physics-guided reconstructions."
- [76] B. Yaman et al., "Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data," Magn. Reson. Med., vol. 84, no. 6, pp. 3172–3191, 2020.
- [77] K. Ryu et al., "Accelerated multicontrast reconstruction for synthetic MRI using joint parallel imaging and variable splitting networks," Med. Phys., vol. 48, no. 6, pp. 2939–2950, 2021.
- [78] X. Liu et al., "On the regularization of feature fusion and mapping for fast MR multi-contrast imaging via iterative networks," Magn. Reson. Imaging, vol. 77, pp. 159–168, 2021.
- [79] Z. Ramzi et al., "NC-PDNet: A density-compensated unrolled network for 2D and 3D non-Cartesian MRI reconstruction," *IEEE Trans. Med. Imaging*, vol. 41, no. 7, pp. 1625–1638, 2022.
- [80] F. Lam, X. Peng, and Z.-P. Liang, "High-dimensional MR spatiospectral imaging by integrating physics-based modeling and data-driven machine learning: Current progress and future directions," *IEEE Signal Process. Mag.*, vol. 40, no. 2, pp. 101–115, 2023.
- [81] P. L. Combettes and J.-C. Pesquet, "Proximal splitting methods in signal processing," in *Fixed-point algorithms for inverse problems in* science and engineering. Springer, 2011, pp. 185–212.
- [82] S. H. Chan, X. Wang, and O. A. Elgendy, "Plug-and-play ADMM for image restoration: Fixed-point convergence and applications," *IEEE Trans. Comput. Imaging*, vol. 3, no. 1, pp. 84–98, 2016.
- [83] Y. Romano, M. Elad, and P. Milanfar, "The little engine that could: Regularization by denoising (RED)," SIAM J. Imaging Sci., vol. 10, no. 4, pp. 1804–1844, 2017.

- [84] E. Ryu et al., "Plug-and-play methods provably converge with properly trained denoisers," in ICML. PMLR, 2019, pp. 5546–5557.
- [85] F. Liu, L. Feng, and R. Kijowski, "MANTIS: model-augmented neural network with incoherent k-space sampling for efficient MR parameter mapping," Magn. Reson. Med., vol. 82, no. 1, pp. 174–188, 2019.
- [86] M. P. Mani et al., "Model-based deep learning for reconstruction of joint k-q under-sampled high resolution diffusion MRI," in *IEEE ISBI*. IEEE, 2020, pp. 913–916.
- [87] K. C. Tezcan et al., "MR image reconstruction using deep density priors," *IEEE Trans. Med. Imaging*, vol. 38, no. 7, pp. 1633–1642, 2019.
- [88] G. Raut and A. Singh, "Generative AI in vision: A survey on models, metrics and applications," arXiv:2402.16369, 2024.
- [89] D. Narnhofer et al., "Inverse GANs for accelerated MRI reconstruction," in Proc. SPIE, vol. 11138. SPIE, 2019, pp. 381–392.
- [90] Y. Korkmaz et al., "Unsupervised MRI reconstruction via zero-shot learned adversarial transformers," *IEEE Trans. Med. Imaging*, vol. 41, no. 7, pp. 1747–1763, 2022.
- [91] A. Jalal et al., "Robust compressed sensing MRI with deep generative priors," in *NeurIPS*, vol. 34, 2021, pp. 14938–14954.
- [92] H. Chung and J. C. Ye, "Score-based diffusion models for accelerated MRI," Med. Image Anal., vol. 80, p. 102479, 2022.
- [93] G. Luo et al., "Bayesian MRI reconstruction with joint uncertainty estimation using diffusion models," Magn. Reson. Med., vol. 90, no. 1, pp. 295–311, 2023.
- [94] Y. Korkmaz, T. Cukur, and V. M. Patel, "Self-supervised MRI reconstruction with unrolled diffusion models," in *MICCAI*. Springer, 2023, pp. 491–501.
- [95] A. Güngör et al., "Adaptive diffusion priors for accelerated MRI reconstruction," Med. Image Anal., vol. 88, p. 102872, 2023.
- [96] S. Liang et al., "Analysis of deep image prior and exploiting self-guidance for image reconstruction," IEEE Trans. Comput. Imaging, 2025.
- [97] K. Kwon, D. Kim, and H. Park, "A parallel MR imaging method using multilayer perceptron," *Med. Phys.*, vol. 44, no. 12, pp. 6209–6224, 2017
- [98] M. Akçakaya et al., "Scan-specific robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction: database-free deep learning for fast imaging," Magn. Reson. Med., vol. 81, no. 1, pp. 439–453, 2019.
- [99] A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in ICLR, 2021.
- [100] P. Guo et al., "Reconformer: Accelerated MRI reconstruction using recurrent transformer," IEEE Trans. Med. Imaging, 2024.
- [101] J. Huang et al., "Swin transformer for fast MRI," Neurocomputing, vol. 493, pp. 281–304, 2022.
- [102] C.-M. Feng et al., "Task transformer network for joint MRI reconstruction and super-resolution," in MICCAI. Springer, 2021, pp. 307–317.
- [103] C. Qin et al., "Convolutional recurrent neural networks for dynamic MR image reconstruction," *IEEE Trans. Med. Imaging*, vol. 38, no. 1, pp. 280–290, 2018.
- [104] S. A. H. Hosseini et al., "Dense recurrent neural networks for accelerated MRI: History-cognizant unrolling of optimization algorithms," IEEE J. Sel. Top. Signal Process., vol. 14, no. 6, pp. 1280–1291, 2020.
- [105] P. Guo et al., "Over-and-under complete convolutional RNN for MRI reconstruction," in MICCAI. Springer, 2021, pp. 13–23.
- [106] A. Gu and T. Dao, "Mamba: Linear-time sequence modeling with selective state spaces," arXiv:2312.00752, 2023.
- [107] J. Huang et al., "Enhancing global sensitivity and uncertainty quantification in medical image reconstruction with Monte Carlo arbitrary-masked Mamba," Med. Image Anal., vol. 99, p. 103334, 2025.
- [108] Y. Korkmaz and V. M. Patel, "MambaRecon: MRI reconstruction with structured state space models," in *IEEE/CVF WACV*, 2025, pp. 4142– 4152
- [109] B. Kabas et al., "Physics-driven autoregressive state space models for medical image reconstruction," arXiv:2412.09331, 2024.
- [110] O. F. Atli *et al.*, "I2I-Mamba: Multi-modal medical image synthesis via selective state space modeling," *arXiv:2405.14022*, 2024.
- [111] Y. Han, L. Sunwoo, and J. C. Ye, "k-space deep learning for accelerated MRI," *IEEE Trans. Med. Imaging*, vol. 39, no. 2, pp. 377–386, 2020.
- [112] M. U. Mirza et al., "Learning Fourier-constrained diffusion bridges for MRI reconstruction," arXiv:2308.01096, 2023.
- [113] T. Eo et al., "KIKI-net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images," Magn. Reson. Med., vol. 80, no. 5, pp. 2188–2201, 2018.

- [114] S. Wang et al., "DIMENSION: Dynamic MR imaging with both k-space and spatial prior knowledge obtained via multi-supervised network training," NMR Biomed., p. e4131, 2022.
- [115] Q. Liu et al., "Highly undersampled magnetic resonance imaging reconstruction using autoencoding priors," Magn. Reson. Med., vol. 83, no. 1, pp. 322–336, 1 2020.
- [116] V. A. Nezhad et al., "Generative autoregressive transformers for modelagnostic federated MRI reconstruction," arXiv:2502.04521, 2025.
- [117] S. Wang et al., "Knowledge-driven deep learning for fast MR imaging: Undersampled MR image reconstruction from supervised to unsupervised learning," Magn. Reson. Med., vol. 92, no. 2, pp. 496–518, 2024.
- [118] E. Gong et al., "PROMISE: parallel-imaging and compressed-sensing reconstruction of multicontrast imaging using SharablE information," Magn. Reson. Med., vol. 73, no. 2, pp. 523–535, 2015.
- [119] B. Bilgic, V. K. Goyal, and E. Adalsteinsson, "Multi-contrast reconstruction with Bayesian compressed sensing," *Magn. Reson. Med.*, vol. 66, no. 6, pp. 1601–1615, 2011.
- [120] L. Xiang et al., "Deep-learning-based multi-modal fusion for fast MR reconstruction," *IEEE Trans. Biomed. Eng.*, vol. 66, no. 7, pp. 2105– 2114, 2019.
- [121] W.-J. Do et al., "Reconstruction of multicontrast MR images through deep learning," Med. Phys., vol. 47, no. 3, pp. 983–997, 2020.
- [122] E. Kopanoglu et al., "Simultaneous use of individual and joint regularization terms in compressive sensing: Joint reconstruction of multichannel multi-contrast MRI acquisitions," NMR Biomed., vol. 33, no. 4, p. e4247, 2020.
- [123] S. U. Dar et al., "Prior-guided image reconstruction for accelerated multi-contrast MRI via generative adversarial networks," *IEEE J. Sel. Top. Signal Process.*, vol. 14, no. 6, pp. 1072–1087, 2020.
- [124] K. H. Kim, W. Do, and S. Park, "Improving resolution of MR images with an adversarial network incorporating images with different contrast," *Med. Phys.*, vol. 45, no. 7, pp. 3120–3131, 2018.
- [125] X. Liu et al., "Multi-contrast MR reconstruction with enhanced denoising autoencoder prior learning," in ISBI. IEEE, 2020, pp. 1–5.
- [126] G. Rizzuti, A. Sbrizzi, and T. Van Leeuwen, "Joint retrospective motion correction and reconstruction for brain MRI with a reference contrast," *IEEE Trans. Comput. Imaging*, vol. 8, pp. 490–504, 2022.
- [127] T. Oved et al., "Deep learning of personalized priors from past MRI scans enables fast, quality-enhanced point-of-care MRI with low-cost systems," arXiv:2505.02470, 2025.
- [128] F. Knoll et al., "Deep-learning methods for parallel magnetic resonance imaging reconstruction: A survey of the current approaches, trends, and issues," *IEEE Signal Process. Mag.*, vol. 37, no. 1, pp. 128–140, 2020.
- [129] J. I. Tamir, S. X. Yu, and M. Lustig, "Unsupervised deep basis pursuit: Learning reconstruction without ground-truth data," in *Proc. ISMRM*, 2019, p. 0660.
- [130] M. Blumenthal et al., "Self-supervised learning for improved calibrationless radial MRI with NLINV-Net," Magn. Reson. Med., vol. 92, no. 6, pp. 2447–2463, 2024.
- [131] E. K. Cole et al., "Fast unsupervised MRI reconstruction without fully-sampled ground truth data using generative adversarial networks," in *ICCVW*, 2021, pp. 3988–3997.
- [132] C. Millard and M. Chiew, "A theoretical framework for self-supervised MR image reconstruction using sub-sampling via variable density Noisier2Noise," *IEEE Trans. Comput. Imaging*, vol. 9, pp. 707–720, 2023
- [133] A. D. Desai et al., "Noise2Recon: Enabling SNR-robust MRI reconstruction with semi-supervised and self-supervised learning," Magn. Reson. Med., vol. 90, no. 5, pp. 2052–2070, 2023.
- [134] T. H. Kim, P. Garg, and J. P. Haldar, "LORAKI: Autocalibrated recurrent neural networks for autoregressive MRI reconstruction in kspace," arXiv:1904.09390, 2019.
- [135] Y. Arefeen et al., "Scan-specific artifact reduction in k-space (SPARK) neural networks synergize with physics-based reconstruction to accelerate MRI," Magn. Reson. Med., vol. 87, no. 2, pp. 764–780, 2022.
- [136] K. H. Jin et al., "Time-dependent deep image prior for dynamic MRI," arXiv:1910.01684, 2019.
- [137] S. Arora, V. Roeloffs, and M. Lustig, "Untrained modified deep decoder for joint denoising and parallel imaging reconstruction," in *Proc.* ISMRM, 2020, p. 3585.
- [138] M. Z. Darestani and R. Heckel, "Accelerated MRI with un-trained neural networks," *IEEE Trans. Comput. Imaging*, vol. 7, pp. 724–733, 2021.
- [139] K. P. Slavkova et al., "An untrained deep learning method for reconstructing dynamic MR images from accelerated model-based data," Magn. Reson. Med., vol. 89, no. 4, pp. 1617–1633, 2023.

- [140] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep Image Prior," Int. J. Comput. Vis., vol. 128, no. 7, pp. 1867–1888, 2020.
- [141] B. Yaman, S. A. H. Hosseini, and M. Akcakaya, "Zero-shot self-supervised learning for MRI reconstruction," in ICLR, 2022.
- [142] H. Wang et al., "Early stopping for Deep Image Prior," Trans. Mach. Learn. Res., 2021.
- [143] Y. U. Alçalar and M. Akçakaya, "Zero-shot adaptation for approximate posterior sampling of diffusion models in inverse problems," in ECCV. Springer, 2024, pp. 444–460.
- [144] D. C. Levin and V. M. Rao, "Factors that will determine future utilization trends in diagnostic imaging," *J. Am. Coll. Radiol.*, vol. 13, no. 8, pp. 904–908, 2016.
- [145] S. Kiryu et al., "Clinical impact of deep learning reconstruction in MRI," Radiographics, vol. 43, no. 6, 2023.
- [146] W. Peng et al., "Prospective and multi-reader evaluation of deep learning reconstruction-based accelerated rectal MRI: image quality, diagnostic performance, and reading time," Eur. Radiol., vol. 34, no. 11, pp. 7438–7449, 2024.
- [147] S. W. Mooney, P. J. Marlow, and B. L. Anderson, "The perception and misperception of optical defocus, shading, and shape," *Elife*, vol. 8, Jul. 2019.
- [148] H. Yoo et al., "Deep learning-based reconstruction for acceleration of lumbar spine MRI: a prospective comparison with standard MRI," Eur. Radiol., vol. 33, no. 12, pp. 8656–8668, 2023.
- [149] O. Afacan et al., "Evaluation of motion and its effect on brain magnetic resonance image quality in children," *Pediatr. Radiol.*, vol. 46, no. 12, pp. 1728–1735, Nov. 2016.
- [150] J. B. Andre *et al.*, "Toward quantifying the prevalence, severity, and cost associated with patient motion during clinical MR examinations," *J. Am. Coll. Radiol.*, vol. 12, no. 7, pp. 689–695, Jul. 2015.
- [151] M. Zaitsev, J. Maclaren, and M. Herbst, "Motion artifacts in MRI: A complex problem with many partial solutions," *J. Magn. Reson. Imaging*, vol. 42, no. 4, pp. 887–901, 2015.
- [152] X. V. Nguyen et al., "Cost economy of motion," in Advances in Magnetic Resonance Technology and Applications. Elsevier, 2022, pp. 25–34.
- [153] N. Seiberlich et al., Eds., Quantitative magnetic resonance imaging: Volume 1, ser. Advances in Magnetic Resonance Technology and Applications. San Diego, CA: Academic Press, 2020.
- [154] L. Feng et al., "XD-GRASP: Golden-angle radial MRI with reconstruction of extra motion-state dimensions using compressed sensing," Magn. Reson. Med., vol. 75, no. 2, pp. 775–788, 2016.
- [155] A. Towfighi et al., "Bridging the gap between research, policy, and practice: Lessons learned from academic-public partnerships in the CTSA network," J. Clin. Transl. Sci., vol. 4, no. 3, pp. 201–208, 2020.
- [156] "What will it take to make precision health a global reality," Nat. Med., vol. 30, no. 7, pp. 1793–1794, 2024.
- [157] V. M. Rao et al., "Multimodal generative AI for medical image interpretation," *Nature*, vol. 639, no. 8056, pp. 888–896, 2025.
- [158] A. Hagiwara et al., "Multiparametric MRI: From simultaneous rapid acquisition methods and analysis techniques using scoring, machine learning, radiomics, and deep learning to the generation of novel metrics," *Invest. Radiol.*, 2023.
- [159] J. E. van Timmeren et al., "Radiomics in medical imaging-"how-to" guide and critical reflection," *Insights Imaging*, vol. 11, no. 1, p. 91, 2020.
- [160] A. Hagiwara et al., "Variability and standardization of quantitative imaging: Monoparametric to multiparametric quantification, radiomics, and artificial intelligence," *Invest. Radiol.*, vol. 55, no. 9, pp. 601–616, 2020.
- [161] H. Paverd et al., "Radiology and multi-scale data integration for precision oncology," NPJ Precis. Oncol., vol. 8, no. 1, p. 158, Jul. 2024.
- [162] D. J. Lin et al., "Artificial intelligence for MR image reconstruction: An overview for clinicians," J. Magn. Reson. Imaging, vol. 53, no. 4, pp. 1015–1028, 2021.
- [163] I. Conti et al., "Magnetic resonance imaging in pediatric epilepsy patients: two different protocols," J. Med. Imaging Interv. Radiol., vol. 11, no. 1, 2024.
- [164] D. V. Flores et al., "Distal radioulnar joint: Normal anatomy, imaging of common disorders, and injury classification," *Radiographics*, vol. 43, no. 1, p. e220109, Jan. 2023.
- [165] P. M. Johnson *et al.*, "Deep learning reconstruction enables prospectively accelerated clinical knee MRI," *Radiology*, vol. 307, no. 2, p. e220425, 2023.
- [166] A. S. Purysko et al., "RadioGraphics update: PI-RADS version 2.1-a pictorial update," Radiographics, vol. 40, no. 7, pp. E33–E37, 2020.

- [167] P. S. Rajiah, C. J. François, and T. Leiner, "Cardiac MRI: State of the art," *Radiology*, vol. 307, no. 3, p. e223008, 2023.
- [168] J. A. Oscanoa et al., "Deep learning-based reconstruction for cardiac MRI: a review," Bioengineering, vol. 10, no. 3, p. 334, 2023.
- [169] M. A. Morales, W. J. Manning, and R. Nezafat, "Present and future innovations in AI and cardiac MRI," *Radiology*, vol. 310, no. 1, p. e231269, 2024.
- [170] F. E. Lecouvet *et al.*, "Present and future of whole-body MRI in metastatic disease and myeloma: How and why you will do it," *Skeletal Radiol.*, vol. 53, no. 9, pp. 1815–1831, 2024.
- [171] J. M. Pollard *et al.*, "The future of image-guided radiotherapy will be MR guided," *Br. J. Radiol.*, vol. 90, no. 1073, p. 20160667, 2017.
- [172] S. De Pietro *et al.*, "The role of MRI in radiotherapy planning: a narrative review "from head to toe"," *Insights Imaging*, vol. 15, no. 1, p. 255, 2024.
- [173] R. Otazo et al., "MRI-guided radiation therapy: An emerging paradigm in adaptive radiation oncology," *Radiology*, vol. 298, no. 2, pp. 248– 260, 2021.
- [174] S. M. Thompson *et al.*, "Body interventional MRI for diagnostic and interventional radiologists: Current practice and future prospects," *Radiographics*, vol. 41, no. 6, pp. 1785–1801, 2021.
- [175] A. E. Campbell-Washburn *et al.*, "Opportunities in interventional and diagnostic imaging by using high-performance low-field-strength MRI," *Radiology*, vol. 293, no. 2, pp. 384–393, 2019.
- [176] E. Shimron et al., "Accelerating Low-field MRI: Compressed Sensing and AI for fast noise-robust imaging," arXiv:2411.06704, 2024.
- [177] F. Gannon, "The steps from translatable to translational research," EMBO Rep., vol. 15, no. 11, pp. 1107–1108, 2014.
- [178] D. G. Fort et al., "Mapping the evolving definitions of translational research," J. Clin. Transl. Sci., vol. 1, no. 1, pp. 60–66, 2017.
- [179] R. Reisenhofer et al., "A Haar wavelet-based perceptual similarity index for image quality assessment," Signal Process. Image Commun., vol. 61, pp. 33–43, 2018.
- [180] C. Karner et al., "Parameter choices in HaarPSI for IQA with medical images," in IEEE ISBI. IEEE, 2025, pp. 1–5.
- [181] F. Knoll et al., "Advancing machine learning for MR image reconstruction with an open competition: Overview of the 2019 fastMRI challenge," Magn. Reson. Med., vol. 84, no. 6, pp. 3054–3070, 2020.
- [182] Center for Devices and Radiological Health, "Methods and tools for effective postmarket monitoring of artificial intelligence (AI)-enabled medical devices," https://www.fda.gov/.../ai-enabled-medical-devices, 2024, [Online]. Accessed on Apr. 29, 2025.
- [183] V. Antun et al., "On instabilities of deep learning in image reconstruction and the potential costs of AI," Proc. Natl. Acad. Sci., vol. 117, no. 48, pp. 30 088–30 095, 2020.
- [184] E. Shimron et al., "Implicit data crimes: Machine learning bias arising from misuse of public data," Proc. Natl. Acad. Sci, vol. 119, no. 13, p. e2117203119, 2022.
- [185] A. Sriram et al., "GrappaNet: Combining parallel imaging with deep learning for multi-coil MRI reconstruction," in *IEEE CVPR*, 2020, pp. 14315–14322.
- [186] J. Zhang and B. Ghanem, "ISTA-Net: Interpretable optimizationinspired deep network for image compressive sensing," in *IEEE CVPR*, 2018, pp. 1828–1837.
- [187] J. Cheng et al., "Model learning: Primal dual networks for fast MR imaging," in MICCAI, 2019, pp. 21–29.
- [188] S. U. H. Dar et al., "Parallel-stream fusion of scan-specific and scan-general priors for learning deep MRI reconstruction in low-data regimes," Comput. Biol. Med., vol. 167, p. 107610, 2023.
- [189] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL-MUSSELS: model-based deep learning for multishot sensitivity-encoded diffusion MRI," *IEEE Trans. Med. Imaging*, vol. 39, no. 4, pp. 1268–1277, 2020.
- [190] D. S. Weller, M. Salerno, and C. H. Meyer, "Content-aware compressive magnetic resonance image reconstruction," *Magn. Reson. Imaging*, vol. 52, pp. 118–130, 2018.
- [191] V. Edupuganti et al., "Uncertainty quantification in deep MRI reconstruction," *IEEE Trans. Med. Imaging*, vol. 40, no. 1, pp. 239–250, 2020.
- [192] S. Bhadra et al., "On hallucinations in tomographic image reconstruction," *IEEE Trans. Med. Imaging*, vol. 40, no. 11, pp. 3249–3260, 2021.
- [193] C.-C. Chan and J. P. Haldar, "Local perturbation responses and checker-board tests: Characterization tools for nonlinear MRI methods," *Magn. Reson. Med.*, vol. 86, no. 4, pp. 1873–1887, 2021.
- [194] S. Bach et al., "On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation," PloS one, vol. 10, no. 7, p. e0130140, 2015.

- [195] H. A. Bedel and T. Çukur, "DreaMR: Diffusion-driven counterfactual explanation for functional MRI," *IEEE Trans. Med. Imaging*, 2024.
- [196] M. Sundararajan, A. Taly, and Q. Yan, "Axiomatic attribution for deep networks," in *ICML*. PMLR, 2017, pp. 3319–3328.
- [197] E. A. Stanley, M. Wilms, and N. D. Forkert, "Disproportionate subgroup impacts and other challenges of fairness in artificial intelligence for medical image analysis," in *EPIMI*. Springer, 2022, pp. 14–25.
- [198] Y. Du et al., "Unveiling fairness biases in deep learning-based brain MRI reconstruction," in CLIP EPIMI FAIMI, 2023, pp. 102–111.
- [199] P. Guo et al., "Multi-institutional collaborations for improving deep learning-based magnetic resonance image reconstruction using federated learning," in CVPR, vol. 13, 2021, p. 2423–2432.
- [200] N. K. Dinsdale, M. Jenkinson, and A. I. Namburete, "Deep learning-based unlearning of dataset bias for MRI harmonisation and confound removal," *NeuroImage*, vol. 228, p. 117689, 2021.
- [201] O. Dalmaz et al., "One model to unite them all: Personalized federated learning of multi-contrast MRI synthesis," Med. Image Anal., p. 103121, 2024.
- [202] Z. Xu et al., "Addressing fairness issues in deep learning-based medical image analysis: a systematic review," NPJ Digit. Med., vol. 7, no. 1, p. 286, 2024.
- [203] F. Liu et al., "SANTIS: sampling-augmented neural network with incoherent structure for MR image reconstruction," Magn. Reson. Med., vol. 82, no. 5, pp. 1890–1904, 2019.
- [204] Y. Han et al., "Deep learning with domain adaptation for accelerated projection-reconstruction MR," Magn. Reson. Med., vol. 80, no. 3, pp. 1189–1205, 2018.
- [205] F. Knoll et al., "fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning," Radiol. Artif. Intell, vol. 2, no. 1, p. e190007, 2020.
- [206] G. Oh et al., "Unpaired deep learning for accelerated MRI using optimal transport driven CycleGAN," *IEEE Trans. Comput. Imaging*, vol. 6, pp. 1285–1296, 2021.
- [207] K. Lei et al., "Wasserstein GANs for MR imaging: from paired to unpaired training," *IEEE Trans. Med. Imaging*, vol. 40, no. 1, pp. 105– 115, 2021.
- [208] M. Ozbey et al., "Unsupervised medical image translation with adversarial diffusion models," *IEEE Trans. Med. Imaging*, vol. 42, no. 12, pp. 3524–3539, 2023.
- [209] G. Elmas et al., "Federated learning of generative image priors for MRI reconstruction," *IEEE Trans. Med. Imaging*, vol. 42, no. 7, pp. 1996–2009, 2023.
- [210] B. R. Levac, M. Arvinte, and J. I. Tamir, "Federated end-to-end unrolled models for magnetic resonance image reconstruction," *Bio-engineering*, vol. 10, no. 3, p. 364, 2023.
- [211] Z. Wang et al., "One for multiple: Physics-informed synthetic data boosts generalizable deep learning for fast MRI reconstruction," Med. Image Anal., p. 103616, 2025.
- [212] Y. Sun et al., "A foundation model for enhancing magnetic resonance images and downstream segmentation, registration and diagnostic tasks," Nat. Biomed. Eng., vol. 9, no. 4, pp. 521–538, 2025.
- [213] H. Dong *et al.*, "MRI-CORE: A foundation model for magnetic resonance imaging," *arXiv:2506.12186*, 2025.
- [214] T. Zhang et al., "Coil compression for accelerated imaging with Cartesian sampling," Magn. Reson. Med., vol. 69, no. 2, pp. 571–582, 2013.
- [215] S. Schaetz et al., "Accelerated computing in magnetic resonance imaging: real-time imaging using nonlinear inverse reconstruction," Comput. Math. Methods Med., vol. 2017, no. 1, p. 3527269, 2017.