



Deep learning for accelerated and robust MRI reconstruction

Reinhard Heckel¹ · Mathews Jacob² · Akshay Chaudhari^{3,4} · Or Perlman^{5,6} · Efrat Shimron^{7,8} 

Received: 3 April 2024 / Revised: 24 May 2024 / Accepted: 28 May 2024 / Published online: 23 July 2024
© The Author(s) 2024, corrected publication 2024

Abstract

Deep learning (DL) has recently emerged as a pivotal technology for enhancing magnetic resonance imaging (MRI), a critical tool in diagnostic radiology. This review paper provides a comprehensive overview of recent advances in DL for MRI reconstruction, and focuses on various DL approaches and architectures designed to improve image quality, accelerate scans, and address data-related challenges. It explores end-to-end neural networks, pre-trained and generative models, and self-supervised methods, and highlights their contributions to overcoming traditional MRI limitations. It also discusses the role of DL in optimizing acquisition protocols, enhancing robustness against distribution shifts, and tackling biases. Drawing on the extensive literature and practical insights, it outlines current successes, limitations, and future directions for leveraging DL in MRI reconstruction, while emphasizing the potential of DL to significantly impact clinical imaging practices.

Keywords MRI · Deep learning · Machine learning · Image reconstruction

Introduction

Magnetic resonance imaging (MRI) has long been at the forefront of medical imaging, in that it offers unparalleled ability to visualize the human body's internal structures and functions. Because it is noninvasive and has exceptional soft tissue contrast, it has become an indispensable tool in modern diagnostic medicine. From detecting subtle pathological changes to guiding therapeutic interventions, MRI's versatility is unmatched. However, the full potential of MRI is

often constrained by inherent limitations in imaging speed and resolution, which are crucial for accurate diagnosis and patient comfort [1].

Long acquisition times represent a major barrier in clinical MRI, primarily due to the inherent trade-off between image quality and speed of imaging [2–8]. High-resolution images, which are crucial for accurate diagnosis, require extended scan durations that can be uncomfortable for patients and increase the risk of motion artifacts [9]. This challenge is particularly significant in dynamic imaging,

✉ Efrat Shimron
efrat.s@technion.ac.il

Reinhard Heckel
reinhard.heckel@tum.de

Mathews Jacob
mathews-jacob@uiowa.edu

Akshay Chaudhari
akshaysc@stanford.edu

Or Perlman
orperlman@tauex.tau.ac.il

¹ Department of computer engineering, Technical University of Munich, Munich, Germany

² Department of Electrical and Computer Engineering, University of Iowa, Iowa 52242, IA, USA

³ Department of Radiology, Stanford University, Stanford 94305, CA, USA

⁴ Department of Biomedical Data Science, Stanford University, Stanford 94305, CA, USA

⁵ Department of Biomedical Engineering, Tel Aviv University, Tel Aviv, Israel

⁶ Sagol School of Neuroscience, Tel Aviv University, Tel Aviv, Israel

⁷ Department of Electrical and Computer Engineering, Technion-Israel Institute of Technology, Haifa 3200004, Israel

⁸ Department of Biomedical Engineering, Technion-Israel Institute of Technology, Haifa 3200004, Israel

such as in cardiac and abdominal studies, where fast physiological movements can lead to blurring and other distortions. Moreover, sequences with long repetition times (TRs) like those used in diffusion MRI, extensive field-of-view (FOV) coverage, and protocols that require multiple contrasts significantly extend the duration of MRI scans. Additionally, the acquisition of 3D images, which provide comprehensive spatial detail for better clinical evaluation, also necessitates longer scan times due to the increased volume of data being collected. These complexities highlight the pressing need for advancements in faster imaging techniques to reduce MRI acquisition times while maintaining high image quality.

Significant research efforts have been dedicated to accelerating MRI. Central to these endeavors is the development of methods for image reconstruction from under-sampled data. Techniques based on parallel imaging (PI) introduced in the 1990s were an important watershed [3–8]. These techniques leverage the spatial diversity of multiple coil arrays to reconstruct images, thus allowing for reduced scan times by acquiring less data. The early 2000s saw the emergence of compressed sensing (CS) methods, which constituted a novel approach to MRI reconstruction [10–17]. CS overcomes the sparsity of MRI images by enabling the reconstruction of high-quality images from a much smaller set of measurements than traditionally required. The development of these techniques represented major milestones in MRI, and contributed to substantial reductions in scan duration and improvements in image quality.

However, both parallel imaging and CS methods have their limitations, each affecting their practicality in clinical settings. Parallel imaging is highly dependent on the geometry and sensitivity of the coil arrays, with suboptimal configurations leading to uneven image quality and potential artifacts. It also faces a practical limit on acceleration, beyond which significant noise can degrade the image quality. CS, on the other hand, is computationally demanding due to the complex optimization problems it involves, particularly with non-Cartesian trajectories, and often involve excessive compute times. It also relies heavily on hand-crafted priors for image reconstruction, which may not be applicable across different types of scans, thus limiting its adaptability. Furthermore, in low-rank CS MRI, the challenge of choosing an appropriate rank to balance image fidelity and computational efficiency often leads to a trade-off between reconstruction accuracy and speed. These challenges underscore the need for continued advancements in MRI technology to balance speed, image quality, and usability in clinical environments.

In recent years, the advent of deep learning (DL) has heralded a new era in MRI by offering promising solutions to these longstanding challenges. DL, a subset of machine learning characterized by algorithms based on computational neural networks, has had remarkable success in extracting complex patterns from large datasets [18–21]. In the realm

of MR image reconstruction, DL methods focus on learning from vast amounts of data to transform under-sampled or noisy data into high-fidelity images. These methods have demonstrated their ability to mitigate artifacts, enhance resolution, and accelerate the imaging process [22–31]. There are already multiple public datasets curated to enable training of DL models on MRI, e.g. [32–36], as well as community challenges related to MRI reconstruction problems [37–41].

This review provides a comprehensive overview of recent advances and applications of deep learning (DL) to the reconstruction of magnetic resonance (MR) images. Given the rapid pace at which this field is evolving, encapsulating the entirety of the published literature is a formidable challenge. Previous reviews have laid the groundwork by detailing the fundamental components of DL architectures and providing theoretical analysis [26–30, 42–49]. Here we cover a broad range of approaches, and highlight emerging methods such as self-supervised learning and diffusion models. Furthermore, we also review closely related topics, such as DL methods for *k-space* trajectory optimization, pulse sequence design, quantitative MRI, motion correction, and multi-task pipelines. Moreover, we address areas where DL encounters significant hurdles, including susceptibility to distribution shifts, instabilities, and inherent biases. Finally, building on our hands-on experience, we propose actionable strategies for effectively enhancing the robustness of DL models to such challenges.

Background on MRI reconstruction

In this section, we describe the image formation process and forward model, and discuss conventional, optimization-based, non-DL image reconstruction methods.

Image formation and forward model

The acquisition process in many imaging schemes can be modeled by an operator \mathbf{A} applied on the continuous domain image \mathbf{x} , where the process of collecting measurements is described by $\mathbf{y} = \mathbf{A}(\mathbf{x}) + \mathbf{n}$. In MRI, the measurement operator \mathbf{A} commonly corresponds to a multi-coil Fourier sampling operator. Although the acquisition is continuous, the general practice is to discretize the problem. Thus, we consider the reconstruction of an image vector \mathbf{x} from linear measurements, modeled by a matrix \mathbf{A} , by:

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

The above equation is a numerical model for the imaging device, and is often referred to as the *forward model*. In many imaging methods, the forward model is known

precisely. However, there are many applications where the forward model is unknown or only partially known. Examples include imaging in the presence of motion during the acquisition, trajectory errors, and field inhomogeneity effects in MRI acquisitions.

Due to the MRI’s long scan duration, many scans are accelerated by sampling k -space at a sub-Nyquist rate. In these cases, the forward model \mathbf{A} is often rank-deficient, making the recovery of \mathbf{x} an ill-posed problem. In MRI, the forward map \mathbf{A} includes undersampling, the Fourier transform, and sensitivity maps.

Conventional model-based image recovery

When the recovery of \mathbf{x} is ill-posed, many MRI schemes such as SENSE [5] pose the recovery as an optimization problem $\mathbf{x} = \arg \min_{\mathbf{x}} \mathcal{L}(\mathbf{x})$ with an objective function

$$\mathcal{L}(\mathbf{x}) = \underbrace{\|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2}_{\text{data consistency}} + \lambda \underbrace{\mathcal{R}(\mathbf{x})}_{\text{regularization}}, \tag{2}$$

where the first term is often called a *data consistency* term, and the second term is called a *regularization* prior. The objective (2) is sometimes called a *variational objective* [24].

The prior $\mathcal{R} : \mathbb{C}^n \rightarrow \mathbb{R}_+$ is used to restrict the solutions to the space of desirable images. The prior $\mathcal{R}(\mathbf{x})$ has a large value when \mathbf{x} is an undesirable image and is small for a desirable image. A common prior used in compressive sensing methods is wavelet-domain sparsity, where the number of non-zero wavelet coefficients or their surrogates are used as priors [10, 11]. In this case, the optimization algorithm facilitates the recovery of an image \mathbf{x} that has few non-zero wavelet coefficients.

From a Bayesian perspective, the above formulation can be viewed as an a-posterior estimate [28, 50], where the goal is to find an image \mathbf{x} that maximizes the posterior distribution $p(\mathbf{x}|\mathbf{y}) = p(\mathbf{y}|\mathbf{x}) * p(\mathbf{x})/p(\mathbf{y})$. The estimate is obtained by minimizing the negative log posterior

$$-\log p(\mathbf{x}|\mathbf{y}) = \underbrace{-\log p(\mathbf{y}|\mathbf{x})}_{\text{data consistency}} - \underbrace{\log p(\mathbf{x})}_{\text{prior}}. \tag{3}$$

Here, the first term is the data consistency term. It yields the mean-squared error in Eq. (1) if the noise vector \mathbf{n} has i.i.d. Gaussian entries. Data consistency terms appear in both compressed sensing [10] and DL methods [51], as they ensure that the reconstructed images adhere closely to the acquired data. The second term incorporates prior information on the images [50, 52].

Over the past few decades, substantial research efforts have been dedicated to crafting effective priors. Tikhonov regularization, for instance, employs a Gaussian prior on

\mathbf{x} , resulting in a regularization term $\mathcal{R}(\mathbf{x}) = \|\mathbf{x}\|^2$, and compressed sensing methods mentioned earlier promote sparsity.

Optimization algorithms

The loss in Eq. (2) is typically minimized by applying iterative first-order optimization algorithms such as gradient descent. Starting from \mathbf{x}_0 with a stepsize of η_t , iteration $t + 1$ of gradient descent is described by

$$\begin{aligned} \mathbf{x}_{t+1} &= \mathbf{x}_t - \eta_t \nabla \mathcal{L}(\mathbf{x}_t) \\ &= \mathbf{x}_t - \eta_t (\mathbf{A}^H (\mathbf{A}\mathbf{x}_t - \mathbf{y}) + \nabla \mathcal{R}(\mathbf{x}_t)). \end{aligned} \tag{4}$$

where H denotes a Hermitian transpose (i.e., conjugate transpose). The above algorithm depends on the gradient of the regularizer $\nabla \mathcal{R}$. When $\mathcal{R}(\mathbf{x}) = \log p(\mathbf{x})$ is the log-prior as mentioned earlier, $\nabla \mathcal{R}(\mathbf{x}) = \nabla \log p(\mathbf{x})$, which is often referred to as the *score* of the distribution. This term moves the estimate towards a signal with higher likelihood.

Other popular fast iterative algorithms for minimizing the objective (2) include the alternating direction method of multipliers (ADMM) [53] and the fast iterative shrinkage thresholding algorithm (FISTA) [54]. For example, the ADMM scheme considers the equivalent problem,

$$\mathbf{x} = \arg \min_{\mathbf{x}} \min_{\mathbf{v}} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \lambda \mathcal{R}(\mathbf{v}) \text{ such that } \mathbf{v} = \mathbf{x} \tag{5}$$

The above problem is solved by alternating between the following steps

$$\mathbf{x}_{t+1} = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \beta \|\mathbf{x} - (\mathbf{v}_t - \mathbf{u}_t)\|^2 \tag{6}$$

$$\mathbf{v}_{t+1} = \arg \min_{\mathbf{v}} \beta \|\mathbf{v}_t - \underbrace{(\mathbf{x}_{t+1} - \mathbf{u}_t)}_{\bar{\mathbf{x}}}\|^2 + \lambda \mathcal{R}(\mathbf{v}) \tag{7}$$

$$\mathbf{u}_{t+1} = \mathbf{u}_t + (\mathbf{x}_{t+1} - \mathbf{v}_{t+1}) \tag{8}$$

where \mathbf{u} is the Lagrange multiplier and \mathbf{v} is an auxiliary variable introduced in the ADMM algorithm to split the original optimization problem into smaller, more manageable sub-problems. The second step of the above optimization scheme

$$\mathbf{v} = \arg \min_{\mathbf{v}} \beta \|\mathbf{v} - \bar{\mathbf{x}}\|^2 + \lambda \mathcal{R}(\mathbf{v}) = \mathcal{D}_{\beta}(\bar{\mathbf{x}}) \tag{9}$$

can be viewed as a *denoising* step to clean the current solution $\bar{\mathbf{x}}$, thus yielding \mathbf{v} . For many penalties (e.g. ℓ_1 norm), the solution to (9) can be evaluated as proximal mapping. Here, β is a continuation parameter that can be interpreted as $1/\sigma_{\beta}^2$, where σ_{β} is the variance of *noise* in $\bar{\mathbf{x}}$ that decreases with the iterations. The first step (67) involves an inversion step to reduce the cost function composed of the linear combination of the data consistency error and the deviation from

the *denoised* image \mathbf{v} . This provides an iterative denoising interpretation, which is used in plug-and-play algorithms (discussed below, in “Pretrained plug-and-play (PnP) methods”). One challenge with the above convex optimization schemes is their high computational complexity, which is due to the numerous iterations required for convergence. In particular, the data consistency step (67) involves the evaluation of the forward model and its adjoint, which is often computationally expensive.

DL reconstruction: approaches and architectures

In this section we provide an overview of the main approaches and architectures in the MRI reconstruction landscape. While many different methods are available, and those often incorporate elements from other techniques, we classify them into five main categories: (i) neural networks trained end-to-end; (ii) approaches based on pre-trained denoisers, often called *plug-and-play (PnP)* methods; (iii) approaches based on generative models; and (iv) un-trained methods; (v) self-supervised methods. Additionally, we identify several recent architectures, e.g. transformers and dual-domain networks, that are used in various classes.

Interestingly, the top performing models in both the 2020 FastMRI challenge and the 2024 CMRxRecon challenge all used neural networks trained end-to-end [41, 55]. However, relative comparisons of algorithms depend on the problem

setup and metrics. In addition, the field is progressing rapidly, so that new and diverse benchmarking studies would be valuable (Fig. 1).

Neural networks trained end-to-end

Neural networks trained end-to-end are commonly trained to map the acquired data, which is often noisy and degraded by undersampling artifacts, to a target, ground-truth image. Their training hence commonly requires such paired data.

Let f_θ be a neural network, which receives the measurements as input and produces clean, reconstructed images as output. Given a training set consisting of pairs of measurements and target images $\{(\mathbf{y}_1, \mathbf{x}_1), \dots, (\mathbf{y}_n, \mathbf{x}_n)\}$, the network is trained by minimizing the loss between the prediction of the network and the target images, i.e.,

$$\mathcal{L}(\theta) = \sum_{i=1}^n \text{loss}(f_\theta(\mathbf{y}_i), \mathbf{x}_i). \quad (10)$$

Networks mapping a noisy measurement to a clean image

Many architectures have been developed for mapping an image with noise and undersampling artifacts to a clean image. Here we review some of the most well-known approaches. One of the early architectures is known as AUTOMAP (Automated Transform by Manifold




Fig. 1 Overview of the DL-based MRI reconstruction landscape. While many different methods are available, and those often incorporate elements from other techniques, we classify them into five main categories

Approximation) [22]. It utilizes a fully connected network followed by a convolutional network as the architecture f_θ in Eq. (10). This architecture does not incorporate the known forward model directly. Instead, it maps data from k -space to image domain, and learns the forward model from the data. A representative AUTOMAP application is shown in Fig. 2.

Other architectures commonly utilize image-to-image neural networks and incorporate the forward map. This class of approaches maps a coarse reconstruction of the image, for example an image generated from the zero-filled k -space, to a target image. In the notation above, the architecture f_θ consists of a linear map computing the zero-filled image followed by application of an image-to-image neural network.

The architecture of the image-to-image network is most commonly a convolutional neural network (CNNs). One of the pioneering works in this area was by Wang et al. [57], which demonstrated that using a CNN as the image-to-image network enables substantial improvements in both speed and image quality. Another pioneering work was by Jin et al. [58], which showed that this approach is applicable to a wide range of linear forward models.

Additionally, numerous other studies have embedded UNET, ResNet or recurrent neural networks as the backbone architecture [26–30, 42–45]. More recently, vision transformers have been utilized instead of a CNN as the image-to-image network. Several studies demonstrated that transformers can provide improvements [59–61]. However, they are computationally more expensive. In practice, a UNET architecture is often chosen as a simple starting point, as it provides a good trade-off between image quality and computational performance. However, the choice of

architecture depends on the problem at hand, the database size, the acceleration rate, and the desired reconstruction accuracy.

Architectures of unrolled networks

Currently, some of the best-performing neural networks are based on unrolled architectures [55]. These networks are obtained by unrolling an iterative algorithm such as gradient descent. The idea of unrolled networks was first introduced by [62], and several pioneering works applied it in the context of MRI reconstruction [24, 25, 63, 64]. These architectures iterate between two types of blocks: (i) data-consistency blocks, which can be computed using different algorithms [51], and (ii) blocks that remove noise and artifacts, which are commonly implemented by a deep neural network.

One of the early works in this context was by Hammernik et al. [24], who introduced the *variational network* (Fig. 3). This approach relies on a gradient descent algorithm to minimize the variational objective (2) where the regularizer \mathcal{R} is taken as the total-variation norm. In this case, the gradient of the regularizer in the gradient descent iterations (4) takes the form of a convolution. Thus, the gradient descent iterations can be interpreted as a neural network that applies data consistency operations (originating from the gradient of the least-squares loss) and the application of a convolutional network (originating from the gradient of the regularizer). Motivated by this observation, these so-called *variational networks* initialize $\mathbf{x}_0 = \mathbf{A}^H \mathbf{y}$ and then perform the following computations with a neural network:




Fig. 2 Elimination of hardware artifacts at low field MRI (6.5 mT) using AUTOMAP. Two slices from a 3D bSSFP (NA = 50) are shown. When reconstructed with IFFT (a, b), a vertical artifact (red arrows) is present across slices. When the same raw data was reconstructed with AUTOMAP (c, d), the artifacts are eliminated. The

error maps of each slice with respect to a reference scan (NA = 100) is shown for both IFFT and AUTOMAP reconstruction. e–g Uncorrupted k -space (NA = 50) was reconstructed with AUTOMAP (e) and IFFT (f). Adapted and modified from Koonjoo, N. et al. *Sci Rep* 11, 8248 (2021). <https://doi.org/10.1038/s41598-021-87482-7> [56]




Fig. 3 Variational network (VN) training procedure. The objective is learning a set of VN parameters during an offline training procedure. For this purpose, the current reconstruction of the VN is compared to an artifact-free reference using a similarity measure. This yields the

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \eta_t \mathbf{A}^H (\mathbf{A} \mathbf{x}_t - \mathbf{y}) + \text{CNN}_t(\mathbf{x}_t). \quad (11)$$

Here, both the parameter η_t and the parameters of the CNN are learnable. The original variational network [24] used a relatively shallow CNN, inspired by the parameterization provided by the total variation norm and its generalizations; specifically, the fields-of-experts-model. This yields a well-performing network with very few parameters. Later work has also shown that using a UNET within the unrolled network can improve the overall performance. [25, 65].

Other unrolled methods adopted alternate algorithms to minimize the variational loss (4) (described in “[Optimization algorithms](#)”). Those replaced the CNN in Eq. (11) with other image-to-image architectures; see [25, 63, 64, 66, 67] for a few examples. Furthermore, other architectures replaced the image-domain CNN with either a *k-space* CNN, a dual-domain (*k-space* and image domain) network [68–70] or a transformer [66].

Computational considerations The unrolling step requires multiple physical realizations of the CNN block during training, which translates into a high memory demand during training. This restricts their application in higher dimensional (e.g. 3D, 4D) applications. Programming solutions such as gradient check-pointing are now available to reduce the memory demand, at the expense of increased computational complexity. An alternative approach relies on deep equilibrium models [71, 72]. These models use a single CNN block and iterate the steps (67) and (8) until convergence to a fixed point,

reconstruction error which is propagated back to the VN to compute a new set of parameters. Reproduced with permission from Hammernik, K. et al. (2018) *Magn. Reson. Med.*, 79: 3055–3071. <https://doi.org/10.1002/mrm.26977>. [24]

similar to PnP methods. These methods then implement the fixed point iterations for back propagation. These methods thus enable the evaluation of the forward and back propagation using a single physical CNN block, thus reducing the memory demand. The MOL [73] method also imposes a local Lipschitz constraint on the CNN block, which offers theoretical guarantees and robustness without sacrificing performance.

Pretrained plug-and-play (PnP) methods

Early CS methods relied on convex priors $\mathcal{R}(\mathbf{x})$, such as the total-variation norm. Plug-and-play (PnP) methods make it possible to solve inverse problems with pre-trained denoisers. Another benefit is that they work with arbitrary forward models, where the prior incorporates information about the image.

ADMM and FISTA are iterative optimization methods for solving the a regularized least-squares problem and involve evaluations of the proximal operator proximal operator \mathcal{D}_β in (9) of the regularizer. One class of PnP methods replaces the proximal operator with a pre-trained denoiser; two well-known examples are PnP-ADMM and PnP-FISTA [74, 75]. While early methods relied on off-the-shelf image denoisers such as BM3D [76], pre-trained CNN denoisers are now considered to be more effective [74, 75, 77, 78]. Note that the proximal step in (9)

$$\mathcal{D}_{1/2\sigma^2}(\bar{\mathbf{x}}) = \arg \min_{\mathbf{v}} \frac{1}{2\sigma^2} \|\mathbf{v} - \bar{\mathbf{x}}\|^2 + \lambda \mathcal{R}(\mathbf{v}) \quad (12)$$

can be seen as the maximum a-posteriori (MAP) estimate of $\bar{\mathbf{x}}$ from its noise corrupted measurements

$$\mathbf{v} = \bar{\mathbf{x}} + \sigma \mathbf{n}. \quad (13)$$

Here, $\mathbf{n} \sim \mathcal{N}(0, \mathbf{I})$ is a sample from a Gaussian distribution with variance σ^2 . The CNN modules are hence pre-learned from training data as MAP denoisers, where noise-corrupted images are fed as input and the model is trained to yield noise-free images. During inference, steps (6)-(7) and (8) are iterated until the algorithm converges to a fixed point. Similar to CS methods, several iterations are often needed for convergence, which translates into higher computational complexity than the unrolled approaches described in “Pre-trained plug-and-play (PnP) methods”.

Another PnP framework is known as *regularization by denoising (RED)* [79]. This framework is more general than the above because it does not rely on any specific optimization algorithm, i.e. it enables using other methods, not only ADMM and FISTA. Furthermore, it offers great flexibility in choosing the denoising algorithm, as it enables incorporating almost any denoiser. Further information can be found in recent reviews of PnP methods [74, 75].

Generative priors

Another successful approach to DL-based MRI reconstruction is to learn an image prior parameterized by a generative neural network. Several major classes of generative methods have emerged, based on variational autoencoders [80, 81], Generative Adversarial Networks (GANs) [82–88], and very recently, diffusion models [89–95]. Here we focus on the two latter ones, which have attracted substantial attention.

One of the major advantages of generative approaches for image reconstruction is that they are flexible with regard to changes of the forward model, and at the same time perform well for reconstructing high-quality images from undersampled data. Furthermore, their probabilistic nature provides measures for uncertainty quantification, which is highly important for clinical imaging [96–98].

GANs

GANs [82] are a framework for generative modeling. A GAN consists of two competing neural networks: a generator, which aims to produce data indistinguishable from a given dataset of real images, and a discriminator, whose role is to distinguish between the generator’s output and the real data. GANs are trained using an adversarial loss [82]; this process enables the generator to learn to generate high-quality realistic images. After training, the generator can be used either to generate images that look similar to those in the training set, or as a prior for image reconstruction.

In the context of MRI reconstruction, GANs have attracted substantial attention over the last few years [83–88]. For example, DAGAN (Deep De-Aliasing Generative Adversarial Networks) [83] was a pioneering work that proposed a conditional GAN with a refinement-learning stage, and used a loss function comprised of an adversarial and a perceptual component. Mardani et al. [99] proposed a reconstruction framework where GANs were used for learning the low-dimensional manifold that underlies high-quality MR images. However, images generated by the generator are not necessarily consistent with the acquired measurements. To ensure such consistency, they included an affine projection operation, conducted by a layer placed between the generator and discriminator. Another approach for tackling this was proposed by Quan et al. [85], who introduced a novel cyclic loss in their GAN architecture to enforce data consistency. These methods, and many others [87, 100] showcased the potential of GANs to produce clinically viable MRI reconstructions.

Diffusion models

Diffusion models, a class of generative models that have garnered substantial attention in recent years, are making an impact in a variety of fields, including MRI reconstruction [89–95, 101]. These models operate by learning to reverse a diffusion process that gradually transforms random noise into structured images, and have shown a remarkable capability to generate high-quality, detailed images.

Diffusion models have been derived using different approaches [89], including discretized corruptions, e.g., denoising diffusion probabilistic models (DDPMs) [93], denoising score matching [102], and continuous formulations based on stochastic differential equations (SDEs) [103].

For a general probability density function $p(x)$, these approaches approximate the *score function*, defined by $\nabla_x \log p(x)$, using a neural network $s_\theta(x)$. To do so, the network is used to approximate a series of conditional score functions, $s_t(x_t) = \nabla_{x_t} \log p(x_t|x_{t+1})$, which guide the denoising process from pure noise, i.e., x_T drawn from a normal Gaussian distribution for some maximum iteration value T , to a clean sample $x_0 \sim p(x)$. Once trained, these models can be used to sample unconditionally from the prior distribution by running the reverse diffusion process, and hence generate new samples.

In the context of inverse problems in general, and MRI reconstruction in particular, the diffusion process can be hijacked to approximately sample from the conditional posterior distribution, $p(y|x)$ instead. One method involves conditioning on the k -space measurements y and applying Bayes’ rule to the series of score functions, i.e.,

$$\nabla \log p(x_t|y, x_{t+1}) = \nabla \log p(y|x_t) + \nabla \log p(x_t|x_{t+1}). \quad (14)$$

The second term, corresponding to the prior conditioned on the denoising process, is unchanged from the original diffusion model and can be learned by training on clean, fully sampled images. The first term, corresponding to the likelihood conditioned on the denoising process, can be approximated through various approaches [93, 94, 104]. In a naive approximation,

$$\nabla \log p(y|x_t) \approx A^H(Ax_t - y), \quad (15)$$

given the MRI forward model.

A growing body of work demonstrates that diffusion models work well for accelerated MRI and exhibit flexibility when handling various sampling patterns [94–96, 98, 105, 106]. For example, in a pioneering work, Jalal et al. [107] demonstrated that training a score-based generative model using Langevin dynamics, without making any assumptions on the measurement system, could yield competitive reconstruction results for both in-distribution and out-of-distribution data. Chung et al. [96] demonstrated that score-based diffusion models trained solely on magnitude images can be utilized for reconstructing complex-valued data. Luo et al. [97] described a comprehensive approach using data-driven Markov chains for MRI reconstruction which not only facilitates efficient image reconstruction across variable sampling schemes, but also enables the generation of uncertainty maps.

The flexibility afforded by explicitly decoupling the image prior (which is learned with diffusion models) and the statistical measurement model has also enabled other extensions. These include incorporating errors into the forward model, e.g., due to motion [108] and field inhomogeneity [109] and incorporating multiple image contrasts [110].

Un-trained neural networks

Un-trained methods are DL models that do not rely on training data apart from hyper-parameter tuning. Instead of conventional training on large datasets, these methods are typically based on fitting a randomly initialized neural network to a specific measurement. Here we discuss two types of methods: un-trained neural networks based on CNNs, and methods based on coordinate-wise implicit neural networks.

Un-trained CNNs for single image recovery

CNNs can be used as an image prior by fitting a randomly initialized CNN with gradient descent to a measurement. This approach, termed the deep image prior (DIP), was introduced in a pioneering work by Ulyanov et al. [111]. The optimization problem is formulated by,

$$\mathbf{x} = \arg \min_{\theta} \|\mathbf{Ax} - \mathbf{y}\|^2 \text{ such that } \mathbf{x} = \mathcal{G}_{\theta}(\mathbf{z}), \quad (16)$$

where \mathcal{G}_{θ} is a CNN generator whose input \mathbf{z} is a noise vector drawn from some noise distribution. The optimization is performed using gradient descent or ADAM [112], starting with random initialization of the network weights, and early stopping is used for regularization. The image quality first improves with the number of iterations, and then degrades as the network begins to fit the measurement noise in \mathbf{y} . This behavior is caused by the implicit bias of CNN networks to natural images: when trained with gradient descent, CNNs fit the smooth images before the noise, as formalized in [113].

Un-trained networks perform very well for denoising [111, 114] and compressive sensing (e.g., accelerated MRI). These methods can provably denoise smooth signals [113] and can provably reconstruct undersampled smooth images [115]. Un-trained networks also work quite well for accelerated MRI; they provide significant improvement over sparsity-based methods for 2D accelerated MRI [116].

One key benefit of un-trained networks is that they do not need training data. However, this benefit comes at the expense of performance; the images produced by DIP methods are commonly not comparable to those from the pre-trained networks discussed above. In addition, DIP often suffers from longer run times compared to the unrolled and direct inversion approaches because of the need for ADAM or gradient descent optimization during reconstruction.

Un-trained CNNs for joint recovery of multiple images

Recently, the DIP framework was extended to dynamic imaging applications [117, 118] where the images in a time series are modeled as the output of a generator

$$\gamma_t(\mathbf{r}) = \mathcal{G}_{\theta}[\mathbf{z}_t]. \quad (17)$$

Unlike the fixed noisy input used in the original DIP work [111], here \mathbf{z}_t are low dimensional latent vectors at a specific time point t . \mathcal{G}_{θ} is a deep CNN generator, whose weights θ are independent of t . For example, in a free-breathing cardiac MRI, the images in the time series at a specific time t can be viewed as non-linear functions of cardiac and respiratory phases captured by \mathbf{z}_t . This model (17) can be viewed as a non-linear mapping/lifting from a low-dimensional subspace \mathbf{Z} to the image space. The low-dimensional nature of the latent vectors enables the exploitation of the non-local redundancies between images at different time points, thus facilitating the fusion of information between them as in [119, 120].

The network parameters θ and the latent variables \mathbf{z} are jointly optimized for by minimizing the cost function

$$\mathcal{C}(\mathbf{z}, \theta) = \sum_{t=1}^N \|\mathbf{A}_t \mathcal{G}_\theta[\mathbf{z}_t] - \mathbf{y}_t\|^2 + \lambda_1 \underbrace{\|\nabla_{\mathbf{z}} \mathcal{G}_\theta\|^2}_{\text{network regularization}} + \lambda_2 \underbrace{\mathcal{R}(\mathbf{z})}_{\text{latent regularization}}. \quad (18)$$

The network regularization is an ℓ_2 penalty on the weights θ , which was shown to minimize the need for early stopping and provide improved performance. The latent vector regularization term involves a smoothness regularization to capitalize on the temporal smoothness of the images in the time series.

The above approach can also be extended to 3D applications, where the joint alignment and recovery of data from different slices obtained using different acquisitions may differ in cardiac/respiratory motion. Different sets of latent vectors are used for different slices to account for differences in breathing patterns and cardiac motion. In this case, a Kullback–Leibler divergence term is used to encourage the latent vectors of all the slices to follow a zero-mean Gaussian distribution, thus facilitating the alignment of data from different slices.

Coordinate-based networks

Coordinate-based neural representations, also known as implicit neural representations (NeRF-type networks), have recently emerged as an efficient way to represent and work with images, 3D shapes, and other signals [121]. They are commonly used for representing scenes and performing view syntheses in vision [122, 123]. To represent a 2D or 3D object, these models map a coordinate input (e.g., (x, y) coordinates for 2D or (x, y, z) coordinates for 3D) to a pixel value, for example to a real number for a gray-scale image and to two real numbers for a complex-valued image.

Coordinate-based networks can be used in an analogous fashion to un-trained CNNs to reconstruct an image [124–129]. Specifically, they can replace the CNN in an un-trained network and can be fitted to measurement data. Networks with Fourier-feature input (like NeRF [122], SIREN [130], and Fourier Feature inputs [125]) impose a smoothness prior similar to the un-trained CNN discussed in the previous section.

Networks with Fourier-feature inputs take a coordinate input (e.g., an (x, y) -coordinate \mathbf{z}) and map it to a feature representation with the map $[\sin(\mathbf{C}_0 \mathbf{z}), \cos(\mathbf{C}_0 \mathbf{z})] \in \mathbb{R}^{2m}$, where \mathbf{C}_0 are parameters initialized randomly. The parameters \mathbf{C}_0 can be fixed or trainable. Those features are then mapped

with a standard MLP with trainable parameters to an output. If two coordinates are close, the Fourier features are close, and how close can be controlled by the scale (variance) of the initialization.

If used as an image prior, compared to un-trained CNNs, coordinate networks with Fourier features perform worse in terms of image quality. In the context of MRI reconstruction, coordinate networks have been shown to be useful for representing high-dimensional objects such as 3D volumes and scenarios with motion. For example, [124] used coordinate networks to perform cardiac MRI reconstruction by fitting a network to the k -space data. This can be computationally efficient since the undersampled k -space data is sparse. [131] also used coordinate networks for free-breathing cardiac MRI reconstruction, by fitting a coordinate in the image domain.

Self-supervised methods

Neural networks, such as the end-to-end networks discussed in “[Neural networks trained end-to-end](#)”, are usually trained in a supervised manner (see Eq. (10)). This requires pairs of measurement and target (ground-truth) images. However, in practice, such pairs cannot always be acquired, e.g., due to scan time constraints, signal decay effects along echo trains, or physiological motion. Therefore, self-supervised methods are attracting increased research interest. These methods make it possible to train networks without target or ground-truth data by either making assumptions on the measurements or using additional noisy or partial measurements. A plethora of approaches has been developed, including methods for learning from under-sampled data [132, 133], unpaired data [134], or limited-resolution data [135]. Here we describe several approaches that are architecture-agnostic. For recent reviews on this topic see [134, 136, 137].

Learning of algorithms based on Stein’s Unbiased Risk Estimate (SURE)

We start with a method that is based on assumptions on the noise distribution, called Stein’s Unbiased Risk Estimate (SURE) [138]. We consider the estimation of \mathbf{x} , denoted by $\hat{\mathbf{x}}$ from its noisy measurements $\mathbf{v} = \mathbf{x} + \mathbf{n}$. Here \mathbf{n} is zero-mean Gaussian noise with a variance of σ^2 . In practice, the estimate $\hat{\mathbf{x}}$ is derived from the noisy measurements \mathbf{v} using a deep network as $\hat{\mathbf{x}} = f_\theta(\mathbf{v})$. When the noiseless reference image \mathbf{x} is available, the true mean-square error (MSE), denoted by

$$\text{MSE} = \mathbb{E}_{\mathbf{x}} \|\hat{\mathbf{x}} - \mathbf{x}\|^2 \quad (19)$$

can be used.

By contrast, the SURE [138] approach uses the loss function

$$\text{SURE}(f_{\theta}(\mathbf{v}), \mathbf{v}) = \|f_{\theta}(\mathbf{v}) - \mathbf{v}\|_2^2 + 2\sigma^2 \nabla_{\mathbf{v}} \cdot f_{\theta}(\mathbf{v}) - N\sigma^2, \quad (20)$$

which is an unbiased estimate of (19). Note that the expression in (20) does not depend on the noise-free images \mathbf{x} ; it only depends on the noisy images \mathbf{v} and the network parameters θ . In (20), $\nabla_{\mathbf{u}} \cdot f_{\theta}(\mathbf{u})$ represents the network divergence, which is often estimated using Monte-Carlo simulations [139]. Several researchers have adapted SURE as a loss function for the unsupervised training of deep image denoisers [140, 141] and demonstrated performance approaching that of supervised methods.

The SURE approach was extended to inverse problems with a rank-deficient measurement operator known as the generalized SURE (GSURE) [142]. The GSURE provides an unbiased estimate of the projected MSE, which is the expected error of the projections in the range space of the measurement operator. The GSURE approach was recently used for inverse problems in [140]. The experiments in [140] showed that the GSURE-based projected MSE was a poor approximation of the actual MSE in the highly undersampled setting. To improve performance, the authors trained the denoisers at each iteration in a message-passing algorithm in a layer-by-layer fashion using classical SURE, which was termed LDAMP-SURE [140]. This approach approximates the residual aliasing errors at each iteration to be Gaussian random noise. As this assumption is violated in many inverse problems, the performance of this layer-by-layer training approach is not as good as supervised methods.

The ENSURE framework circumvents the poor approximation of the true MSE by GSURE by considering different sampling operators for different images. Similar to classical SURE metrics [142, 143], the ENSURE loss metric has a data consistency term and a divergence term. The data consistency term in ENSURE is the sum of the weighted projected losses [142] from multiple subjects; the weighting depends on the class of sampling operators. When different sampling patterns from different subjects fully cover k -space, the ENSURE metric is an unbiased estimate of the true image-domain MSE and hence is a superior loss function than projected SURE [142]. The comparison of the above methods shows that the ENSURE approach can provide performance comparable to that of supervised training.

Self-supervised DL based on Noise2noise

Noise2noise [144] is a well-established framework, which constructs a self-supervised loss based on independent noisy

measurements of the same object. Recall that for single-coil accelerated MRI, the forward map is $\mathbf{A} = \mathbf{M}\mathbf{F}$, where \mathbf{M} is an undersampling mask and \mathbf{F} is the Fourier transform.

Suppose we are given two measurements $\mathbf{y} = \mathbf{M}\mathbf{F}\mathbf{x}$ and $\mathbf{y}' = \mathbf{M}'\mathbf{F}\mathbf{x}$, where \mathbf{M} and \mathbf{M}' are two different random undersampling masks. From these measurements, we can construct the self-supervised loss

$$\ell_{\text{SS}}(f_{\theta}(\mathbf{y}), \mathbf{y}') = \|\mathbf{M}'\mathbf{F}f_{\theta}(\mathbf{y}) - \mathbf{y}'\|^2. \quad (21)$$

It can be shown that in expectation over the random measurements, a minimizer of the self-supervised loss is also a minimizer of the expectation of the supervised loss (see Prop. 2 in [145]). Thus, with enough training examples, such self-supervised training can approach the performance of supervised training [145].

One notable method that has implemented this approach successfully for MRI reconstruction is Self-Supervised Learning via Data Undersampling (SSDU) [132]. This method partitions the available k -space measurements into two disjoint sets; the first set is used in the data consistency units of the unrolled network, i.e., for the forward pass, and the other one is used for computing the loss, i.e., for supervision. SSDU can hence be trained using under-sampled data alone. In their work, Yaman et al. [132] demonstrated that SSDU achieved comparable performance to fully supervised learning methods while offering practical advantages in real-world MRI applications.

Recently, Millard and Chiew [146] introduced a general theoretical framework that extends Noiser2Noise [147] and also explains SSDU. Unlike the SSDU formulation, where one set is recovered from the other, they applied two sub-sampling masks to the data. They proposed a weighted ℓ_2 loss, computed in k -space, with a weighting that compensates for the sampling and sampling-partitioning densities. They derived the framework analytically and showed that when the weighting matrix W is rank-deficient and fulfils certain conditions, the method boils down to SSDU. They showed analytically that SSDU with an ℓ_2 k -space loss approximates fully sampled reconstruction, on expectation. It is worth mentioning that their analysis was done for an ℓ_2 k -space loss, while the original SSDU method was trained with a mixed ℓ_1/ℓ_2 loss.

Self-supervised DL using k -space bands

The self-supervised methods described above focused on learning from under-sampled data acquired with variable-density or parallel-imaging schemes. Although such datasets have undersampling artifacts, they effectively constitute high-resolution data, because the sampling masks commonly cover the entire k -space extent (note that undersampling creates artifacts but does not necessarily reduce

the resolution). However, the acquisition of high-resolution data can be challenging. In dynamic MRI, for example, there is often a trade-off between the spatial and temporal resolutions, which requires acquisition compromises.

Recently, the *k-band* framework was proposed for self-supervised learning from partial, *limited-resolution* data [135]. This framework is based on the acquisition of *k-space* bands, where each band acquires data with high resolution in the MRI readout dimension and limited resolution in the phase encoding (PE) dimension. The authors suggested acquiring different bands from different subjects, and randomizing the bands' orientation across subjects; fundamentally, this randomization serves to expose the network to all *k-space* areas across the training iterations (Fig. 4). Thus, even though the network does not get a full *k-space* from any single subject, it can learn connections across all the *k-space* regions. To enable self-supervised learning from limited-resolution data without limiting the resolution during inference, the authors introduced an optimization method dubbed *stochastic gradient descent (SGD) over k-space subsets*.

In this framework, the loss is computed in *k-space* and formulated by

$$\ell_{k_band} = \|\mathbf{WB}(\mathbf{F}f_{\theta}(\mathbf{y}) - \mathbf{F}\mathbf{x})\|_1 \quad (22)$$

where $\mathbf{B} \in \{\mathbf{B}_i\}_{i=1,\dots,180}$ is a binary band sampling operator that samples a band with angle i , and \mathbf{W} is a loss weighting mask

$$\mathbf{W} = 180 \left(\sum_{i=1}^{180} \mathbf{B}_i \right)^{-1}. \quad (23)$$

This loss-weighting compensates for the over-exposure of the network to low-frequency *k-space* data and enhances learning in the *k-space* periphery. This is beneficial because in the *k-band* acquisition setting, the center of *k-space* is included in all bands (Fig. 4), unlike the periphery. The authors showed analytically that when this loss-weighting mask is applied, the self-supervised training process stochastically approximates fully supervised training, on expectation. They demonstrated that learning from limited-resolution data can hence result in performance comparable to supervised and self-supervised methods trained on high-resolution data, and hence offers a practical solution for cases where such data are unavailable.

Loss-weighting

Several recent studies have independently proposed applying *spatially varying weighting* to *k-space* loss functions and demonstrated that such weighting can improve the performance of self-supervised DL methods [133, 135]. Interestingly, this general concept emerged even though the studies analyzed different sampling schemes and loss functions. For example, Millard and Chiew [133] analyzed SSDU with variable-density sampling masks that cover the entire *k-space* area, and an ℓ_2 *k-space* loss function. In contrast, the authors of *k-band* [135] explored training on band-limited data using an ℓ_1 *k-space* loss function. Despite these differences, these studies arrived at similar conclusions: they derived loss-weighting masks that weigh down the loss in the center of *k-space* and enhance it in the periphery. These masks hence inhibit the learning of low-frequency data and facilitate learning of high-frequency




Fig. 4 Example of the input *training* data for three DL reconstruction methods. The fully-supervised MoDL method [25] receives var-dens sampled data as input and uses the entire *k-space* for supervision. The self-supervised SSDU method [132] receives var-dens data as input, splits it into two subsets, and uses one set for data consistency and the other for supervision. In this example, the var-dens data were sampled from parallel-imaging (equispaced) acquired data, as in [132]. The

k-band method [135] receives var-dens sampled data from a *k-space* band, and uses data from the whole band for supervision, without any supervision outside the band. Different bands are acquired from different subjects, with random orientations. At *inference*, the input to all three methods is var-dens data from the entire *k-space*, similar to that shown here for MoDL

details, so that eventually all frequencies are weighted equally.

A related concept was proposed by Huang et al., who developed a neural implicit k -space representation model for cardiac MRI [124]. To account for the large variations in k -space values they proposed a log transform that inhibits high-magnitude k -space data to make their magnitude values similar to those of low-magnitude data. However, because such a non-linear transform has an undesired effect on the noise distribution, the authors proposed an approximation using a linear function [148]. Altogether, this is an alternative approach for balancing the different parts of k -space.

Recent architectures

Transformers and dual-domain networks

In addition to the training methods described above, much progress has also been made in the development of advanced architectures. For example, two architectures that recently garnered substantial attention are transformers and dual-domain networks. Transformers [149, 150] have powerful computational capabilities due to their use of an attention mechanism [151] that makes it possible to weigh the importance of different parts of the input data and capture long-range dependencies. Transformers first made a substantial impact in the field of natural language processing [152, 153] and then became highly influential in computer vision [149].

In the context of MRI reconstruction, recent studies demonstrated that transformers offer excellent performance and ability to deliver improved structural and textural fidelity. For example, Korkmaz et al. [154] developed an unsupervised MRI reconstruction method based on a generative vision transformer. Their method utilizes cross-attention transformer blocks, which receive both global and local latent variables as input and progressively map them to MR images with increasing spatial resolution. This style-generative architecture enhances representational learning and improves model invertibility. Feng et al. [155] introduced the T^2 Net for simultaneous MRI reconstruction and super-resolution. This network has two branches dedicated to these two tasks, and incorporates a task transformer module to facilitate effective feature sharing between them. Guo et al. [61] introduced the ReconFormer, an architecture that leverages recurrent pyramid transformer layers and scale-wise attention mechanisms. It effectively captures multi-scale information and deep feature correlations, leading to efficient, high-quality image reconstruction and computational efficiency.

Another emerging type of architecture is known as dual-domain networks, which commonly integrate information from the image and k -space domains [69, 156–159]. This approach, exemplified by MD-Recon-Net [158], leverages

the complementary strengths of these two domains to enhance reconstruction quality. A study by Souza et al. [156] demonstrated the effectiveness of such networks in multi-channel MRI reconstructions. Singh et al. [159] demonstrated that layers utilizing joint learning of image and frequency domain features can directly replace standard convolutional layers. This is useful for numerous tasks, including image reconstruction, motion correction and denoising.

Transformers and dual-domain networks have recently been integrated, leading to state-of-the-art architectures. For example, Zhao et al. [160] introduced SwinGAN, a dual-domain Swin Transformer-based GAN. This network combines frequency-domain and image-domain generators, both utilizing Swin Transformer backbones. This design allows for effective capture of long-distance dependencies in MR images. SwinGAN also features a contextual image relative position encoder, which enhances its ability to capture local information. Wang et al. introduced DCT-Net, a dual-domain transformer network for MRI reconstruction [70], which integrates image and frequency domain information through its cross-attention and fusion-attention blocks. DCT-Net is designed to enhance MRI reconstruction performance, particularly under low sampling rates, by leveraging the complementary strengths of both domains. In summary, these recent architectures offer high computational power to improve image reconstruction quality.

Recent architectures incorporating diffusion models

Recently, some of the architectures mentioned earlier have been integrated with diffusion models, yielding state-of-the-art methods. For example, Korkmaz et al. [161] introduced the Self-Supervised Diffusion Reconstruction (SSDiffRecon) method, which poses a diffusion model as an unrolled network, with interleaved cross-attention transformer blocks and physics-driven data-consistency steps. Furthermore, Zhao et al. [162] introduced DiffGan, an architecture that combines a local vision transformer with a diffusion model, which mitigates computational challenges in training generative models.

DL for acquisition optimization

The previous chapter discussed various *reconstruction* approaches for retaining data integrity and accuracy when k -space sub-sampling takes place. This chapter will highlight two complementary deep-learning-based interventions at the *acquisition* step that further assure optimized performance, while suggesting additional acceleration. First, we explore methods that optimize the k -space sampling trajectories in tandem with the reconstruction. Next, we describe

recent advances in harnessing DL to design and refine MRI pulse sequences.

Optimizing k -space trajectories

The computational design of sampling patterns has a long history in MRI. Generally, two types of approaches have been taken. *Algorithm-agnostic* methods, e.g. [163–167], consider specific image properties (e.g., the Cramér-Rao bound or image support) and optimize the sampling pattern to improve the measurement diversity for that class. *Algorithm-dependent* methods, on the other hand, e.g. [168–171], optimize the sampling pattern assuming specific reconstruction algorithms. These are typically CS algorithms, which employ regularizers such as TV, wavelet-domain sparsity, or pre-trained diffusion models [171].

The main challenge with the above computational approaches is their high computational complexity. In particular, algorithm-dependent schemes need to solve the CS problem for each image in the dataset, to evaluate the loss for a specific sampling pattern. The design of sampling pattern thus involves a nested optimization strategy; the optimization of the sampling patterns is performed in an outer loop, while image recovery is performed in the inner loop to evaluate the cost associated with the sampling pattern.

DL provides an opportunity to speed up the computational design, because DL inference schemes enable rapid evaluation of the loss for each sampling pattern. This enables a joint strategy that simultaneously optimizes the acquisition scheme and the reconstruction algorithm. Early DL-based joint optimization schemes solved for a binary sampling mask [172, 173]. The PILOT method, for example, solves for a hardware-constrained k -space trajectory [172]. The LOUPE method, on the other hand, learns the optimal sampling density in tandem with a reconstruction network [173]. It was first developed for 2D Cartesian imaging [173] and later extended to non-Cartesian sampling [174]. Other studies have focused on 3D Cartesian sampling with a variational reconstruction network [175].

More recent work represents the sampling locations ϕ as continuous variables and jointly solves for them and for the parameters of the DL algorithms. These methods consider a forward model \mathbf{A}_ϕ , where ϕ denotes the sampling locations. This forward model may be represented either by using an analytical Fourier transform [176] or a non-uniform Fourier transform [177]. We denote the reconstruction algorithm (which can be unrolled, direct inversion, or plug-and-play) by

$$\hat{\mathbf{x}} = \mathcal{M}_{\theta, \phi}(\mathbf{y}), \quad (24)$$

where θ denotes the parameters of the reconstruction algorithm and ϕ are the sampling locations corresponding to the forward model. Joint optimization schemes, e.g. [176, 177], are designed to optimize the sampling pattern Φ and the CNN parameters ϕ in tandem, i.e.,

$$\{\theta^*, \phi^*\} = \arg \min_{\theta, \phi} \sum_{i=1}^N \|\mathcal{M}_{\theta, \phi}(\mathbf{A}_\phi(\mathbf{x}_i)) - \mathbf{x}_i\|_2^2. \quad (25)$$

Several methods have been developed within this framework. For example, J-MODL focuses on a model-based reconstruction and utilizes an unrolled network [176]. In a different work, Wang et al. [177] parameterized trajectories with quadratic B-spline kernels, and performed optimization under penalties describing realistic MRI hardware constraints, e.g. the slew rate and gradient amplitude. This work was later extended to a generalized Stochastic optimization framework for 3D Non-Cartesian sampling trajectory (SNOPY) [178], which can accommodate several optimization objectives. Chaithya and Ciuciu [179] introduced the PROJeCTOR framework, which enables joint learning of non-Cartesian trajectories and reconstruction networks by using a projected gradient descent algorithm. Alkan et al. [180] introduced joint sampling and reconstruction optimization through variations in information maximization, where they used an encoder to represent non-uniform sampling and a decoder in an unrolled neural network. Xie et al. [181] introduced the PUERT method for learning probabilistic sampling patterns along with an interpretable reconstruction method; their learning module incorporated a dynamic gradient estimation strategy. Finally, Zou et al. [182] demonstrated that joint optimization can reduce the bias and uncertainty of pharmacokinetic parameter estimation in dynamic contrast enhanced MRI, and hence contribute to higher diagnostic value. Altogether, these methods have shown significant benefits from jointly optimizing the sampling pattern and the reconstruction algorithm.

Pulse sequence design

The previous section focused on accelerating MRI scans via k -space sub-sampling. The complementary element to this effort is the optimization of the remaining pulse sequence parameters by a set of radio-frequency (RF) powers, shapes, and duration that enable the shortest possible scan time, while retaining sufficient contrast, SNR, and consistency with conventional (and lengthy) alternatives.

The pulse sequence design task was traditionally hand-crafted by MR experts, who combined strong intuitions and an understanding of spin physics with mathematical solutions of the Bloch equations. While a remarkable number of contrast mechanisms and imaging schedules have been developed since the invention of MRI, the reliance on

solvable differential equations severely limits our ability to reach a globally optimized schedule and reduce the scan time. Recent developments in DL architectures and computational frameworks have created new opportunities for the automatic and efficient optimization of rapid acquisition protocols.

Zhou et al. [183, 184] introduced the representation of the Bloch equations as a computational graph. By treating each of the acquisition parameters as a neural network node weight, an efficient gradient-descent-based optimization was realized, where simulated signal trajectories were fed into the network, enabling an automatic generation of pulse sequences. The resulting protocols were characterized by non-intuitive gradient waveforms, where continuous off-resonant excitation applied as the receive channel was continuously and simultaneously recorded. This approach yielded an ultra-short scan time for T_1/T_2 mapping at 1D. By expanding for 2D imaging, Lee et al. [185] used automatic differentiation to optimize the Cramér-Rao Lower Bound (CRLB) of multiple-echo spin echo T_2 mapping, driven equilibrium single pulse observation of T_1 (DESPOT1) mapping, and the MRF IR-FISP sequence.

Loktyushin et al. [186] developed a supervised learning framework termed MR-Zero where a target contrast of interest is used for learning the optimal set of RF events, the gradient moment, and the delay times (Fig. 5). One important

feature of this approach is the use of a task-driven cost function that provides the user with the flexibility to prioritize the characteristics required from the output protocol, such as high data fidelity, short scan time, or the specific absorption rate (SAR) limits. In a later study, the same group used this approach to optimize the refocusing flip angles and minimize T_2 -induced blurring in accelerated spin echo sequences [187].

In the molecular MRI field, an end-to-end DL-based framework was developed for the discovery of rapid, quantitative chemical exchange saturation transfer (CEST), semisolid magnetization transfer acquisition and reconstruction protocols [188]. The system was based on a computational graph representation of the Bloch-McConnell analytical solution which receives the molecular imaging scenario of interest as input, and outputs an optimized set of acquisition parameters and the corresponding reconstruction network that translates the raw data into quantitative parameter maps. In vivo experiments showed it could acquire in-vivo data in merely 35 s and reconstruct parameter maps in less than 1 s. The use of recurrent neural networks and training over a wide range of saturation pulse frequency offsets has further increased the robustness of this conceptual approach for B_0 and B_1 inhomogeneity [189].

All these approaches exploit DL-based strategies to optimize and derive novel acquisition routes *offline*. Recently, a different optimization paradigm was suggested




Fig. 5 Automated discovery of MRI acquisition protocols using supervised learning. A differentiable MR scanner utilizes the Bloch equations for in-silico signal generation and the later reconstruction

of the target contrast of interest from real, acquired data. Reproduced from Loktyushin et al. *Magn. Reson. Med.* 2021; 86: 709-724 [186]

where the acquisition parameters are modified and adapted on the fly, *during* data acquisition [190]. By combining a Bayesian framework, CRLB calculation, and model-based reconstruction, the acquisition parameters for a series of images can be optimized in real-time based on the previous image history. This concept has demonstrated up to a 3.3 fold acceleration of multi-echo sequences in human subjects and molecular imaging phantoms.

As an intermediate conclusion, while the concept of machine-learning-based pulse sequence design is relatively young and not heavily explored, the first reports suggest a promising new avenue for optimizing image contrast, shortening the scan time, and finding new acquisition schemes beyond human intuition.

Advanced techniques and applications

In this section we discuss DL methods for quantitative MRI and dynamic MRI.

DL methods for quantitative MRI

The goal of quantitative MRI is to extract one or more tissue parameter maps from a series of qualitative images [191]:

$$\mathbf{I}_m = \Phi_m(T_{param})\rho \quad (26)$$

where \mathbf{I}_m denotes the contrast-weighted images for $m=1, \dots, M$ acquisitions, ρ denotes the spin density, T_{param} denotes the tissue parameters (T_1 , T_2 , etc.), and Φ_m is the biophysical function connecting the acquisition parameters with the resulting contrast-weighted images. The mapping of tissue properties enables de-biasing imaging protocols and harmonization of the final diagnosis across sites, vendors, and physicians. It thus provides sensitive and standardized tools for reproducible interpretation of MRI-based information [192]. The classical approach to MR property mapping mandates repeated acquisition, where all the protocol parameters are held fixed, and only a single parameter is slowly and gradually varied (e.g., the flip angle or the repetition time across different M acquisitions). The resulting long scan time hinders the widespread use of quantitative MRI in clinical settings [193]. The reconstruction of the acquired raw data series demands a lengthy parameter-fitting procedure that is computationally intensive and slow.

The development of powerful DL architectures such as CNNs, UNets, GANs, ResNets and recurrent neural networks has been leveraged to accelerate and enhance the performance of quantitative MRI [194]. To accelerate relaxometry studies, Liu et al. [195] developed a model-augmented neural network that receives a series of incoherently-sampled

multi-echo images and uses a CNN to reconstruct the T_2 parameter maps. The supervised learning was guided by a parameter-space loss, which compares the reconstructed T_2 maps to the ground truth reference, and a k -space loss. The latter was designed to ensure that the physical-model-based synthetic undersampled k -space measurements matched the originally acquired k -space information. In a later work, the same group developed a model-guided self-supervised DL framework for rapid T_1/T_2 mapping [196], to obviate the need for fully sampled training references.

While many DL-based quantitative mapping strategies are focused on k -space sub-sampling, a further acceleration potential lies in reducing the number of contrast-weighted images acquired. In a very recent work, Li et al. [197] trained a deep residual CNN network to receive just three k -space under-sampled contrast weighted images and output the corresponding $T_{1\rho}$ and T_2 parametric maps (which are particularly useful for the study of osteoarthritis).

Magnetic resonance fingerprinting (MRF), which was first reported in a 2013 Nature paper [198] and increasingly studied since then, constitutes a paradigm shift in MRI-based tissue characterization. Unlike traditional relaxometry studies, MRF starts with the acquisition of tens or hundreds of images, using a pseudo-random acquisition pattern accommodating a series of short repetition times, small flip angles, and heavily under-sampled k -space data (e.g., via a single variable density spiral trajectory). Although each of the resulting raw images is extremely noisy, the temporal evolution of the signal at each pixel entails a unique fingerprint. By comparing the experimental trajectory to a Bloch-equation-derived dictionary of simulated signals, the inverse problem can be solved to uncover the associated parameter maps (T_{param}).

While the resulting acquisition times are incredibly short (e.g., <13 s [198, 200]) the quantitative image reconstruction step (via pattern matching) may take hours, because the similarity between each acquired signal trajectory and *all* possible dictionary entries needs to be calculated.

In recent years, several DL-based strategies have been suggested to overcome this challenge. Cohen et al. [201] trained a fully connected neural network using synthetic signal dictionaries to reconstruct MRF data in less than 100 ms. To take advantage of the inherent dependencies between adjacent image pixels, Balsiger et al. [202] designed a spatiotemporal CNN where the time-evolution dimension is the third dimension of the CNN patch kernel. In-vivo human brain validation studies using this approach demonstrated improved performance compared to alternative MRF networks [201, 203]. To accommodate 3D imaging, Gomez et al. [204] combined fully connected reconstruction networks with radial and spiral readout trajectories, and achieved whole-brain reconstruction in less than 7 min (compared to > 1.5 h using traditional reconstruction).

Another MRF-associated bottleneck relates to the time required to generate the synthetic signal dictionary, which increases exponentially with the number of simulated parameters [205]. Even when high-end computer clusters are used for this task, the computation time may reach hours/days for complex multi-pool imaging [206, 207]. Recently, this challenge was addressed by training a fully connected neural network using a variety of dictionaries to learn the nonlinear relations embedded in the physical model. The resulting system enabled the rapid generation of simulated signals for various protocols and imaging scenarios [208]. NN-based simulators can be further combined with reconstruction networks to provide a unified rapid method for MRF analysis [209]. A different approach to circumvent the need for exhaustive dictionary generation involves the direct synthesis of multi-contrast images (e.g., T_1 -weighted, T_2 -weighted, and

FLAIR) from raw MRF data. While synthetic images can be derived from quantitative MRF data by forward model activation using the desired acquisition parameters [210, 211], Wang et al. [212] showed that a dictionary-free conditional GAN (trained on MRF raw data and paired ground truth weighted images) can perform the same task much faster. For cases where full quantitative information is required, a different work demonstrated that multi-parameter maps can still be extracted with GANs, even when merely 30% of the acquired MRF data is used [199] (Fig. 6).

Dynamic MRI

Deep learning has become a transformative force in the realm of dynamic MRI, particularly in addressing the challenges related to limited acquisitions and motion correction,




Fig. 6 A demonstration of two MRI quantification strategies/architectures. **a** Deep learning reconstruction of quantitative magnetic resonance fingerprinting (MRF) information. A fully connected neural network is trained using simulated signal trajectories. During inference, it receives a series of raw MRF images pixel-wise (gray-scale images, left), as well as auxiliary maps (color, top left), yielding

quantitative parameter maps (top right). **b** A further acceleration in quantitative MRI scan time can be achieved by training a generative adversarial network (GAN) using a smaller subset of raw input data to yield the same quantitative output maps. Reproduced and modified from Weigand-Whittier et al. [199]

which constitute substantial hurdles in clinical imaging [9, 45, 213, 214]. DL excels at learning signal evolution [157, 215], a critical factor when aiming to accurately visualize and interpret dynamic changes in the body.

Motion-resolved reconstruction

Motion-resolved algorithms can effectively learn spatio-temporal correlations and reconstruct images from highly undersampled sequential data [216, 217]. These methods have primarily been developed in the context of cardiac MRI [64, 120, 213, 216, 218, 219]. For instance, supervised unrolled algorithms have been used to recover cardiac cine MRI from breath-held MRI using using 4D (3D+Time spatial) convolutions [64, 213]. In other studies, architectures included unrolled algorithms that combine manifold [120] or low-rank priors [216, 218], and joint learning of motion estimation and segmentation in cardiac MRI [219]. From a clinical perspective, DL has been found useful for measuring myocardial displacement [220], noninvasive diagnosis of myocardial ischemia [221], and evaluation of cardiac function in pediatric imaging [222].

To tackle the scarcity of training data, *unsupervised* implicit learning approaches have recently been introduced for dynamic MRI [217, 223] (for more information see “Untrained CNNs for joint recovery of multiple images” and “Coordinate-based networks”). These methods were also extended to multi-slice dynamic MRI data [224], where the dynamic data from slice $\mathbf{x}_i(t)$ are acquired sequentially at different time points. The above model has been generalized to recover a pseudo-3D reconstruction by modeling the data as $\mathbf{x}_i(t) = \mathcal{G}_\theta(\mathbf{z}_i(t))$, where $\mathbf{z}_i(t)$ are allowed to vary for different slices.

Motion-compensated reconstruction

The development of DL techniques for motion estimation and correction is a highly active research field, as DL can accurately detect and compensate for both rigid and non-rigid motion artifacts, which leads to more diagnostically valuable images. A thorough review of motion estimation and correction techniques is beyond the scope of this manuscript. Here we highlight some of the main applications, and more information can be found in recent reviews [44, 45, 225–227].

One of the main applications where DL is highly effective for motion correction is brain MRI, which is characterized by rigid-body motion [226–229]. One of the early works in this field, by Johnson and Dragnova [230], proposed conditional GANs to infer clean images from motion-corrupted data. More recent techniques include co-optimization for

jointly estimating the motion parameters and reconstructed image [134, 228], methods for detection and correction of motion-corrupted *k-space* lines [229, 231], and the use of score-based generative models [232].

DL approaches are also highly useful for tackling *non-rigid*, irregular motion. Applications include imaging of the body trunk [233], fetal MRI [234, 235], abdominal MRI [236, 237], and MR angiography [238]. Unsupervised implicit learning methods that recover the deformable motion fields at each time point have also been introduced and found to be effective in motion-compensated recovery [239]. DL is also making strides in the field of real-time interventional MRI. Here, the rapid processing capabilities of DL algorithms enable real-time feedback and guidance during medical procedures, thus enhancing both the safety and efficacy of interventions. [240–242]

Multi-task pipelines

The fundamental goal motivating the acquisition of diagnostic-quality MR images is to extract clinically useful insights to further clinical care or to interrogate disease activity. Consequently, efficient, high-quality image acquisition is just the first step (typically referred to as *upstream DL*) in the imaging workflow, which is followed by image analysis and insight extraction (typically referred to as *downstream DL*) [243]. In many applications, these upstream and downstream processes are disconnected, leading to insufficient insights as to whether a novel MRI acceleration and reconstruction technique can reliably produce the requisite diagnostic information [244]. As a result, there is a substantial need to combine the upstream and downstream processes to ultimately harness advances in MRI physics, hardware, and DL for end-to-end acquisition-to-analysis workflows.

Conventional and DL-based reconstruction techniques can potentially be combined with downstream task of clinical utility to guide useful model development. Specifically, MRI reconstruction workflows can be combined with three different downstream tasks that use whole images as inputs: (i) *image classification*, which performs binary identification (via a yes/no) to identify the presence of one or multiple disorders; (ii) *abnormality detection*, which performs localization via bounding boxes to accurately depict where one or multiple disorders are present in images; (iii) *image segmentation*, which performs image classification at the voxel level to distinguish voxels that belong to a particular tissue or disease class. The DL sub-field of multi-task learning can learn multiple tasks simultaneously with *positive task transfer*, where learning one task improves the performance of other tasks.

In the context of accelerated MRI, combining the upstream task of MRI reconstruction with the downstream tasks of classification, detection, or segmentation can

improve performance on all tasks. It also contributes to optimizing reconstruction techniques with clinically informed metrics. One of the greatest challenges in doing so, however, is the lack of available datasets that can merge both sets of tasks. The fastMRI raw-data dataset was recently supplemented with the fastMRI+ dataset that includes classification and detection bounding box annotations for knee and brain abnormalities at the slice level [245]. Such datasets can enable the design of end-to-end techniques to optimize reconstruction, subject to high performance on lesion detection [246]. Similarly, even beyond end-to-end methods, such abnormality labels can be used to design clinical task-specific undersampling trajectories [247].

Beyond fastMRI+, SKM-TEA datasets include raw *k-space* data as well as classification labels, detection bounding boxes, segmentation masks, and quantitative T2 relaxation time maps [35]. The original work profiled how different reconstruction approaches combined with different segmentation tools affected a common musculoskeletal biomarker of cartilage T2 relaxation time. Despite differences in the performance of the individual DL blocks, the overall impact on regional cartilage T2 values was small, a surprising finding that has been replicated for cartilage morphology and T2 tasks [248, 249]. Recent work has evaluated new approaches that combine generic pre-training tasks such as image reconstruction with fine-tuning for different clinically-relevant downstream tasks [250]. This approach achieves high performance in image acceleration as well as segmentation. Similar to these findings, the K2S challenge at MICCAI 2022 combined knee MRI reconstruction with bone/cartilage segmentation and bone shape analysis [40]. Yet again, there were only weak correlations between the metrics of reconstruction and segmentation quality, with one of the best segmentation models producing highly artifactual reconstructions but high quality segmentations.

Joint estimation of sensitivity maps and reconstruction

The power of DL has also been harnessed for improving parallel multi-coil MRI, where the coil sensitivity maps must be estimated and incorporated in the image reconstruction process. Many DL methods utilize the popular ESPIRiT algorithm [251] for computing the sensitivity maps prior to the reconstruction process. However, joint estimation of the sensitivity maps and reconstructed data could contribute to improving image quality, as indicated in different studies, first with classical approaches [252, 253] and later using DL [254–257]. DL frameworks were hence recently developed for joint estimation of the sensitivity maps and reconstruction data. For example, the well-known E2E-VarNet method [254] included a module for sensitivity maps estimation and incorporated it into a larger unrolled network,

trained end-to-end. A similar approach was taken by Jun et al., who proposed the IC-Net [255]. Luo et al. suggested using a deep image prior [256], and Zhang et al. proposed a zero-shot learning method which is trained solely on data from a specific subject and jointly estimates the sensitivity maps and temporal data [258]. Most recently, Hu et al. [257] introduced the self-supervised SPICER method, which enables joint reconstruction and sensitivity maps estimation with training only on noisy data.

Other applications

The powerful capabilities of DL have also been exploited for other computational tasks in the MRI workflow aside from image reconstruction. A detailed review of these applications is beyond the scope of this manuscript, which focuses on MRI reconstruction. Specific examples include the joint recovery of multi-contrast MRI data [110, 259], the synthesis of missing contrasts or synthesis of quantitative maps based on anatomical data [259, 260], super-resolution [261, 262], B0 estimation and off-resonance correction [263], enhancement of low-field MRI data, where the low SNR degrades image quality [56, 264], and automated scan prescription [265].

Datasets and software

Datasets

The availability of public datasets and open-source code repositories has played a crucial role in the rapid development of DL techniques [266]. In the MRI reconstruction field, several major databases such as fastMRI [32], SKM-TEA [35], mridata.org [33], and Calgary-Campinas [34] catalyzed development by making available large amounts of raw *k-space* data, which are useful for developing and benchmarking methods [51, 267]. Other resources provide valuable data for specific applications, including MR imaging of speech production [268], cardiovascular imaging [269], and low-field MRI [270]. Many other MRI datasets are also available on the web, but those were commonly designed for downstream, non-reconstruction tasks, hence they do not always contain raw *k-space* data. Examples include the Human Connectome project [271], IXI [272], BRaTS [273, 274], ADNI [275], UK-biobank [276] and OASIS [277]

Open-source software

The adoption of open-source frameworks has significantly accelerated the development of DL methods as they provide researchers robust, flexible platforms to develop new algorithms. The two most prominent general-purpose DL

frameworks are PyTorch [278] and TensorFlow [279], which offer extensive libraries that facilitate the design, training, and deployment of DL models.

Several open-source software frameworks have been developed specifically for MRI. These offer useful computational tools for handling raw *k-space* data, implementation of algorithms and computation of MRI-related metrics. For example, BART (Berkeley Advanced Reconstruction Toolbox) [280] is a large and highly popular software package. It enables efficient data processing and contains implementations of different iterative reconstruction algorithms. The recent versions of BART also contain general-purpose tools that are highly useful for the development of DL reconstruction models, e.g., an automated differentiation framework compatible with complex-valued data, and implementations of well-established DL models [281]. Gadgetron [282] is another popular package, which offers extensive tools for image reconstruction, data management, and implementations of iterative solvers. A different package is MRIReco.jl [283], which is written entirely in Julia and utilizes the ISMRMRD file format. This package offers many building blocks for data management, simulations, and image reconstruction. Another example is Sigpy [284], which offers a set of operators, blocks, and algorithms that are highly suitable for iterative reconstruction. Unlike other toolboxes, Sigpy is written entirely in Python and can hence be integrated easily into frameworks such as PyTorch.

While the above packages focused on data management and image reconstruction, SNOPI (non-Cartesian sampling trajectory) [178] is a framework that offers practical tools for optimizing *k-space* sampling trajectories. These tools include a differentiable MRI system model and loss functions corresponding to constraints on image quality, e.g. hardware (e.g., maximum slew rate and gradient strength) and peripheral nerve stimulation (PNS) constraints. Finally, a different framework is Yarra (<https://cai2r.net/resources/yarra/>), which provides tools for automated collection of raw *k-space* data. These can facilitate acquisition of datasets and the creation of new datasets.

Another key area is pulse sequence development. One of the main challenges in reproducing complex pulse sequences across different sites and scanners is the dedicated prototyping environment and software used by each vendor. Pulseseq [285] is a rapid, hardware-independent pulse sequence prototyping framework, which enables intuitive high-level programming of acquisition protocols in Matlab or Python. It enables easy deployment across different field strengths and hardware. Importantly, the spin physics associated with the specific acquisition protocol compiled at the scanner can be accurately simulated as part of the Pulseseq framework or its derivatives [286]. Techniques for data harmonization can help mitigate challenges in transferring protocols across different systems [287].

In the context of quantitative imaging, the qMRLab software was developed to facilitate reproducibility across MRI systems [288]. It consists of practical tools for analyzing and processing quantitative MRI data acquired by different vendors. The user-friendly interface and modular design of qMRLab enable researchers to easily implement and share quantitative MRI techniques.

Many other open-source codes can also be found on online platforms such as GitHub, articles with code (<https://paperswithcode.com/>), and the two dedicated websites of the ISMRM: MR-Hub (<https://ismrm.github.io/mrhub/>), which hosts toolboxes, and MR-Pub (<https://ismrm.github.io/mrpub/>), which hosts git repositories published together with articles.

The toolboxes and platforms described above are essential for enhancing reproducible research in the MRI community. The development of a unified data format, the ISMRMRD [289] can also facilitate easy translation of datasets and methods across sites and research groups. The recent development of techniques for federated learning [290, 291] are also useful for training algorithms collaboratively without sharing the data; this can help address data-privacy issues.

Robustness challenges

In this section we discuss challenges related to developing, evaluating, and benchmarking DL reconstruction methods, and suggest approaches for mitigating them.

Distribution shifts

In deep learning, *generalization* refers to the ability of a trained model to accurately reconstruct images that it has never seen before, particularly when these new images differ substantially from the training data. Good generalization of DL-based MRI reconstruction models is critical for clinical workflows. However, achieving a good generalization is challenging because MRI data can vary substantially in terms of different factors, e.g. MRI hardware, vendor-specific scanning protocols, patient populations, and the anatomical regions being imaged. This variability can lead to a model that performs well on data from one source but poorly on data from another, a phenomenon known as *domain shift* or *distribution shift*. Here we review some of the main challenges related to this issue.

Domain shifts have been studied from several perspectives. Johnson et al. [292] analyzed the robustness of the models submitted to the 2019 fastMRI challenge to distribution shifts, e.g. small structural changes, addition of noise to *k-space* data, and changes in the number of coils. The study found that many of these models were sensitive

to the distribution shifts. Darestani et al. [293] evaluated the robustness of DL reconstruction methods with regard to out-of-distribution data, and found that both trained and untrained networks were affected by distribution shifts. Avidan et al. studied another type of distribution shift [294], related to sampling; methods trained on specific sampling schemes may not generalize well to other schemes. Altogether, distribution shifts can lead to substantial performance drops in MRI and can hence be a major limiting factor in practice.

Potential mitigation strategies. In cases where only a few training examples from a target domain are available, pre-training a network on other data and fine-tuning it to the target domain can improve performance [295–297]. In the challenging case where no target data are available for fine-tuning, test-time-training, which involves adapting to a single training at inference, is a viable performance-enhancing alternative [298]. Another good strategy that can help mitigate the performance drop due to distribution shifts is to train on broad and diverse data [299].

A growing body of work has pointed out the advantages of diffusion models in robustifying networks to distribution shifts. As described above (“Diffusion models”), diffusion models decouple the image prior from the statistical measurement model. They can hence generalize easily to various anatomies and sampling patterns [94–96, 98, 105, 106]. Another technique to robustify networks to shifts in sampling patterns is to provide the network with the under-sampling mask, and train the network to generalize to various sampling masks [294]. In addition, generative networks can be trained solely on magnitude images and applied to complex-valued data with different sub-sampling patterns and out-of-distribution data [98].

Bias and “data crimes”

In the field of AI, the term *bias* is often associated with gender-related or population-related bias. This can occur when models are trained on datasets that do not contain equal distributions of subjects having different genders, different ethnicities, or even data that only contain a narrow set of medical conditions [300]. This training can lead to algorithmic failure for under-served populations [301, 302] or rare conditions [300].

However, when solving inverse problems, bias can also arise from a naive, seemingly-appropriate use of open-access datasets. One of the primary challenges in the development of DL reconstruction methods is the need for raw *k-space* data, which are scarce and difficult to acquire due to the high cost of MRI scans and the long scan duration. While several databases offer such data, e.g. [32–35, 250, 268], there are many other databases that offer non-raw MRI data.

Those are generally designed for downstream tasks, e.g., segmentation and classification, hence they are frequently preprocessed. Nevertheless, due to their high availability, researchers sometimes download and use them for synthesizing *k-space* data, for training DL reconstruction models. This has been referred to as “off-label” data use, because those datasets are used for a different task than the one they were designed for [303].

Surprisingly, training DL reconstruction algorithms using “off-label” data could give rise to good-looking results; nevertheless, those are often biased and *overly optimistic*, i.e. they are “too good to be true” [303]. This is due to subtle preprocessing steps which, although imperceptible to the naked eye, impact algorithm performance. Common preprocessing steps include *k-space* zero-padding, coil combination, and JPEG data compression. The authors of [303] demonstrated that CS, dictionary learning, and DL algorithms are all sensitive to these preprocessing steps and yield biased results. This underscores their potential risk, as these algorithms are aimed for clinical purposes. The findings also demonstrated that DL algorithms trained on preprocessed data may not be able to generalize well to real-world clinical data, and could potentially eliminate crucial clinical details [303].

Another concerning finding is that popular error metrics, e.g., the normalized root mean square error (NRMSE) and SSIM could be *blind* to the preprocessing [303], and miss the bias. This is because those metrics compare the reconstructed and reference image, but those come from the same underlying data, i.e. both are preprocessed. The error metrics thus cannot measure the true image quality. Therefore, DL algorithms may achieve “good” NRMSE and SSIM scores even when their performance is poor. This makes the head-to-head comparison of results across papers very difficult, because some papers report experiments with raw *k-space* data while others report results for preprocessed data, where the metrics tend to be better. To raise awareness of this phenomenon, the authors dubbed the publication of misleading results “data crimes”.

Potential mitigation strategies. In the context of gender-related or population-related biases, the best strategy is to train on large, diverse datasets [300–302]. In the context of bias that stems from training on preprocessed data, the optimal strategy is to train solely on raw *k-space* data. When raw, zero-padded *k-space* data are available, and such data have not been subject to any other preprocessing steps, one simple correction step is to crop *k-space* to its original size. However, this scenario is quite rare. In practice, other preprocessing steps are often applied, e.g. coil-combination (e.g. using root-sum-of-squares operation) and JPEG compression. Those steps are irreversible, hence there is no simple technique to remedy them, and the bias cannot be easily prevented.

Techniques for data synthesis. When raw k -space data are unavailable, one possible strategy is to train on synthetic data [304–306]. When magnitude data are available, synthetic complex-valued data can be obtained, for example, by adding a synthetic phase to the magnitude images and training a generative model to learn priors of complex-valued images [306]. Multi-coil data can be synthesized by multiplying the phase-enhanced magnitude data with sensitivity maps [304]. Another approach is to leverage the Bloch equations to simulate realistic data [305, 307]. However, it should be emphasized that the synthesis of complex-valued data does not guarantee good performance for real-world data; i.e., it cannot automatically prevent the bias described above. Models trained on synthetic data must therefore be tested using real-world, raw k -space data.

Hallucinations

The term *hallucinations* refers to the generation of false, realistic-looking features which are not present in the actual data. This can arise from the use of inaccurate priors [308], e.g., when there is a distribution shift between the training and test data, as described above. Strikingly, the team that organized the second FastMRI challenge found that many of the top-performing models produced hallucinations, and that these hallucinations were not captured by image quality metrics such as SSIM [37]. They also noticed that hallucinations could morph abnormal structures into seemingly normal ones.

Several studies have also highlighted and explored manifestations of hallucinations. For example, Cohen et al. (2018) [309] discussed how distribution matching losses in medical image translation can lead to hallucinated features. Bhadra et al. [308] reported hallucinations in the context of tomographic image reconstruction and introduced the concept of hallucination maps to identify and understand the impact of prior information in regularized reconstruction methods. Gottschling et al. [310] explored hallucinations from a theoretical perspective, and highlighted problematic scenarios of *in-distribution* hallucinations. The issue of hallucinations is a critical concern in DL-based image reconstruction, as it can lead to misleading results and false diagnoses.

Potential mitigation strategies. At present, there is a pressing need for techniques to mitigate hallucinations. Currently, the best strategy is to have radiology experts evaluate the reconstructed images, in the hope that they will be able to detect hallucinations. However, more research and development are required.

Adversarial attacks and instabilities

Neural networks for image classification are known to be sensitive to adversarial attacks, i.e. small, imperceptible, adversarially chosen perturbations. Generally, such attacks are not a major concern for clinical MRI systems, as those systems are typically closed and require password access. However, adversarial attacks have been utilised for analysing robustness of DL reconstruction methods and this topic attracted substantial attention [309–312]. For example, Antun et al. [312] provided simulations demonstrating that small, adversarially selected perturbations in undersampled measurements can result in severe reconstruction artifacts. They also showed that there is a trade-off between robustness and performance, and that classical sparsity-based reconstruction methods are also sensitive to adversarially selected perturbations.

Robustness to adversarial attacks is still an open issue. For example, at present there is no conclusive evidence indicating that DL reconstruction methods are more sensitive than classical methods such as sparsity-based reconstruction to worst-case perturbations. Krainovic et al. [313] provided theoretical results on a worst-case optimal estimator. Morshuis et al. [314] found that simple end-to-end variational networks are as sensitive to perturbations as U-nets, and Darestani et al. [293] showed that both un-trained and trained networks are sensitive to such perturbations.

Strategies for enhancing robustness. Goujon et al. [315] demonstrated that the robustness of reconstruction algorithms can be improved by constraining the CNN module to be convex or using a monotone constraint [72]. However, a global convexity or monotone constraint often translates into reduced performance [72, 315], as predicted by Antun et al. [312]. A recent work showed that the global monotone constraint described in [72] can be replaced by a local constraint around the image manifold [316], to achieve improved robustness without compromising on performance. Following the approach implemented in non-convex algorithms, this algorithm is theoretically guaranteed to converge to a minimum, provided it is initialized with SENSE reconstructions. From a different perspective, [317] and [318] suggested the use of adversarial attacks during training to reduce false negatives.

Benchmarking challenges

When evaluating various techniques comparing the performance of different reconstruction and image analysis methods, it is crucial to have appropriate benchmarks that can characterize true clinical utility. Conventionally, image quality metrics such as the mean square error (MSE), peak signal-to-noise ratio (PSNR), and the structural similarity

metric (SSIM) [319] are used to assess the image quality of reconstructed MR images. These metrics have been described extensively in the computer vision literature and have a moderate correspondence with human-perceived image quality [319]. However, as there is a substantial domain shift between natural images and medical images, studies have shown that such traditional image quality metrics do not correlate well with radiologist-perceived metrics [320]. One likely reason for this phenomenon is that not all pixels in a given image have similar diagnostic value. Consider, for example, knee MRI scans, e.g. those contained in the popular fastMRI database. Most abnormalities in knee scans are likely to be located in small, subtle regions in the cartilage, meniscus, and ligaments, while the tissues that are visible in the scans, such as bone and muscle, are likely to have substantially fewer abnormalities. Computing traditional image quality metrics that weigh all image pixels similarly may therefore not reflect how a radiologist perceives the images. Thus, it is essential to define metrics that correspond to downstream clinical utility so that the same metrics can be used for benchmarking models.

Another issue of concern is that the performance of DL reconstruction algorithms is typically evaluated using a relatively narrow test set, which is similar in nature to the training data. The evaluation results can thus be misleading, since they do not yield a reliable estimation of the model's generalization ability; i.e., the ability to perform well on test data that deviate from the training data, which is of critical importance in clinical settings.

Yet another challenge is related to the common train/test data split. Even for a single dataset, different research groups or studies may apply different splits, and hence train and test the model with different data subsets. For example, even if two identical studies used the same number of training images, it is not always possible to verify whether they used the same images. Previous work has shown that there can be a significant variation in how well the same deep learning models perform on different imaging exams within the same larger dataset [293]. Consequently, even if the number of training examples is maintained across studies, some examples may be easier or more challenging to train and evaluate. This makes a head-to-head comparison of published literature difficult to interpret.

Strategies for mitigating benchmarking challenges. Initiatives such as the large-scale fastMRI database and challenge [32] allow multiple different methods to be compared on identical training, validation, and testing data splits. This type of consistent evaluation platform allows for head-to-head comparisons across different methods and can shed light on their pros and cons.

More recent studies dealing with image quality assessment have gone beyond PSNR and SSIM metrics and used perceptual metrics of image quality instead of relying on

handcrafted metrics [321]. Perceptual metrics utilize representations extracted from pre-trained neural networks and have been used for evaluating and optimizing MR reconstruction quality [322, 323]. For example, Learned Perceptual Image Patch Similarity (LPIPS) is commonly used for assessing the quality of natural images. A similar analog of Self-Supervised Feature Distance (SSFD) has been proposed for assessing MR image quality [324, 325]. When using these perceptual metrics, a reference image and an evaluation image can be fed into the same pre-trained network, and the differences in their representations is compared and taken as the similarity between the two images. The impact of common image perturbations on such deep feature metrics is shown in Fig. 7. These perceptual metrics can be computed using networks that were pre-trained either on natural images or medical images. A recent study demonstrated that although there may not be considerable differences between these two types of networks, i.e., between the networks used for generating the representations, using perceptual metrics is better than using traditional image quality metrics when evaluating against radiologists' perceived image quality [322]. However, due to the challenges in developing image quality metrics that reflect radiological assessment, further research is needed in this context.

Beyond the metrics of image quality, another approach for benchmarking reconstruction models aims to directly benchmark the downstream value provided by the MR image. For example, the SKM-TEA dataset is designed to enable the evaluation of downstream tasks such as clinical classification, segmentation, and articular cartilage T2 quantification [35]. This dataset emphasizes the combination of automated cartilage segmentation with T2 quantification, a known metric with clinical and research significance [326, 327]. Similarly, the K2S dataset aims to optimize cartilage volume and thickness quantification, whereas the fastMRI+ dataset assesses classification and detection tasks [40, 246]. Such benchmarking metrics, which relate to clinical significance, provide a promising path forward for assessing image quality when used in conjunction with traditional or perceptual image quality metrics. Doing so can contribute to optimizing both image quality and image value for radiologists.

Uncertainty estimation

Connecting both upstream and downstream DL tools with clinical practice first requires promoting confidence among the users of these tools. Robustly characterizing the uncertainty of DL models may help increase the trust of downstream users in the outputs of DL algorithms prior to their integration into routine clinical practice. Uncertainty quantification, a popular subfield of DL research, can encourage building such trust among downstream users, regarding




Fig. 7 Impact of common image perturbations on image quality metrics. A variety of image perturbations applied to a sample image from the fastMRI dataset (top row: noise addition, image blurring, pixel *rolling* (where an image is shifted by a number of pixels), and physics-based subject motion). The impact of these corruptions is shown for conventional image quality metrics (SSIM, PSNR) and

deep feature distance metrics (LPIPS—made for natural images, SSFD—made for MR images). The deep feature metrics exhibit a larger dynamic range to the noise, blurring, and motion corruptions, but present very little change due to pixel rolling, since the image quality does not change. These qualities of deep feature metrics are ideal for assessing MRI reconstruction quality

both image reconstruction and automated image analysis algorithms.

Several works have proposed methods to evaluate the uncertainty of accelerated MRI reconstructions [97, 328–332]. Given that MRI reconstruction is an ill-posed problem where several high-quality images can correspond to the same low-quality image, quantifying the uncertainty of a sample reconstruction can help guide the fidelity of the overall reconstruction process. The developed uncertainty approaches leverage the variational formulation of inverse image recovery for sampling not only a single output of reconstruction, but a variety of different outputs using a learned model of the posterior distribution. Sampling a multitude of image outputs makes it possible to evaluate the variance at the voxel-level to compute uncertainties within

a given image. This technique of uncertainty quantification can be used to determine whether specific regions of interest have high uncertainty values around abnormal image findings.

For downstream image analysis tasks, techniques like Monte Carlo Dropout have gained popularity for estimating output uncertainty [333]. Although dropout is typically employed for model regularization and reducing overfitting during training by randomly setting parameters to zero, it can also be applied at inference time. This involves producing multiple outputs with different neural network weights set to zero. This Monte Carlo Dropout approximates a Gaussian process, which facilitates the computation of uncertainty across the variance of all the generated inputs. Although

straightforward to implement, this approach requires multiple forward passes during inference.

Looking forward, despite our ability to generate voxel-level uncertainties, the optimal utilization of these generated uncertainty maps still remains unclear. This elicits a number of intriguing research questions concerning the ways in which these uncertainties can be utilized beyond simply presenting them to end-users. For instance, it is unlikely that a radiologist would review both the output of a deep learning reconstruction and the corresponding uncertainty map since doing so would nearly double clinical read durations, and the radiologist may not necessarily know how to contextualize regions of high and low uncertainty.

Alternatively, these uncertainty maps could be applied in iterative reconstruction techniques [334]. In this approach, the uncertainty associated with a specific step in an image reconstruction pipeline could serve as input for the subsequent step. The reconstruction network's objective would be to reconstruct the same image but with the new constraint of reducing the underlying uncertainty. This type of formulation holds potential for adaptive, case-based sampling schemes tailored to individual patients.

Another promising area of research involves directly integrating uncertainties from MR image reconstruction with those of the relevant downstream parameters of interest [335]. This methodology could serve to estimate maximum acceleration rates without compromising quantitative clinical parameters. Achieving such end-to-end analysis necessitates datasets that can provide raw case-based data alongside downstream image analysis datasets (e.g. SKM-TEA, fastMRI+, K2S, etc).

Implementation issues

To practically reproduce, implement, and further develop the DL-based strategies discussed in the review, one must have access to the appropriate hardware resources. This section will highlight the main computational challenges expected, especially for large networks. For purposes of illustration, we focus on the case of unrolled networks.

Memory demands of unrolled algorithms. PnP and score-based algorithms that pre-train deep learning modules as denoisers are associated with low memory demand. In contrast, unrolled algorithms, which offer state-of-the-art performance compared to PnP methods, involve a number of iterations and their training is restricted by the memory of the GPU devices during training. This often limits the applicability of unrolled algorithms to large-scale multi-dimensional problems.

Strategies for computational efficiency. Several strategies have been introduced to overcome the memory limitations of unrolled methods. For an unrolled network with N

iterations and shared CNN modules across iterations, the computational complexity and memory demands of backpropagation are $\mathcal{O}(N)$ and $\mathcal{O}(N)$, respectively. The forward steps can be recomputed during backpropagation, which reduces the memory demand to $\mathcal{O}(1)$, while the computational complexity increases to $\mathcal{O}(N^2)$. Forward checkpointing [336] saves the variables for every K layers during forward propagation, which reduces the computational demand to $\mathcal{O}(NK)$, while the memory demand is $\mathcal{O}(N/K)$. Reverse recalculation has been proposed to reduce the memory demand to $\mathcal{O}(1)$ and computational complexity to $\mathcal{O}(N)$ [337]. However, the approach in [337] requires multiple iterations to invert each CNN block, resulting in high computational complexity in practical applications. The deep equilibrium (DEQ) model [338] was recently adapted to inverse problems to significantly improve the memory demand [71, 72] of unrolled methods. Unlike unrolled methods, DEQ schemes run the iterations until convergence, similar to PnP algorithms. This property makes it possible to perform forward and backward propagation using fixed-point iteration involving a single physical layer, which reduces the memory demand to $\mathcal{O}(1)$, while the computational complexity is $\mathcal{O}(N)$; this offers better tradeoffs than the alternatives discussed above [336, 337]. The runtime of DEQ methods that are iterated until convergence are variable compared to unrolled methods, which use a finite number of iterations. In addition, the convergence of the iterative algorithm is crucial for the accuracy of backpropagation steps in DEQ, unlike in unrolled methods. Convergence guarantees were introduced in [71, 72].

In summary, while computational infrastructure is increasingly evolving and expanding, both within research institutes and as cloud services, deep-learning-based reconstruction algorithms may require tremendous storage, RAM, and parallelization capabilities, to facilitate the training and inference of massive and complex models. Computational efficiency should, therefore, be considered an important aspect, when assessing the applicability of a new ML approach.

Conclusion

The introduction and rapid development of deep-learning-based strategies for MRI reconstruction have brought about a dramatic acceleration in acquisition time, paved the way for rapid and accurate parameter mapping, and facilitated automatic schedule optimization. While several key challenges still lie ahead, especially in terms of robust generalization, careful considerations of the training data diversity, ongoing model validation, and potentially, the development

of adaptive or continuous learning systems are expected to enable adjustment to new data distributions over time.

Acknowledgements The authors thank Jonathan Tamir for useful comments regarding diffusion models. M.J acknowledges support from NIH grants R01 AG067078, R01EB019961, R01 EB031169, and Canon Medical Research. O.P. acknowledges funding support from the Ministry of Innovation, Science and Technology, Israel, the Tel Aviv University Center for AI and Data Science (TAD), and the European Union under the Horizon Europe program (ERC, BabyMagnet, project no. 101115639). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. E.S. is a Horev Fellow and acknowledges funding support from the Technion's Leaders in Science and Technology program. A.C. receives support from NIH grants R01 AR077604, R01 EB002524, and R01 AR079431; and from GE Healthcare, Philips, Amazon, Microsoft, and Stability AI.

Funding Open access funding provided by Technion - Israel Institute of Technology.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Constantine G, Shan K, Flamm SD, Sivananthan MU (2004) Role of MRI in clinical cardiology. *Lancet* 363(9427):2162–2171
- Hollingsworth KG (2015) Reducing acquisition time in clinical MRI by data undersampling and compressed sensing reconstruction. *Phys Med Biol* 60(21):297
- Sodickson DK, Manning WJ (1997) Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magn Reson Med* 38:591–603
- Griswold MA et al (2002) Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med* 47(6):1202–1210
- Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P (1999) SENSE: sensitivity encoding for fast MRI. *Magn Reson Med* 42(5):952–962
- Seiberlich N, Breuer FA, Blaimer M, Barkauskas K, Jakob PM, Griswold MA (2007) Non-Cartesian data reconstruction using GRAPPA operator gridding (GROG). *Magn Reson Med* 58:1257–1265
- Deshmane A, Gulani V, Griswold MA, Seiberlich N (2012) Parallel MR imaging. *J Magn Reson Imaging* 36(1):55–72
- Lustig M, Pauly JM (2010) SPIRiT: iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magn Reson Med* 64(2):457–471
- Zaitsev M, Maclaren J, Herbst M (2015) Motion artifacts in MRI: a complex problem with many partial solutions. *J Magn Reson Imaging* 42:887–901
- Lustig M, Donoho D, Pauly JM (2007) Sparse MRI: the application of compressed sensing for rapid MR imaging. *Magn Reson Med* 58(6):1182–1195
- Donoho DL (2006) Compressed sensing. *IEEE Trans Inf Theory* 52(4):1289–1306
- Jung H, Sung K, Nayak KS, Kim EY, Ye JC (2009) k-t FOCUSS: a general compressed sensing framework for high resolution dynamic MRI. *Magn Reson Med* 61(1):103–116
- Feng L, Axel L, Chandarana H, Block KT, Sodickson DK, Otazo R (2016) XD-GRASP: golden-angle radial MRI with reconstruction of extra motion-state dimensions using compressed sensing. *Magn Reson Med* 75(2):775–788
- Duarte MF, Eldar YC (2011) Structured compressed sensing: from theory to applications. *IEEE Trans Signal Process* 59:4053–4085
- Ye JC (2019) Compressed sensing MRI: a review from signal processing perspective. *BMC Biomed Eng* 1(1):1–17
- Lingala SG, Jacob M (2013) Blind compressive sensing dynamic MRI. *IEEE Trans Medical Imaging* 32(6):1132–1145
- Fessler JA (2020) Optimization methods for magnetic resonance image reconstruction: key models and optimization algorithms. *IEEE Signal Process Mag* 37(1):33–40
- LeCun Y, Bengio Y, Hinton G (2015) Deep learning. *Nature* 521(7553):436–444
- Alzubaidi L, Zhang J, Humaidi AJ, Al-Dujaili A, Duan Y, Al-Shamma O, Santamaria J, Fadhel MA, Al-Amidie M, Farhan L (2021) Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. *J Big Data* 8:1–74
- Shrestha A, Mahmood A (2019) Review of deep learning algorithms and architectures. *IEEE Access* 7:53040–53065
- Chandra SS, Bran Lorenzana M, Liu X, Liu S, Bollmann S, Crozier S (2021) Deep learning in magnetic resonance image reconstruction. *J Med Imaging Radiat Oncol* 65(5):564–577
- Zhu B, Liu JZ, Cauley SF, Rosen BR, Rosen MS (2018) Image reconstruction by domain-transform manifold learning. *Nature* 555(7697):487–492
- Wang S, Cheng H, Ying L, Xiao T, Ke Z, Zheng H, Liang D (2020) DeepcomplexMRI: exploiting deep residual network for fast parallel MR imaging with complex convolution. *Magn Reson Imaging* 68:136–147
- Hammernik K et al (2018) Learning a variational network for reconstruction of accelerated MRI data. *Magn Reson Med* 79(6):3055–3071
- Aggarwal HK, Mani MP, Jacob M (2018) MoDL: model-based deep learning architecture for inverse problems. *IEEE Trans Med Imaging* 38(2):394–405
- Lundervold AS, Lundervold A (2019) An overview of deep learning in medical imaging focusing on MRI. *Zeitschrift für Medizinische Physik* 29(2):102–127
- Mazurowski MA, Buda M, Saha A, Bashir MR (2019) Deep learning in radiology: an overview of the concepts and a survey of the state of the art with focus on MRI. *J Magn Reson Imaging* 49(4):939–954
- Ravishanker S, Ye JC, Fessler JA (2019) Image reconstruction: from sparsity to data-adaptive methods and machine learning. *Proc IEEE* 108(1):86–109
- Knoll F, Hammernik K, Zhang C, Moeller S, Pock T, Sodickson DK, Akcakaya M (2020) Deep-learning methods for parallel magnetic resonance imaging reconstruction: a survey of the current approaches, trends, and issues. *IEEE Signal Process Mag* 37(1):128–140
- Liang D, Cheng J, Ke Z, Ying L (2020) Deep magnetic resonance image reconstruction: inverse problems meet neural networks. *IEEE Signal Process Mag* 37(1):141–151

31. Shimron E, Perlman O (2023) AI in MRI: computational frameworks for a faster, optimized, and automated imaging workflow. *Bioengineering (Basel)* 10:492
32. Knoll F et al (2020) fastMRI: a publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning. *Radiol Artif Intell* 2(1):e190007
33. Ong F, Amin S, Vasanaawala S, Lustig M (2018) Mridata.org: an open archive for sharing MRI raw data. In: *International society for magnetic resonance in medicine*, vol 26, p 1
34. Souza R, Lucena O, Garrafa J, Gobbi D, Saluzzi M, Appenzeller S, Rittner L, Frayne R, Lotufo R (2018) An open, multi-vendor, multi-field-strength brain MR dataset and analysis of publicly available skull stripping methods agreement. *NeuroImage* 170:482–494
35. Desai AD et al (2021) SKM-TEA: a dataset for accelerated MRI reconstruction with dense image labels for quantitative clinical evaluation. In: *Thirty-fifth conference on neural information processing systems datasets and benchmarks track*
36. Wang C, Lyu J, Wang S, Qin C, Guo K, Zhang X, Yu X, Li Y, Wang F, Jin J et al (2023) CMRxRecon: an open cardiac MRI dataset for the competition of accelerated image reconstruction. *arXiv preprint arXiv:2309.10836*
37. Muckley MJ, Riemenschneider B, Radmanesh A, Kim S, Jeong G, Ko J, Jun Y, Shin H, Hwang D, Mostapha M et al (2021) Results of the 2020 fastMRI challenge for machine learning MR image reconstruction. *IEEE Trans Med Imaging* 40(9):2306–2317
38. Knoll F, Murrell T, Sriram A, Yakubova N, Zbontar J, Rabbat M, Defazio A, Muckley MJ, Sodickson DK, Zitnick CL et al (2020) Advancing machine learning for MR image reconstruction with an open competition: overview of the 2019 fastMRI challenge. *Magn Reson Med* 84(6):3054–3070
39. Muckley MJ, Riemenschneider B, Radmanesh A, Kim S, Jeong G, Ko J, Jun Y, Shin H, Hwang D, Mostapha M et al (2020) State-of-the-art machine learning MRI reconstruction in 2020: Results of the second fastMRI challenge. *arXiv preprint arXiv:2012.06318* 2(6):7
40. Tolpadi AA, Bharadwaj U, Gao KT, Bhattacharjee R, Gassert FG, Luitjens J, Giesler P, Morshuis JN, Fischer P, Hein M et al (2023) K2S challenge: from undersampled k-space to automatic segmentation. *Bioengineering* 10(2):267
41. Lyu J, Qin C, Wang S, Wang F, Li Y, Wang Z, Guo K, Ouyang C, Tänzer M, Liu M et al (2024) The state-of-the-art in cardiac MRI reconstruction: results of the CMRxRecon challenge in MICCAI 2023. *arXiv preprint arXiv:2404.01082*
42. Arghya Pal, Yogesh Rathi (2022) A review and experimental evaluation of deep learning methods for MRI reconstruction. *J Mach Learn Biomed Imaging* 1:001
43. Lin DJ, Johnson PM, Knoll F, Lui YW (2021) Artificial intelligence for MR image reconstruction: an overview for clinicians. *J Magn Reson Imaging* 53(4):1015–1028
44. Oscanoa JA, Middione MJ, Alkan C, Yurt M, Loecher M, Vasanaawala SS, Ennis DB (2023) Deep learning-based reconstruction for cardiac MRI: a review. *Bioengineering* 10:334
45. Spieker V, Eichhorn H, Hammernik K, Rueckert D, Preibisch C, Karampinos DC, Schnabel JA (2023) Deep learning for retrospective motion correction in MRI: a comprehensive review. *IEEE Trans Med Imaging* 43:846–859
46. Scarlett J, Heckel R, Rodrigues MR, Hand P, Eldar YC (2022) Theoretical perspectives on deep learning methods in inverse problems. *IEEE J Select Areas Inf Theory* 3(3):433–453
47. Hammernik K, Kustner T, Rueckert D (2022) Machine learning for MRI reconstruction. *Advances in magnetic resonance technology and applications*, vol 7. Elsevier, Amsterdam, pp 281–323
48. Bustin A, Fuin N, Botnar RM, Prieto C (2020) From compressed-sensing to artificial intelligence-based cardiac MRI reconstruction. *Front Cardiovasc Med* 7:17
49. Hammernik K, Kustner T, Yaman B, Huang Z, Rueckert D, Knoll F, Akcakaya M (2023) Physics-driven deep learning for computational magnetic resonance imaging: combining physics and machine learning for improved medical imaging. *IEEE Signal Process Mag* 40:98–114
50. Bilgic B, Goyal VK, Adalsteinsson E (2011) Multi-contrast reconstruction with Bayesian compressed sensing. *Magn Reson Med* 66(6):1601–1615
51. Hammernik K, Schlemper J, Qin C, Duan J, Summers RM, Rueckert D (2021) Systematic evaluation of iterative deep neural networks for fast parallel MRI reconstruction with sensitivity-weighted coil combination. *Magn Reson Med* 86(4):1859–1872
52. Luo G, Zhao N, Jiang W, Hui ES, Cao P (2020) MRI reconstruction using deep Bayesian estimation. *Magn Reson Med* 84(4):2246–2261
53. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J et al (2011) Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found Trends Mach Learn* 3(1):1–122
54. Beck A, Teboulle M (2009) A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J Imaging Sci* 2(1):183–202
55. Muckley MJ, Riemenschneider B, Radmanesh A, Kim S, Jeong G, Ko J, Jun Y, Shin H, Hwang D, Mostapha M et al (2021) Results of the 2020 fastMRI challenge for machine learning MR image reconstruction. *IEEE Trans Med Imaging* 40:2306–2317
56. Koonjoo N, Zhu B, Bagnall GC, Bhutto D, Rosen M (2021) Boosting the signal-to-noise of low-field MRI with deep learning image reconstruction. *Sci Rep* 11(1):1–16
57. Wang S et al (2016) Accelerating magnetic resonance imaging via deep learning, pp 514–517
58. Jin KH, McCann MT, Froustey E, Unser M (2017) Deep convolutional neural network for inverse problems in imaging. *IEEE Trans Image Process* 26(9):4509–4522
59. Lin K, Heckel R (2022) Vision transformers enable fast and robust accelerated MRI. *PMLR*. <https://proceedings.mlr.press/v172/lin22a.html>
60. Lin K, Heckel R (2022) Vision transformers enable fast and robust accelerated MRI. In: *International conference on medical imaging with deep learning*. PMLR, pp 774–795
61. Guo P, Mei Y, Zhou J, Jiang S, Patel VM (2023) ReconFormer: accelerated MRI reconstruction using recurrent transformer. *IEEE Trans Med Imaging* 43:582–593
62. Gregor K, LeCun Y (2010) Learning fast approximations of sparse coding. In: *Proceedings of the 27th international conference on neural information processing systems*. MIT Press, USA, pp 399–406
63. Sun J, Li H, Xu Z et al (2016) Deep ADMM-Net for compressed sensing MRI. In: *Advances in neural information processing systems* 29
64. Schlemper J, Caballero J, Hajnal JV, Price AN, Rueckert D (2018) A deep cascade of convolutional neural networks for dynamic MR image reconstruction. *IEEE Trans Med Imaging* 37:491–503
65. Sriram A, Zbontar J, Murrell T, Defazio A, Zitnick CL, Yakubova N, Knoll F, Johnson P (2020) End-to-end variational networks for accelerated MRI reconstruction. *arXiv:2004.06688*
66. Fabian Z, Tinaz B, Soltanolkotabi M (2022) HUMUS-Net: hybrid unrolled multi-scale network architecture for accelerated MRI reconstruction. *Adv Neural Inf Process Syst* 35:25306–25319
67. Darestani MZ, Nath V, Li W, He Y, Roth HR, Xu Z, Xu D, Heckel R, Zhao C (2024) IR-FRestormer: iterative refinement

- with fourier-based Restormer for accelerated MRI reconstruction. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision, pp 7655–7664
68. Pramanik A, Aggarwal HK, Jacob M (2020) Deep generalization of structured low-rank algorithms (Deep-SLR). *IEEE Trans Med Imaging* 39(12):4186–4197
 69. Eo T, Jun Y, Kim T, Jang J, Lee H-J, Hwang D (2018) KIKI-net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images. *Magn Reson Med* 80(5):2188–2201
 70. Wang B, Lian Y, Xiong X, Zhou H, Liu Z, Zhou X (2024) DCT-net: dual-domain cross-fusion transformer network for MRI reconstruction. *Magn Reson Imaging* 107:69–79
 71. Gilton D, Ongie G, Willett R (2021) Deep equilibrium architectures for inverse problems in imaging. *IEEE Trans Comput Imaging* 7:1123–1133
 72. Pramanik A, Zimmerman MB, Jacob M (2023) Memory-efficient model-based deep learning with convergence and robustness guarantees. *IEEE Trans Comput Imaging* 9:260–275
 73. Pramanik A, Jacob M (2023) Accelerated parallel MRI using memory efficient and robust monotone operator learning (MOL). In: 2023 IEEE 20th international symposium on biomedical imaging (ISBI), IEEE. pp 1–4
 74. Ahmad R, Bouman CA, Buzzard GT, Chan S, Liu S, Reehorst ET, Schniter P (2020) Plug-and-play methods for magnetic resonance imaging: using denoisers for image recovery. *IEEE Signal Process Mag* 37(1):105–116
 75. Kamilov US, Bouman CA, Buzzard GT, Wohlberg B (2023) Plug-and-play methods for integrating physical and learned models in computational imaging: theory, algorithms, and applications. *IEEE Signal Process Mag* 40(1):85–97
 76. Dabov K, Foi A, Katkovnik V, Egiazarian K (2007) Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Trans Image Process* 16(8):2080–2095
 77. Ryu E, Liu J, Wang S, Chen X, Wang Z, Yin W (2019) Plug-and-play methods provably converge with properly trained denoisers. In: International conference on machine learning, PMLR. pp 5546–5557
 78. Hong T, Xu X, Hu J, Fessler JA (2024) Provable preconditioned plug-and-play approach for compressed sensing MRI reconstruction. *arXiv preprint arXiv:2405.03854*
 79. Romano Y, Elad M, Milanfar P (2017) The little engine that could: regularization by denoising (RED). *SIAM J Imaging Sci* 10(4):1804–1844
 80. Kingma DP, Welling M (2013) Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*
 81. Tezcan KC, Baumgartner CF, Luechinger R, Pruessmann KP, Konukoglu E (2018) MR image reconstruction using deep density priors. *IEEE Trans Med Imaging* 38(7):1633–1642
 82. Goodfellow IJ, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Generative adversarial networks. *arXiv preprint arXiv:1406.2661*
 83. Yang G, Yu S, Dong H, Slabaugh G, Dragotti PL, Ye X, Liu F, Arridge S, Keegan J, Guo Y et al (2017) DAGAN: deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction. *IEEE Trans Med Imaging* 37(6):1310–1321
 84. Mardani M et al (2018) Deep generative adversarial neural networks for compressive sensing MRI. *IEEE Trans Med Imaging* 38(1):167–179
 85. Quan TM, Nguyen-Duc T, Jeong W-K (2018) Compressed sensing MRI reconstruction using a generative adversarial network with a cyclic loss. *IEEE Trans Med Imaging* 37(6):1488–1497
 86. Cole EK, Pauly JM, Vasanawala SS, Ong F (2020) Unsupervised MRI reconstruction with generative adversarial networks. *arXiv preprint arXiv:2008.13065*
 87. Lv J, Zhu J, Yang G (2021) Which GAN? A comparative study of generative adversarial network-based fast MRI reconstruction. *Philos Trans R Soc A* 379(2200):20200203
 88. Korkmaz Y, Dar SU, Yurt M, Özbey M, Cukur T (2022) Unsupervised MRI reconstruction via zero-shot learned adversarial transformers. *IEEE Trans Med Imaging* 41:1747–1763
 89. Yang L, Zhang Z, Song Y, Hong S, Xu R, Zhao Y, Zhang W, Cui B, Yang M-H (2023) Diffusion models: a comprehensive survey of methods and applications. *ACM Comput Surv* 56(4):1–39
 90. Croitoru F-A, Hondru V, Ionescu RT, Shah M (2023) Diffusion models in vision: a survey. *IEEE Trans Pattern Anal Mach Intell* 45:10850–10869
 91. Kawar B, Elad M, Ermon S, Song J (2022) Denoising diffusion restoration models. In: Neural information processing systems (NeurIPS)
 92. Bora A, Jalal A, Price E, Dimakis AG (2017) Compressed sensing using generative models. In: International conference on machine learning
 93. Ho J, Jain A, Abbeel P (2020) Denoising diffusion probabilistic models. *Adv Neural Inf Process Syst* 33:6840–6851
 94. Jalal A, Arvinte M, Daras G, Price E, Dimakis AG, Tamir JI (2021) Robust compressed sensing MRI with deep generative priors. In: style="font-family: -apple-system, BlinkMacSystemFont, "Segoe UI", Roboto, Oxygen, Ubuntu, Cantarell, "Fira Sans", "Droid Sans", "Helvetica Neue", sans-serif;">Advances in Neural Information Processing Systems 34: 14938–14954.
 95. Güngör A, Dar SU, Öztürk Ş, Korkmaz Y, Bedel HA, Elmas G, Ozbey M, Çukur T (2023) Adaptive diffusion priors for accelerated MRI reconstruction. *Med Image Anal* 88:102872
 96. Chung H, Ye JC (2022) Score-based diffusion models for accelerated MRI. *Med Image Anal* 80:102479
 97. Luo G, Blumenthal M, Heide M, Uecker M (2023) Bayesian MRI reconstruction with joint uncertainty estimation using diffusion models. *Magn Reson Med* 90(1):295–311
 98. Zach M, Knoll F, Pock T (2023) Stable deep MRI reconstruction using generative priors. *IEEE Trans Med Imaging* 42:3817–3832
 99. Mardani M et al (2018) Deep generative adversarial neural networks for compressive sensing MRI. *IEEE Trans Med Imaging* 38(1):167–179
 100. Ali H, Biswas MR, Mohsen F, Shah U, Alamgir A, Mousa O, Shah Z (2022) The role of generative adversarial networks in brain MRI: a scoping review. *Insights Imaging* 13(1):98
 101. Bouter ML, Ippolito G, O'Reilly TPA, Remis RF, Gijzen MB, Webb AG (2022) Deep learning-based single image super-resolution for low-field MR brain images. *Sci Rep* 12:6362
 102. Song Y, Ermon S (2020) Improved techniques for training score-based generative models. In: Advances in neural information processing systems
 103. Song Y, Sohl-Dickstein J, Kingma DP, Kumar A, Ermon S, Poole B (2020) Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*
 104. Chung H, Kim J, Mccann MT, Klasky ML, Ye JC (2022) Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*
 105. Daras G, Chu W-S, Kumar A, Lagun D, Dimakis AG (2021) Solving inverse problems with NerfGANs. *arXiv:2112.09061*
 106. Yu C, Guan Y, Ke Z, Lei K, Liang D, Liu Q (2023) Universal generative modeling in dual domains for dynamic MRI. *NMR Biomed* 36(12):5011
 107. Jalal A, Arvinte M, Daras G, Price E, Dimakis AG, Tamir J (2021) Robust compressed sensing MRI with deep generative priors. *Adv Neural Inf Process Syst* 34:14938–14954
 108. Levac B, Jalal A, Tamir JI (2023) Accelerated motion correction for MRI using score-based generative models. In: IEEE international symposium on biomedical imaging (ISBI), pp 1–5

109. Alkan C, Oscanoa J, Abraham D, Gao M, Nurdinova A, Setsompop K, Pauly JM, Mardani M, Vasanawala S (2023) Variational diffusion models for blind MRI inverse problems. In: NeurIPS 2023 workshop on deep learning and inverse problems
110. Levac B, Jalal A, Ramchandran K, Tamir J (2023) MRI reconstruction with side information using diffusion models. arXiv preprint [arXiv:2303.14795](https://arxiv.org/abs/2303.14795)
111. Ulyanov D, Vedaldi A, Lempitsky V (2020) Deep image prior. *Int J Comput Vis* 128:1867–1888
112. Kingma DP, Ba JL (2015) Adam: a method for stochastic optimization. In: International conference on learning representations 2015, pp 1–15
113. Heckel R, Soltanolkotabi M (2020) Denoising and regularization via exploiting the structural bias of convolutional generators. In: International conference on learning representations
114. Heckel R, Hand P (2019) Deep decoder: concise image representations from untrained non-convolutional networks. In: International conference on learning representations
115. Heckel R, Soltanolkotabi M (2020) Compressive sensing with un-trained neural networks: gradient descent finds the smoothest approximation. In: International conference on machine learning
116. Darestani MZ, Heckel R (2021) Accelerated MRI with un-trained neural networks. *IEEE Trans Comput Imaging* 7:724–733
117. Jin KH, Gupta H, Yerly J, Stuber M, Unser M (2019) Time-dependent deep image prior for dynamic MRI. [arXiv:1910.01684](https://arxiv.org/abs/1910.01684)
118. Zou Q, Ahmed AH, Nagpal P, Jacob M (2021) Dynamic imaging using a deep generative SToRM (GenSToRM) model. *IEEE Trans Med Imaging* 40:3102–3112
119. Ahmed AH, Zhou R, Yang Y, Nagpal P, Salerno M, Jacob M (2020) Free-breathing and ungated dynamic MRI using navigator-less spiral storm. *IEEE Trans Med Imaging* 39(12):3933–3943. <https://doi.org/10.1109/TMI.2020.3008329>
120. Poddar S, Mohsin YQ, Ansah D, Thattaliyath B, Ashwath R, Jacob M (2019) Manifold recovery using kernel low-rank regularization: application to dynamic imaging. *IEEE Trans Comput Imaging* 5(3):478–491
121. Mildenhall B, Srinivasan PP, Tancik M, Barron JT, Ramamoorthi R, Ng R (2021) NeRF: representing scenes as neural radiance fields for view synthesis. *Commun ACM* 65(1):99–106
122. Arandjelović R, Zisserman, A (2021) NeRF in detail: learning to sample for view synthesis. [arXiv:2106.05264](https://arxiv.org/abs/2106.05264)
123. Williams F, Schneider T, Silva C, Zorin D, Bruna J, Panozzo D (2019) Deep geometric prior for surface reconstruction. In: Conference on computer vision and pattern recognition
124. Huang W, Li HB, Pan J, Cruz G, Rueckert D, Hammernik K (2023) Neural implicit k-space for binning-free non-cartesian cardiac MR imaging. In: International conference on information processing in medical imaging, Springer, pp 548–560
125. Tancik M, Srinivasan P, Mildenhall B, Fridovich-Keil S, Raghavan N, Singhal U, Ramamoorthi R, Barron J, Ng R (2020) Fourier features let networks learn high frequency functions in low dimensional domains. *Adv Neural Inf Process Syst* 33:7537–7547
126. Dupont E, Goliński A, Alizadeh M, Teh YW, Doucet A (2021) COIN: COmpression with Implicit Neural Representations. [arXiv:2103.03123](https://arxiv.org/abs/2103.03123) [cs, eess]
127. Catalán T, Courdurier M, Osses A, Botnar R, Costabal FS, Prieto C (2023) Unsupervised reconstruction of accelerated cardiac cine MRI using neural fields. arXiv preprint [arXiv:2307.14363](https://arxiv.org/abs/2307.14363)
128. Spieker V, Huang W, Eichhorn H, Stelter J, Weiss K, Zimmer VA, Braren RF, Karampinos DC, Hammernik K, Schnabel JA (2023) ICOnIK: generating respiratory-resolved abdominal MR reconstructions using neural implicit representations in k-space. In: International conference on medical image computing and computer-assisted intervention, Springer, , pp 183–192
129. Feng R, Wu Q, Feng J, She H, Liu C, Zhang Y, Wei H (2023) IMJENSE: scan-specific implicit representation for joint coil sensitivity and image estimation in parallel MRI. *IEEE Trans Med Imaging* 43:1539–1553
130. Sitzmann V, Martel JNP, Bergman AW, Lindell DB, Wetzstein G (2020) Implicit neural representations with periodic activation functions. [arXiv:2006.09661](https://arxiv.org/abs/2006.09661) [cs, eess]
131. Kunz JF, Ruschke S, Heckel R (2023) Implicit neural networks with fourier-feature inputs for free-breathing cardiac MRI reconstruction. arXiv
132. Yaman B, Hosseini SAH, Moeller S, Ellermann J, Ugurbil K, Akcakaya M (2020) Self-supervised physics-based deep learning MRI reconstruction without fully-sampled data. *Magn Reson Med*. <https://doi.org/10.1002/mrm.28378>
133. Millard C, Chiew M (2022) Simultaneous self-supervised reconstruction and denoising of sub-sampled MRI data with Noisier2Noise. arXiv preprint [arXiv:2210.01696](https://arxiv.org/abs/2210.01696)
134. Chen Y, Holmes JH, Corum C, Magnotta V, Jacob M (2023) Deep factor model: a novel approach for motion compensated multi-dimensional MRI. In: IEEE international symposium on biomedical imaging
135. Wang F, Qi H, De Goyeneche A, Heckel R, Lustig M, Shimron E (2023) K-band: self-supervised MRI reconstruction via stochastic gradient descent over k-space subsets. arXiv preprint [arXiv:2308.02958](https://arxiv.org/abs/2308.02958)
136. Akçakaya M, Yaman B, Chung H, Ye JC (2022) Unsupervised deep learning methods for biological image reconstruction and enhancement: an overview from a signal processing perspective. *IEEE Signal Process Mag* 39(2):28–44
137. Zeng G, Guo Y, Zhan J, Wang Z, Lai Z, Du X, Qu X, Guo D (2021) A review on deep learning MRI reconstruction without fully sampled k-space. *BMC Med Imaging* 21(1):195
138. Stein CM (1981) Estimation of the mean of a multivariate normal distribution. *Ann Stat* 9:1135–1151
139. Ramani S, Blu T, Unser M (2008) Monte-Carlo SURE: a black-box optimization of regularization parameters for general denoising algorithms. *IEEE Trans Image Process* 17(9):1540–1554
140. Metzler CA, Mousavi A, Heckel R, Baraniuk RG (2018) Unsupervised learning with Stein’s unbiased risk estimator. arXiv preprint [arXiv:1805.10531](https://arxiv.org/abs/1805.10531)
141. Zhussip M, Soltanayev S, Chun SY (2019) Training deep learning based image denoisers from undersampled measurements without ground truth and without image prior. In: IEEE/CVF conference on computer vision and pattern recognition, pp. 10255–10264
142. Eldar YC (2008) Generalized SURE for exponential families: applications to regularization. *IEEE Trans Signal Process* 57(2):471–481
143. Donoho DL, Johnstone IM (1995) Adapting to unknown smoothness via wavelet shrinkage. *J Am Stat Assoc* 90(432):1200–1224
144. Lehtinen J, Munkberg J, Hasselgren J, Laine S, Karras T, Aittala M, Aila T (2018) Noise2Noise: learning image restoration without clean data. In: International conference on machine learning, pp 2971–2980
145. Klug T, Atik D, Heckel R (2023) Analyzing the sample complexity of self-supervised image reconstruction methods. In: Neural information processing systems (NeurIPS)
146. Millard C, Chiew M (2023) A theoretical framework for self-supervised MR image reconstruction using sub-sampling via variable density Noisier2Noise. *IEEE Trans Comput Imaging* 9:707–720
147. Moran N, Schmidt D, Zhong Y, Coady P (2020) Noisier2Noise: learning to denoise from unpaired noisy data. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 12064–12072

148. Huang W, Li H, Cruz G, Pan J, Rueckert D, Hammernik K (2022) Neural implicit k-space for binning-free non-cartesian cardiac MR imaging. arXiv
149. Han K, Wang Y, Chen H, Chen X, Guo J, Liu Z, Tang Y, Xiao A, Xu C, Xu Y et al (2022) A survey on vision transformer. *IEEE Trans Pattern Anal Mach Intell* 45(1):87–110
150. Han K, Xiao A, Wu E, Guo J, Xu C, Wang Y (2021) Transformer in transformer. *Adv Neural Inf Process Syst* 34:15908–15919
151. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017) Attention is all you need. In: *Advances in neural information processing systems* 30
152. Gillioz A, Casas J, Mugellini E, Abou Khaled O (2020) Overview of the transformer-based models for NLP tasks. In: *2020 15th conference on computer science and information systems (FedCSIS)*, IEEE. pp 179–183
153. Wolf T, Debut L, Sanh V, Chaumond J, Delangue C, Moi A, Cistac P, Rault T, Louf R, Funtowicz M et al (2019) Hugging-face's transformers: state-of-the-art natural language processing. arXiv preprint [arXiv:1910.03771](https://arxiv.org/abs/1910.03771)
154. Korkmaz Y, Yurt M, Dar SUH, Özbey M, Cukur T (2021) Deep MRI reconstruction with generative vision transformers. *Lecture notes in computer science (including subseries lecture notes in artificial intelligence and lecture notes in bioinformatics)* 12964 LNCS, pp 54–64
155. Feng C-M, Yan Y, Fu H, Chen L, Xu Y (2021) Task transformer network for joint MRI reconstruction and super-resolution. In: *Medical image computing and computer assisted intervention—MICCAI 2021: 24th international conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI* 24, Springer. pp 307–317
156. Souza R, Bento M, Nogovitsyn N, Chung KJ, Loos W, Lebel RM, Frayne R (2020) Dual-domain cascade of U-nets for multi-channel magnetic resonance image reconstruction. *Magn Reson Imaging* 71:140–153
157. Wang S, Ke Z, Cheng H, Jia S, Ying L, Zheng H, Liang D (2022) DIMENSION: dynamic MR imaging with both k-space and spatial prior knowledge obtained via multi-supervised network training. *NMR Biomed* 35(4):4131
158. Ran M, Xia W, Huang Y, Lu Z, Bao P, Liu Y, Sun H, Zhou J, Zhang Y (2020) MD-Recon-Net: a parallel dual-domain convolutional neural network for compressed sensing MRI. *IEEE Trans Radiat Plasma Med Sci* 5(1):120–135
159. Singh NM, Iglesias JE, Adalsteinnsson E, Dalca AV, Golland P (2022) Joint frequency and image space learning for MRI reconstruction and analysis. *J Mach Learn Biomed Imaging*. <https://doi.org/10.59275/j.melba.2022-16cc>
160. Zhao X, Yang T, Li B, Zhang X (2023) SwinGAN: a dual-domain Swin transformer-based generative adversarial network for MRI reconstruction. *Comput Biol Med* 153:106513
161. Korkmaz Y, Cukur T, Patel VM (2023) Self-supervised MRI reconstruction with unrolled diffusion models. In: *International conference on medical image computing and computer-assisted intervention*, Springer. pp 491–501
162. Zhao X, Yang T, Li B, Yang A, Yan Y, Jiao C (2024) DiffGAN: an adversarial diffusion model with local transformer for MRI reconstruction. *Magn Reson Imaging* 109:108–119
163. Gao Y, Reeves SJ (2000) Optimal k-space sampling in MRSI for images with a limited region of support. *IEEE Trans Med Imaging* 19(12):1168–1178
164. Xu D, Jacob M, Liang Z (2005) Optimal sampling of k-space with Cartesian grids for parallel MR imaging. In: *Proceedings of the international society for magnetic resonance in medicine*, vol 13. p 2450
165. Haldar JP, Kim D (2019) OEDIPUS: an experiment design framework for sparsity-constrained MRI. *IEEE Trans Med Imaging* 38:1545–1558
166. Levine E, Hargreaves B (2017) On-the-fly adaptive k-space sampling for linear MRI reconstruction using moment-based spectral analysis. *IEEE Trans Med Imaging* 37(2):557–567
167. Senel LK, Kilic T, Gungor A, Kopanoglu E, Guven HE, Saritas EU, Koc A, Çukur T (2019) Statistically segregated k-space sampling for accelerating multiple-acquisition MRI. *IEEE Trans Med Imaging* 38(7):1701–1714
168. Sherry F, Benning M, Reyes JCDI, Graves MJ, Maierhofer G, Williams G, Schönlieb C-B, Ehrhardt MJ (2019) Learning the sampling pattern for MRI. arXiv preprint [arXiv:1906.08754](https://arxiv.org/abs/1906.08754)
169. Gözcü B, Mahabadi RK, Li Y-H, Ilıcak E, Çukur T, Scarlett J, Cevher V (2018) Learning-based compressive MRI. *IEEE Trans Med Imaging* 37(6):1394–1406
170. Lazarus C, Weiss P, Chauffert N, Mauconduit F, El Gueddari L, Destrieux C, Zemmoura I, Vignaud A, Ciuciu P (2019) SPARKLING: variable-density k-space filling curves for accelerated T2*-weighted MRI. *Magn Reson Med* 81(6):3643–3661
171. Ravula S, Levac B, Jalal A, Tamir JI, Dimakis AG (2023) Optimizing sampling patterns for compressed sensing MRI with diffusion generative models. arXiv preprint [arXiv:2306.03284](https://arxiv.org/abs/2306.03284)
172. Weiss T, Senouf O, Vedula S, Michailovich O, Zibulevsky M, Bronstein A (2019) PILOT: physics-informed learned optimized trajectories for accelerated MRI. *Mach Learn Biomed Imaging* 1:1–23
173. Bahadır CD, Wang AQ, Dalca AV, Sabuncu MR (2020) Deep-learning-based optimization of the under-sampling pattern in MRI. *IEEE Trans Comput Imaging* 6:1139–1152
174. Chaithya GR, Ramzi Z, Ciuciu P (2021) Learning the sampling density in 2D SPARKLING MRI acquisition for optimized image reconstruction. In: *European signal processing conference 2021-August*, pp 960–964
175. Zibetti MVW, Knoll F, Regatte RR (2022) Alternating learning approach for variational networks and undersampling pattern in parallel MRI applications. *IEEE Trans Comput Imaging* 8:449–461
176. Aggarwal HK, Jacob M (2020) J-MoDL: joint model-based deep learning for optimized sampling and reconstruction. *IEEE J Select Top Signal Process* 14(6):1151–1162
177. Wang G, Luo T, Nielsen JF, Noll DC, Fessler JA (2022) B-spline parameterized joint optimization of reconstruction and k-space trajectories (BJORK) for accelerated 2D MRI. *IEEE Trans Med Imaging* 41:2318–2330
178. Wang G, Nielsen J-F, Fessler JA, Noll DC (2022) Stochastic optimization of 3D non-cartesian sampling trajectory (SNOPY). arXiv preprint [arXiv:2209.11030](https://arxiv.org/abs/2209.11030)
179. Radhakrishna CG, Ciuciu P (2023) Jointly learning non-cartesian k-space trajectories and reconstruction networks for 2D and 3D MR imaging through projection. *Bioengineering* 10:158
180. Alkan C, Mardani M, Vasanawala SS, Pauly JM (2020) Learning to Sample MRI via Variational Information Maximization. *NeurIPS 2020 Workshop on Deep Learning and Inverse Problems*.
181. Xie J, Zhang J, Zhang Y, Ji X (2022) PUERT: probabilistic under-sampling and explicable reconstruction network for CS-MRI. *IEEE J Select Top Signal Process* 16(4):737–749
182. Zou J, Cao Y (2022) Joint optimization of kt sampling pattern and reconstruction of DCE MRI for pharmacokinetic parameter estimation. *IEEE Trans Med Imaging* 41(11):3320–3331
183. Zhu B, Liu J, Koonjoo N, Rosen BR, Rosen MS (2018) AUTOMated pulse SEquence generation (AUTOSEQ) using Bayesian reinforcement learning in an MRI physics simulation environment. In: *Proceedings of joint annual meeting ISMRM-ESMRMB*
184. Zhu B, Liu J, Koonjoo N, Rosen B, Rosen MS (2019) AUTOMated pulse SEquence generation (AUTOSEQ) and neural network decoding for fast quantitative MR parameter measurement

- using continuous and simultaneous RF transmit and receive. In: ISMRM annual meeting & exhibition, vol 1090
185. Lee PK, Watkins LE, Anderson TI, Buonincontri G, Hargreaves BA (2019) Flexible and efficient optimization of quantitative sequences using automatic differentiation of Bloch simulations. *Magn Reson Med* 82(4):1438–1451
 186. Loktyushin A, Herz K, Dang N, Glang F, Deshmane A, Weinmüller S, Doerfler A, Schölkopf B, Scheffler K, Zaiss M (2021) MRzero-automated discovery of MRI sequences using supervised learning. *Magn Reson Med* 86(2):709–724
 187. Dang HN, Endres J, Weinmüller S, Glang F, Loktyushin A, Scheffler K, Doerfler A, Schmidt M, Maier A, Zaiss M (2023) MR-zero meets RARE MRI: Joint optimization of refocusing flip angles and neural networks to minimize T2-induced blurring in spin echo sequences. *Magn Reson Med* 90(4):1345–1362
 188. Perlman O, Zhu B, Zaiss M, Rosen MS, Farrar CT (2022) An end-to-end AI-based framework for automated discovery of rapid CEST/MT MRI acquisition protocols and molecular parameter quantification (AutoCEST). *Magn Reson Med* 87(6):2792–2810
 189. Kang B, Singh M, Park H, Heo H-Y (2023) Only-train-once MR fingerprinting for B0 and B1 inhomogeneity correction in quantitative magnetization-transfer contrast. *Magn Reson Med* 90(1):90–102
 190. Beracha I, Seginer A, Tal A (2023) Adaptive model-based magnetic resonance. *Magn Reson Med* 90:839–851
 191. Zhao B, Setsompop K, Ye H, Cauley SF, Wald LL (2016) Maximum likelihood reconstruction for magnetic resonance fingerprinting. *IEEE Trans Med Imaging* 35(8):1812–1823
 192. Seiberlich N, Gulani V, Campbell-Washburn A, Sourbron S, Doneva MI, Calamante F, Hu HH (2020) Quantitative magnetic resonance imaging. Elsevier, Amsterdam
 193. Vladimirov N, Perlman O (2023) Molecular MRI-based monitoring of cancer immunotherapy treatment response. *Int J Mol Sci* 24(4):3151
 194. Feng L, Ma D, Liu F (2022) Rapid MR relaxometry using deep learning: an overview of current techniques and emerging trends. *NMR Biomed* 35(4):4416
 195. Liu F, Feng L, Kijowski R (2019) MANTIS: model-augmented neural network with incoherent k-space sampling for efficient MR parameter mapping. *Magn Reson Med* 82(1):174–188
 196. Liu F, Kijowski R, El Fakhri G, Feng L (2021) Magnetic resonance parameter mapping using model-guided self-supervised deep learning. *Magn Reson Med* 85(6):3211–3226
 197. Li H, Yang M, Kim JH, Zhang C, Liu R, Huang P, Liang D, Zhang X, Li X, Ying L (2023) SuperMAP: deep ultrafast MR relaxometry with joint spatiotemporal undersampling. *Magn Reson Med* 89(1):64–76
 198. Ma D, Gulani V, Seiberlich N, Liu K, Sunshine JL, Duerk JL, Griswold MA (2013) Magnetic resonance fingerprinting. *Nature* 495(7440):187–192
 199. Weigand-Whittier J, Sedykh M, Herz K, Coll-Font J, Foster AN, Gerstner ER, Nguyen C, Zaiss M, Farrar CT, Perlman O (2023) Accelerated and quantitative three-dimensional molecular MRI using a generative adversarial network. *Magn Reson Med* 89(5):1901–1914
 200. Panda A, Mehta BB, Coppo S, Jiang Y, Ma D, Seiberlich N, Griswold MA, Gulani V (2017) Magnetic resonance fingerprinting—an overview. *Curr Opin Biomed Eng* 3:56–66
 201. Cohen O, Zhu B, Rosen MS (2018) MR fingerprinting deep reconstruction network (DRONE). *Magn Reson Med* 80(3):885–894
 202. Balsiger F, Shridhar Konar A, Chikop S, Chandran V, Scheidegger O, Geethanath S, Reyes M (2018) Magnetic resonance fingerprinting reconstruction via spatiotemporal convolutional neural networks. In: Machine learning for medical image reconstruction: first international workshop, MLMIR 2018, held in conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Proceedings 1., Springer. pp 39–46
 203. Gómez PA, Molina-Romero M, Ulas C, Buonincontri G, Sperl JI, Jones DK, Menzel MI, Menze BH (2016) Simultaneous parameter mapping, modality synthesis, and anatomical labeling of the brain with MR fingerprinting. In: Medical image computing and computer-assisted intervention-MICCAI 2016: 19th International conference, Athens, Greece, October 17–21, 2016, Proceedings, Part III 19, Springer. pp 579–586
 204. Gómez PA, Cencini M, Golbabaee M, Schulte RF, Pirkel C, Horvath I, Fallo G, Peretti L, Tosetti M, Menze BH et al (2020) Rapid three-dimensional multiparametric MRI with quantitative transient-state imaging. *Sci Rep* 10(1):13769
 205. Cohen O, Yu VY, Tringale KR, Young RJ, Perlman O, Farrar CT, Otazo R (2023) CEST MR fingerprinting (CEST-MRF) for brain tumor quantification using EPI readout and deep learning reconstruction. *Magn Reson Med* 89(1):233–249
 206. Perlman O, Ito H, Herz K, Shono N, Nakashima H, Zaiss M, Chiocca EA, Cohen O, Rosen MS, Farrar CT (2022) Quantitative imaging of apoptosis following oncolytic virotherapy by magnetic resonance fingerprinting aided by deep learning. *Nat Biomed Eng* 6(5):648–657
 207. Perlman O, Herz K, Zaiss M, Cohen O, Rosen MS, Farrar CT (2020) CEST MR-fingerprinting: practical considerations and insights for acquisition schedule design and improved reconstruction. *Magn Reson Medicine* 83(2):462–478
 208. Nagar D, Vladimirov N, Farrar CT, Perlman O (2023) Dynamic and rapid deep synthesis of chemical exchange saturation transfer and semisolid magnetization transfer MRI signals. *Sci Rep* 13(1):18291
 209. Singh M, Jiang S, Li Y, Zijl P, Zhou J, Heo H-Y (2023) Bloch simulator-driven deep recurrent neural network for magnetization transfer contrast MR fingerprinting and CEST imaging. *Magn Reson Med* 90:1518–1536
 210. Blystad I, Warntjes JBM, Smedby O, Landt-blom A-M, Lundberg P, Larsson E-M (2012) Synthetic MRI of the brain in a clinical setting. *Acta Radiol* 53(10):1158–1163
 211. Perlman O, Farrar CT, Heo H-Y (2023) MR fingerprinting for semisolid magnetization transfer and chemical exchange saturation transfer quantification. *NMR Biomed* 36(6):4710
 212. Wang K, Doneva M, Meineke J, Amthor T, Karasan E, Tan F, Tamir JI, Yu SX, Lustig M (2023) High-fidelity direct contrast synthesis from magnetic resonance fingerprinting. *Magn Reson Med* 90:2116–2129
 213. Küstner T et al (2020) CINENet: deep learning-based 3D cardiac CINE MRI reconstruction with multi-coil complex-valued 4D spatio-temporal convolutions. *Sci Rep* 10:1–13
 214. Qin C, Schlemper J, Caballero J, Price AN, Hajnal JV, Rueckert D (2019) Convolutional recurrent neural networks for dynamic MR image reconstruction. *IEEE Trans Med Imaging* 38:280–290
 215. Arefeen Y, Xu J, Zhang M, Dong Z, Wang F, White J, Bilgic B, Adalsteinsson E (2023) Latent signal models: learning compact representations of signal evolution for improved time-resolved, multi-contrast MRI. *Magn Reson Med* 90:483–501
 216. Biswas S, Aggarwal HK, Jacob M (2019) Dynamic MRI using model-based deep learning and STORM priors: MoDL-STORM. *Magn Reson Med* 82:485–494
 217. Yoo J, Jin KH, Gupta H, Yerly J, Stuber M, Unser M (2021) Time-dependent deep image prior for dynamic MRI. *IEEE Trans Med Imaging* 40(12):3337–3348
 218. Sandino CM, Ong F, Iyer SS, Bush A, Vasana-wala S (2021) Deep subspace learning for efficient reconstruction of spatiotemporal imaging data. In: NeurIPS 2021 workshop on deep learning and inverse problems
 219. Qin C, Bai W, Schlemper J, Petersen SE, Piechnik SK, Neubauer S, Rueckert D (2018) Joint learning of motion estimation

- and segmentation for cardiac MR image sequences. In: Medical image computing and computer assisted intervention–MICCAI 2018: 21st international conference, Granada, Spain, September 16–20, 2018, Proceedings, Part II 11, Springer. pp 472–480
220. Ghadimi S, Auger DA, Feng X, Sun C, Meyer CH, Bilchick KC, Cao JJ, Scott AD, Oshinski JN, Ennis DB et al (2021) Fully-automated global and segmental strain analysis of DENSE cardiovascular magnetic resonance using deep learning for segmentation and phase unwrapping. *J Cardiovasc Magn Reson* 23:1–13
 221. Scannell CM, Veta M, Villa AD, Sammut EC, Lee J, Breeuwer M, Chiribiri A (2020) Deep-learning-based preprocessing for quantitative myocardial perfusion MRI. *J Magn Reson Imaging* 51(6):1689–1696
 222. Zucker EJ, Sandino CM, Kino A, Lai P, Vasanaawala SS (2021) Free-breathing accelerated cardiac MRI using deep learning: validation in children and young adults. *Radiology* 300(3):539–548
 223. Zou Q, Ahmed AH, Nagal P, Kruger S, Jacob M (2021) Alignment & joint recovery of multi-slice dynamic MRI using deep generative manifold model. arxiv.org/abs/2101.08196
 224. Zou Q, Ahmed AH, Nagpal P, Priya S, Schulte RF, Jacob M (2022) Variational manifold learning from incomplete data: application to multislice dynamic MRI. *IEEE Trans Med Imaging* 41(12):3552–3561
 225. Sermesant M, Delingette H, Cochet H, Jais P, Ayache N (2021) Applications of artificial intelligence in cardiovascular imaging. *Nat Rev Cardiol* 18(8):600–609
 226. Chang Y, Li Z, Saju G, Mao H, Liu T (2023) Deep learning-based rigid motion correction for magnetic resonance imaging: a survey. *Meta-Radiology* 1:100001
 227. Zhao X, Zhao X-M (2021) Deep learning of brain magnetic resonance images: a brief review. *Methods* 192:131–140
 228. Haskell MW, Cauley SF, Bilgic B, Hossbach J, Splitthoff DN, Pfeuffer J, Setsompop K, Wald LL (2019) Network accelerated motion estimation and reduction (NAMER): convolutional neural network guided retrospective motion correction using a separable motion model. *Magn Reson Med* 82:1452–1461
 229. Singh NM, Dey N, Hoffmann M, Fischl B, Adalsteinsson E, Frost R, Dalca AV, Golland P (2023) Data consistent deep rigid MRI motion correction. *Proc Mac Learn Res* 177:1–14
 230. Johnson PM, Drangova M (2019) Conditional generative adversarial network for 3D rigid-body motion correction in MRI. *Magn Reson Med* 82(3):901–910
 231. Eichhorn H, Hammernik K, Spieker V, Epp SM, Rueckert D, Preibisch C, Schnabel JA (2023) Physics-aware motion simulation for T2*-weighted brain MRI. International Workshop on Simulation and Synthesis in Medical Imaging. Cham: Springer Nature Switzerland, 2023.
 232. Levac B, Jalal A, Tamir JI (2023) Accelerated motion correction for MRI using score-based generative models. In: 2023 IEEE 20th international symposium on biomedical imaging (ISBI), IEEE. pp 1–5
 233. Küstner T, Pan J, Gilliam C, Qi H, Cruz G, Hammernik K, Yang B, Blu T, Rueckert D, Botnar R et al (2020) Deep-learning based motion-corrected image reconstruction in 4D magnetic resonance imaging of the body trunk. In: 2020 Asia-Pacific signal and information processing association annual summit and conference (APSIPA ASC), IEEE. pp 976–985
 234. Wang X, Uecker M, Feng L (2021) Fast real-time cardiac MRI: a review of current techniques and future directions. *Investig Magn Reson Imaging* 25(4):252–265
 235. Singh A, Salehi SSM, Gholipour A (2020) Deep predictive motion tracking in magnetic resonance imaging: application to fetal imaging. *IEEE Trans Med Imaging* 39(11):3523–3534
 236. Shimron E, De Goyeneche A, Wang K, Halgren A, Syed AB, Vasanaawala S, Lustig M (2022) BladeNet: rapid PROPELLER acquisition and reconstruction for high spatio-temporal resolution abdominal MRI. In: Proceedings of the 31st annual international society for magnetic resonance in medicine, London, UK, pp 7–12
 237. Murray V, Siddiq S, Crane C, Homsy ME, Kim TH, Wu C, Otazo R (2024) Movienet: deep space-time-coil reconstruction network without k-space data consistency for fast motion-resolved 4D MRI. *Magn Reson Med* 91:600–614
 238. Malavé MO, Baron CA, Koundinyan SP, Sandino CM, Ong F, Cheng JY, Nishimura DG (2020) Reconstruction of undersampled 3D non-Cartesian image-based navigators for coronary MRA using an unrolled deep learning model. *Magn Reson Med* 84(2):800–812
 239. Zou Q, Torres LA, Fain SB, Higano NS, Bates AJ, Jacob M (2022) Dynamic imaging using motion-compensated smoothness regularization on manifolds (MoCo-SToRM). *Phys Med Biol* 67(14):144001
 240. Freedman JN, Gurney-Champion OJ, Nill S, Shiarli A-M, Bainbridge HE, Mandeville HC, Koh D-M, McDonald F, Kachelrieß M, Oelfke U et al (2021) Rapid 4D-MRI reconstruction using a deep radial convolutional neural network: dracula. *Radiother Oncol* 159:209–217
 241. Terpstra ML, Maspero M, d’Agata F, Stemkens B, Intven MP, Legendijk JJ, Berg CA, Tijssen RH (2020) Deep learning-based image reconstruction and motion estimation from undersampled radial k-space for real-time MRI-guided radiotherapy. *Phys Med Biol* 65(15):155015
 242. Waddington DE, Hindley N, Koonjoo N, Chiu C, Reynolds T, Liu PZ, Zhu B, Bhutto D, Paganelli C, Keall PJ et al (2023) Real-time radial reconstruction with domain transform manifold learning for MRI-guided radiotherapy. *Med Phys* 50(4):1962–1974
 243. Sandino CM, Cole EK, Alkan C, Chaudhari AS, Loening AM, Hyun D, Dahl J, Imran A-A-Z, Wang AS, Vasanaawala SS (2021) Upstream machine learning in radiology. *Radiol Clin N Am* 59(6):967–985. <https://doi.org/10.1016/j.rcl.2021.07.009>
 244. Chaudhari AS, Sandino CM, Cole EK, Larson DB, Gold GE, Vasanaawala SS, Lungren MP, Hargreaves BA, Langlotz CP (2020) Prospective deployment of deep learning in MRI: a framework for important considerations, challenges, and recommendations for best practices. *J Magn Reson Imaging* 54(2):357–371. <https://doi.org/10.1002/jmri.27331>
 245. Zhao R, Yaman B, Zhang Y, Stewart R, Dixon A, Knoll F, Huang Z, Lui YW, Hansen MS, Lungren MP (2022) fastMRI+, clinical pathology annotations for knee and brain fully sampled magnetic resonance imaging data. *Sci Data*. <https://doi.org/10.1038/s41597-022-01255-z>
 246. Zhao R, Zhang Y, Yaman B, Lungren MP, Hansen MS (2021) End-to-end AI-based MRI reconstruction and lesion detection pipeline for evaluation of deep learning image reconstruction. *arXiv preprint arXiv:2109.11524*.
 247. Weber T, Ingrisich M, Bischl B, Rügamer D (2024) Constrained probabilistic mask learning for task-specific undersampled MRI reconstruction. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision (WACV), pp 7665–7674
 248. Desai AD, Caliva F, Iriondo C, Mortazi A, Jambawalikar S, Bagci U, Perslev M, Igel C, Dam EB, Gaj S, Yang M, Li X, Deniz CM, Juras V, Regatte R, Gold GE, Hargreaves BA, Podoia V, Chaudhari AS, Khosravan N, Torigian D, Ellermann J, Akcakaya M, Tibrewala R, Flament I, O’Brien M, Majumdar S, Nakamura K, Pai A (2021) The international workshop on osteoarthritis imaging knee MRI segmentation challenge: a multi-institute evaluation and analysis framework on a standardized dataset. *Radiol Artif Intell* 3(3):200078. <https://doi.org/10.1148/ryai.2021200078>
 249. Schmidt AM, Desai AD, Watkins LE, Crowder HA, Black MS, Mazzoli V, Rubin EB, Lu Q, MacKay JW, Boutin RD, Kogan F, Gold GE, Hargreaves BA, Chaudhari AS (2022) Generalizability

- of deep learning segmentation algorithms for automated assessment of cartilage morphology and MRI relaxometry. *J Magn Reson Imaging* 57(4):1029–1039. <https://doi.org/10.1002/jmri.28365>
250. Wu Z, Yin T, Sun Y, Frost R, Kouwe A, Dalca AV, Bouman KL (2023) Learning task-specific strategies for accelerated MRI. *IEEE Trans Comput Imaging*
 251. Uecker M et al (2014) ESPIRiT—an eigenvalue approach to auto-calibrating parallel MRI: where SENSE meets GRAPPA. *Magn Reson Med* 71(3):990–1001
 252. Uecker M, Hohage T, Block KT, Frahm J (2008) Image reconstruction by regularized nonlinear inversion-joint estimation of coil sensitivities and image content. *Magn Reson Med* 60(3):674–682
 253. Ying L, Sheng J (2007) Joint image reconstruction and sensitivity estimation in SENSE (JSENSE). *Magn Reson Med* 57(6):1196–1202
 254. Sriram A, Zbontar J, Murrell T, Defazio A, Zitnick CL, Yakubova N, Knoll F, Johnson P (2020) End-to-end variational networks for accelerated MRI reconstruction. In: Medical image computing and computer assisted intervention—MICCAI 2020: 23rd international conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II 23, Springer. pp 64–73
 255. Jun Y, Shin H, Eo T, Hwang D (2021) Joint deep model-based MR image and coil sensitivity reconstruction network (joint-ICNet) for fast MRI. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 5270–5279
 256. Luo G, Wang X, Roeloffs V, Tan Z, Uecker M (2021) Joint estimation of coil sensitivities and image content using a deep image prior. In: Proceedings of the 29th annual meeting of ISMRM, p 280
 257. Hu Y, Gan W, Ying C, Wang T, Eldeniz C, Liu J, Chen Y, An H, Kamilov US (2024) SPICER: self-supervised learning for MRI with automatic coil sensitivity estimation and reconstruction. *Magn Reson Med*. <https://doi.org/10.1002/mrm.30121>
 258. Zhang M, Xu J, Arefeen Y, Adalsteinsson E (2023) Zero-shot self-supervised joint temporal image and sensitivity map reconstruction via linear latent space. *arXiv preprint arXiv:2303.02254*
 259. Liu J, Pasumarthi S, Duffy B, Gong E, Datta K, Zaharchuk G (2023) One model to synthesize them all: multi-contrast multi-scale transformer for missing data imputation. *IEEE Trans Med Imaging* 42:2577–2591
 260. Sveinsson B, Chaudhari AS, Zhu B, Koonjoo N, Torriani M, Gold GE, Rosen MS (2021) Synthesizing quantitative T2 maps in right lateral knee femoral condyles from multicontrast anatomic data with a conditional generative adversarial network. *Radiol Artif Intell* 3(5):200122
 261. Zhao C, Shao M, Carass A, Li H, Dewey BE, Ellingsen LM, Woo J, Guttman MA, Blitz AM, Stone M et al (2019) Applications of a deep learning method for anti-aliasing and super-resolution in MRI. *Magn Reson Imaging* 64:132–141
 262. Li Y, Sixou B, Peyrin F (2021) A review of the deep learning methods for medical images super resolution problems. *IRBM* 42(2):120–133
 263. Haskell MW, Nielsen J-F, Noll DC (2023) Off-resonance artifact correction for MRI: a review. *NMR Biomed* 36(5):4867
 264. Ayde R, Senft T, Salameh N, Sarracanie M (2022) Deep learning for fast low-field MRI acquisitions. *Sci Rep* 12(1):11394
 265. Lei K, Syed AB, Zhu X, Pauly JM, Vasanaawala SV (2023) Automated MRI field of view prescription from region of interest prediction by intra-stack attention neural network. *Bioengineering* 10(1):92
 266. Bell LC, Shimron E, Francisco DWSS (2023) Sharing data is essential for the future of AI in medical imaging. *Radiol Artif Intell* 6:e230337
 267. Ramzi Z, Ciuciu P, Starck J-L (2020) Benchmarking deep nets MRI reconstruction models on the FastMRI publicly available dataset. In: ISBI 2020—international symposium on biomedical imaging
 268. Lim Y, Toutios A, Bliesener Y, Tian Y, Lingala SG, Vaz C, Sorensen T, Oh M, Harper S, Chen W et al (2021) A multi-speaker dataset of raw and reconstructed speech production real-time MRI video and 3D volumetric images. *Sci Data* 8(1):1–14
 269. (2020) al., C.C.: OCMR (v1.0)—open-access multi-coil k-space dataset for cardiovascular magnetic resonance imaging. [arXiv:2008.03410v2](https://arxiv.org/abs/2008.03410v2)
 270. Lyu M, Mei L, Huang S, Liu S, Li Y, Yang K, Liu Y, Dong Y, Dong L, Wu EX (2023) M4Raw: a multi-contrast, multi-repetition, multi-channel MRI k-space dataset for low-field MRI research. *Sci Data* 10(1):264
 271. The Human Connectome Project (2011) University of Southern California. <http://www.humanconnectomeproject.org/>. Accessed 22 Mar 2021
 272. IXI dataset (2010) <http://brain-development.org/ixi-dataset/>. Accessed -22 Mar 2021
 273. Menze BH, Jakab A, Bauer S, Kalpathy-Cramer J, Farahani K, Kirby J, Burren Y, Porz N, Slotboom J, Wiest R et al (2014) The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE Trans Med Imaging* 34(10):1993–2024
 274. Ghaffari M, Sowmya A, Oliver R (2019) Automated brain tumor segmentation using multimodal brain scans: a survey based on models submitted to the BraTS 2012–2018 challenges. *IEEE Rev Biomed Eng* 13:156–168
 275. Jack CR Jr, Bernstein MA, Fox NC, Thompson P, Alexander G, Harvey D, Borowski B, Britson PJ, Whitwell JL, Ward C et al (2008) The Alzheimer’s disease neuroimaging initiative MRI (ADNI): methods. *J Magn Reson Imaging* 27(4):685–691
 276. Ollier W, Sprosen T, Peakman T (2005) UK Biobank: from concept to reality. *Pharmacogenomics* 6:639–646
 277. Kurdi B, Lozano S, Banaji MR (2017) Introducing the open affective standardized image set (OASIS). *Behav Res Methods* 49:457–470
 278. Paszke A et al (2019) PyTorch: an imperative style, high-performance deep learning library. In: Advances in neural information processing systems 32
 279. Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A, Dean J, Devin M et al (2015) TensorFlow: large-scale machine learning on heterogeneous systems. <https://www.tensorflow.org/>
 280. Uecker M et al (2015) Berkeley advanced reconstruction toolbox. In: Proceedings of the international society for magnetic resonance in medicine, vol 23
 281. Blumenthal M, Luo G, Schilling M, Holme HCM, Uecker M (2023) Deep, deep learning with BART. *Magn Reson Med* 89(2):678–693
 282. Hansen MS, Sørensen TS (2013) Gadgetron: an open source framework for medical image reconstruction. *Magn Reson Med* 69(6):1768–1776
 283. Knopp T, Grosser M (2021) MRIReco. jl: an MRI reconstruction framework written in Julia. *Magn Reson Medicine* 86(3):1633–1646
 284. Ong F, Lustig M (2019) SigPy: a python package for high performance iterative reconstruction. In: Proceedings of the international society for magnetic resonance in medicine, vol 4819
 285. Layton KJ, Kroboth S, Jia F, Littin S, Yu H, Leupold J, Nielsen J-F, Stöcker T, Zaitsev M (2017) Pulseq: a rapid and hardware-independent pulse sequence prototyping framework. *Magn Reson Med* 77(4):1544–1552
 286. Herz K, Mueller S, Perlman O, Zaitsev M, Knutsson L, Sun PZ, Zhou J, Zijl P, Heinecke K, Schuenke P et al (2021)

- Pulseq-CEST: towards multi-site multi-vendor compatibility and reproducibility of CEST experiments using an open-source sequence standard. *Magn Reson Med* 86(4):1845–1858
287. Ning L, Bonet-Carne E, Grussu F, Sepehrband F, Kaden E, Ver-aart J, Blumberg SB, Khoo CS, Palombo M, Kokkinos I et al (2020) Cross-scanner and cross-protocol multi-shell diffusion MRI data harmonization: algorithms and results. *Neuroimage* 221:117128
288. Karakuzu A, Boudreau M, Duval T, Boshkovski T, Leppert I, Cabana J-F, Gagnon I, Beliveau P, Pike G, Cohen-Adad J et al (2020) qMRLab: quantitative MRI analysis, under one umbrella. *J Open Source Softw* 5(53):2343
289. Inati SJ, Naegele JD, Zwart NR, Roopchansingh V, Lizak MJ, Hansen DC, Liu C-Y, Atkinson D, Kellman P, Kozierke S et al (2017) ISMRM raw data format: a proposed standard for MRI raw datasets. *Magn Reson Med* 77(1):411–421
290. Elmas G, Dar SU, Korkmaz Y, Ceyani E, Susam B, Ozbey M, Avestimehr S, Çukur T (2022) Federated learning of generative image priors for MRI reconstruction. *IEEE Trans Med Imaging* 42:1996–2009
291. Levac BR, Arvinte M, Tamir JI (2023) Federated end-to-end unrolled models for magnetic resonance image reconstruction. *Bioengineering* 10(3):364
292. Johnson PM, Jeong G, Hammernik K, Schlemper J, Qin C, Duan J, Rueckert D, Lee J, Pezzotti N, De Weerd E et al (2021) Evaluation of the robustness of learned MR image reconstruction to systematic deviations between training and test data for the models from the fastMRI challenge. In: *Machine learning for medical image reconstruction: 4th International workshop, MLMIR 2021, held in conjunction with MICCAI 2021, Strasbourg, France, October 1, 2021, Proceedings 4*, Springer. pp 25–34
293. Darestani MZ, Chaudhari A, Heckel R (2021) Measuring robustness in deep learning based compressive sensing. *arXiv preprint arXiv:2102.06103*
294. Avidan N, Freiman M (2022) Physically-primed deep-neural-networks for generalized undersampled MRI reconstruction. *arXiv preprint arXiv:2209.00462*
295. Knoll F, Hammernik K, Kobler E, Pock T, Recht MP, Sodickson DK (2019) Assessment of the generalization of learned image reconstruction and the potential for transfer learning. *Magn Reson Med* 81(1):116–128
296. Huang J, Wang S, Zhou G, Hu W, Yu G (2022) Evaluation on the generalization of a learned convolutional neural network for MRI reconstruction. *Magn Reson Imaging* 87:38–46
297. Dar SUH, Özbey M, Çatlı AB, Çukur T (2020) A transfer-learning approach for accelerated MRI using deep neural networks. *Magn Reson Med* 84(2):663–685
298. Darestani MZ, Liu J, Heckel R (2022) Test-time training can close the natural distribution shift performance gap in deep learning based compressed sensing. In: *International conference on machine learning*, PMLR. pp 4754–4776
299. Lin K, Heckel R (2023) Robustness of deep learning for accelerated MRI: benefits of diverse training data. *arXiv preprint arXiv:2312.10271*
300. Varoquaux G, Cheplygina V (2022) Machine learning for medical imaging: methodological failures and recommendations for the future. *NPJ Digit Med* 5(1):48
301. Seyyed-Kalantari L, Zhang H, McDermott MB, Chen IY, Ghassemi M (2021) Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nat Med* 27(12):2176–2182
302. Chen RJ, Wang JJ, Williamson DF, Chen TY, Lipkova J, Lu MY, Sahai S, Mahmood F (2023) Algorithmic fairness in artificial intelligence for medicine and healthcare. *Nat Biomed Eng* 7(6):719–742
303. Shimron E, Tamir JI, Wang K, Lustig M (2022) Implicit data crimes: machine learning bias arising from misuse of public data. *Proc Natl Acad Sci* 119(13):2117203119
304. Deveshwar N, Rajagopal A, Sahin S, Shimron E, Larson PE (2023) Synthesizing complex-valued multicoil MRI data from magnitude-only images. *Bioengineering* 10(3):358
305. Yang Q, Lin Y, Wang J, Bao J, Wang X, Ma L, Zhou Z, Yang Q, Cai S, He H et al (2022) Model-based synthetic data-driven learning (MOST-DL): application in single-shot T2 mapping with severe head motion using overlapping-echo acquisition. *IEEE Trans Med Imaging* 41(11):3167–3181
306. Luo G, Wang X, Blumenthal M, Schilling M, Rauf EHU, Kotikalapudi R, Focke N, Uecker M (2023) Generative image priors for MRI reconstruction trained from magnitude-only images. *arXiv preprint arXiv:2308.02340*
307. Zhang C, Yang Q, Fan L, Yu S, Sun L, Cai C, Ding X (2023) Towards better generalization using synthetic data: a domain adaptation framework for T2 mapping via multiple overlapping-echo acquisition. *IEEE Trans Med Imaging*. <https://doi.org/10.1109/TMI.2023.3335212>
308. Bhadra S, Kelkar VA, Brooks FJ, Anastasio MA (2021) On hallucinations in tomographic image reconstruction. *IEEE Trans Med Imaging* 40(11):3249–3260
309. Cohen JP, Luck M, Honari S (2018) Distribution matching losses can hallucinate features in medical image translation. In: *Medical image computing and computer assisted intervention—MICCAI 2018: 21st international conference, Granada, Spain, September 16–20, 2018, Proceedings, Part I*, Springer. pp 529–536
310. Gottschling NM, Antun V, Adcock B, Hansen AC (2020) The troublesome kernel: why deep learning for inverse problems is typically unstable. *arXiv preprint arXiv:2001.01258*
311. Huang Y, Würfl T, Breininger K, Liu L, Lauritsch G, Maier A (2018) Some investigations on robustness of deep learning in limited angle tomography. In: *International conference on medical image computing and computer-assisted intervention*, Springer. pp. 145–153
312. Antun V, Renna F, Poon C, Adcock B, Hansen AC (2020) On instabilities of deep learning in image reconstruction and the potential costs of AI. *Proc Natl Acad Sci* 117(48):30088–30095
313. Krainovic A, Soltanolkotabi M, Heckel R (2023) Learning provably robust estimators for inverse problems via jittering. *arXiv preprint arXiv:2308.02340*
314. Morshuis JN, Gatidis S, Hein M, Baumgartner CF (2022) Adversarial robustness of MR image reconstruction under realistic perturbations. In: *International workshop on machine learning for medical image reconstruction*, Springer. pp 24–33
315. Goujon A, Neumayer S, Bohra P, Ducotterd S, Unser M (2023) A neural-network-based convex regularizer for inverse problems. *IEEE Trans Comput Imaging* 9:781–795
316. John M, Chand JR, Jacob M (2023) Local monotone operator learning using non-monotone operators: MnM-MOL. *arXiv preprint arXiv:2312.00386*
317. Calivá F, Cheng K, Shah R, Pedoia V (2020) Adversarial robust training of deep learning MRI reconstruction models. *arXiv preprint arXiv:2011.00070*
318. Cheng K, Calivá F, Shah R, Han M, Majumdar S, Pedoia V (2020) Addressing the false negative problem of deep learning MRI reconstruction models by adversarial attacks and robust training. In: *Medical imaging with deep learning*, PMLR. pp 121–135
319. Wang Z, Bovik AC, Sheikh HR, Simoncelli EP (2004) Image quality assessment: from error visibility to structural similarity. *IEEE Trans Image Process* 13(4):600–612
320. Mason A, Rioux J, Clarke SE, Costa A, Schmidt M, Keough V, Huynh T, Beyea S (2019) Comparison of objective image quality

- metrics to expert radiologists' scoring of diagnostic quality of MR images. *IEEE Trans Med Imaging* 39(4):1064–1072
321. Johnson J, Alahi A, Fei-Fei, L (2016) Perceptual losses for real-time style transfer and super-resolution. In: *Computer vision—ECCV 2016: 14th European conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, Springer. pp 694–711
 322. Adamson PM, Desai AD, Dominic J, Bluethgen C, Wood JP, Syed AB, Boutin RD, Stevens KJ, Vasanawala S, Pauly JM et al (2023) Using deep feature distances for evaluating MR image reconstruction quality. In: *NeurIPS 2023 workshop on deep learning and inverse problems*
 323. Wang K, Tamir JI, Goyeneche AD, Wollner U, Brada R, Yu SX, Lustig M (2022) High fidelity deep learning-based MRI reconstruction with instance-wise discriminative feature matching loss. *Magn Reson Med* 88:476–491
 324. Zhang R, Isola P, Efros AA, Shechtman E, Wang O (2018) The unreasonable effectiveness of deep features as a perceptual metric. *arXiv*
 325. Adamson PM, Gunel B, Dominic J, Desai AD, Spielman D, Vasanawala S, Pauly JM, Chaudhari A (2021) SSFD: self-supervised feature distance as an MR image reconstruction quality metric. In: *NeurIPS 2021 workshop on deep learning and inverse problems*
 326. Chaudhari AS, Kogan F, Pedoia V, Majumdar S, Gold GE, Hargreaves BA (2019) Rapid knee MRI acquisition and analysis techniques for imaging osteoarthritis. *J Magn Reson Imaging* 52(5):1321–1339. <https://doi.org/10.1002/jmri.26991>
 327. Chaudhari AS, Grissom MJ, Fang Z, Sveinsson B, Lee JH, Gold GE, Hargreaves BA, Stevens KJ (2021) Diagnostic accuracy of quantitative multicontrast 5-minute knee MRI using prospective artificial intelligence image quality enhancement. *AJR Am J Roentgenol* 216(6):1614
 328. Tezcan KC, Karani N, Baumgartner CF, Konukoglu E (2022) Sampling possible reconstructions of undersampled acquisitions in MR imaging with a deep learned prior. *IEEE Trans Med Imaging* 41(7):1885–1896
 329. Edupuganti V, Mardani M, Vasanawala S, Pauly J (2020) Uncertainty quantification in deep MRI reconstruction. *IEEE Trans Med Imaging* 40(1):239–250
 330. Narnhofer D, Effland A, Kobler E, Hammernik K, Knoll F, Pock T (2021) Bayesian uncertainty estimation of learned variational MRI reconstruction. *IEEE Trans Med Imaging* 41(2):279–291
 331. Küstner T, Hammernik K, Rueckert D, Hepp T, Gatidis S (2024) Predictive uncertainty in deep learning-based MR image reconstruction using deep ensembles: evaluation on the fastMRI data set. *Magn Reson Med* 2:289–302
 332. Wang K, Angelopoulos A, De Goyeneche A, Kohli A, Shimron E, Yu S, Malik J, Lustig M (2022) Rigorous uncertainty estimation for MRI reconstruction. In: *Proceedings of the proceedings of the 30th annual meeting of ISMRM*, vol 749
 333. Gal Y, Ghahramani Z (2016) Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In: *International conference on machine learning*, PMLR. pp 1050–1059
 334. Zhang Z, Romero A, Muckley MJ, Vincent P, Yang L, Drozdal M (2019) Reducing uncertainty in undersampled MRI reconstruction with active acquisition. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp 2049–2058
 335. Fischer P, Thomas K, Baumgartner CF (2023) Uncertainty estimation and propagation in accelerated MRI reconstruction. In: *International workshop on uncertainty for safe utilization of machine learning in medical imaging*, Springer. pp 84–94
 336. Chen T, Xu B, Zhang C, Guestrin C (2016) Training deep nets with sublinear memory cost. *arXiv preprint [arXiv:1604.06174](https://arxiv.org/abs/1604.06174)*
 337. Kellman M, Zhang K, Markley E, Tamir J, Bostan E, Lustig M, Waller L (2020) Memory-efficient learning for large-scale computational imaging. *IEEE Trans Comput Imaging* 6:1403–1414
 338. Bai S, Kolter JZ, Koltun V (2019) Deep equilibrium models. In: *Advances in neural information processing systems* 32

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.