



Review Article

Reconstruction techniques for accelerating dynamic cardiovascular magnetic resonance imaging

Andrew Phair^{a,1}, René M. Botnar^{a,b,c,d,e}, Claudia Prieto^{a,b,c,*}^a School of Biomedical Engineering and Imaging Sciences, King's College London, London, United Kingdom^b Escuela de Ingeniería, Pontificia Universidad Católica de Chile, Santiago, Chile^c Millennium Institute for Intelligent Healthcare Engineering, Santiago, Chile^d Instituto de Ingeniería Biológica y Médica, Pontificia Universidad Católica de Chile, Santiago, Chile^e Technical University of Munich, Institute of Advanced Study, Munich, Germany

ARTICLE INFO

Keywords:

Dynamic cardiac MRI

MRI cine

MRI reconstruction

Parallel imaging

Spatio-temporal redundancy

Compressed sensing

ABSTRACT

Achieving sufficient spatial and temporal resolution for dynamic applications in cardiovascular magnetic resonance (CMR) imaging is a challenging task due to the inherently slow nature of CMR. In order to accelerate scans and allow improved resolution, much research over the past three decades has been aimed at developing innovative reconstruction methods that can yield high-quality images from reduced amounts of k-space data. In this review, we describe the evolution of these reconstruction techniques, with a particular focus on those advances that have shifted the dynamic reconstruction paradigm as it relates to CMR. This review discusses and explains the fundamental ideas behind the success of modern reconstruction algorithms, including parallel imaging, spatio-temporal redundancies, compressed sensing, low-rank methods and machine learning.

1. Introduction

The primary challenge presented by dynamic cardiovascular magnetic resonance (CMR) imaging is speed; is magnetic resonance imaging (MRI) a fundamentally slow imaging modality, but to be of clinical use dynamic magnetic resonance (MR) images must be acquired with sufficient spatial and temporal resolution. The acceleration of CMR scans, particularly in dynamic imaging applications, is therefore crucial to the success of modern CMR. This has been achieved in three key ways: hardware development, advanced sequence design, and by reducing the amount of data acquired during a scan. It is the third of these approaches that forms the focus of this review, for, as we shall see, when

the amount of acquired data is reduced aliasing and undersampling artifacts are introduced into MR images unless sophisticated image reconstruction techniques are applied.

Reconstructing good-quality images from undersampled data sets is a complicated problem, and has been the focus of much research over more than 30 years. By exploiting spatial and temporal correlations in the data, additional information provided by multi-coil arrays, known prior information about the tissues being scanned or the nature of medical images, estimated motion fields and, most recently, machine learning, image reconstruction algorithms have been devised that continue to push the boundaries of what is possible in dynamic MRI.

Abbreviations: 1D, one-dimensional; 2D, two-dimensional; 3D, three-dimensional; 4D, four-dimensional; 5D, five-dimensional; ACS, auto-calibration signal; BLAST, broad-use linear acquisition speed-up technique; CNN, convolutional neural network; CS, compressed sensing; ECG, electrocardiogram; FFT, fast Fourier transform; FOCUSS, focal underdetermined system solver; FOV, field-of-view; g-factor, geometry factor; GRAPPA, generalized autocalibrating partially parallel acquisitions; iGRASP, iterative golden-angle radial sparse parallel MRI; INRs, implicit neural representations; MASTeR, motion-adaptive spatio-temporal regularization; MC, motion correction; ME, motion estimation; MRF, magnetic resonance fingerprinting; mSENSE, modified SENSE; NUFFT, non-uniform FFT; PCA, principal component analysis; PILS, parallel imaging with localized sensitivities; PSF, point-spread function; RAFT, recurrent all-pairs field transforms; RASCAL, radial self-calibrated; RF, radio-frequency; RIGR, reduced-encoding imaging by generalized-series reconstruction; SENSE, sensitivity encoding; SLR, sparsity and low-rank; SMASH, simultaneous acquisition of spatial harmonics; SNR, signal-to-noise ratio; SPIRiT, iterative self-consistent parallel imaging reconstruction; SSDU, self-supervised via data undersampling; SVD, singular value decomposition; TGRAPPA, temporal GRAPPA; TSENSE, temporal filtering SENSE; TV, total variation; UNFOLD, unaliasing by Fourier-encoding the overlaps using the temporal dimension; XD-GRASP, extra-dimensional iGRASP; CMR, cardiovascular magnetic resonance; PROPELLER, Periodically Rotated Overlapping Parallel Lines with Enhanced Reconstruction

* Corresponding author at: School of Biomedical Engineering and Imaging Sciences, King's College London, London, United Kingdom.

E-mail addresses: andrew.phair@kcl.ac.uk (A. Phair), claudia.prieto@kcl.ac.uk (C. Prieto).

¹ 0000-0003-3527-0724

<https://doi.org/10.1016/j.jocmr.2025.101873>

Received 10 June 2024; Received in revised form 18 December 2024; Accepted 28 February 2025

1097-6647/© 2025 The Authors. Published by Elsevier Inc. on behalf of Society for Cardiovascular Magnetic Resonance. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

In this review, we aim to provide the reader with an understanding of the concepts and ideas that underlie modern reconstruction algorithms. The amount of published literature on this topic is vast. As such, we focus on those particularly prominent methods which have either been incorporated into routine clinical practice or have introduced ideas fundamental to the further development of reconstruction techniques.

The review is organized as follows. In [Section 2](#), the relationship between an MR image and the acquired data is explained and the source of undersampling artifacts is explored. In [Section 3](#), some of the early techniques that were proposed for dynamic image reconstruction are discussed. [Section 4](#) focuses on parallel imaging, and the seminal parallel imaging methods SENSE (sensitivity encoding) and GRAPPA (generalized autocalibrating partially parallel acquisitions). Although these are not specific to dynamic imaging, they are pivotal to many dynamic-specific methods proposed since. [Section 5](#) introduces methods that utilize high-dimensional sampling theory and explicitly exploit spatio-temporal redundancies. In [Section 6](#), iterative reconstructions regularized by sparsity constraints (compressed sensing [CS]) or low-rank assumptions are discussed. In [Section 7](#), the emerging wave of techniques based on deep neural networks is discussed. Finally, some concluding remarks are provided in [Section 8](#). A summary of key methods covered is included in [Table 1](#).

2. Image reconstruction in MRI

During an MRI scan, a radio-frequency (RF) pulse irradiates protons within a slice or volume and causes them to transition to an excited, energetic state. Following the pulse, as the protons relax back to their ground state, they re-emit the energy as RF radiation. It is the strength of this radiation that is measured by the receive RF coils in the scanner, and these measurements form the MRI data set. However, a coil cannot distinguish where the radiation it measures originates from. Rather, it

measures only the combined strength of all the radiation, across the entire imaging volume.

The innovation that made image reconstruction possible in MRI was the use of gradient coils, switched on and off throughout the scan, to Fourier-encode the acquired data. Each data sample still contains information about the entire image, but that information now corresponds to the image modified by a spatial encoding pattern which is induced by the cumulative effect of the gradient coils up to the time at which the sample is acquired. Specifically, the sample records the extent to which a given spatial frequency is present in the image. Large objects that vary smoothly and slowly across the image have low spatial frequency, whereas image texture, with no large-scale structure and rapid variation across the image, has high spatial frequency.

When a Fourier-encoding acquisition scheme is applied, the acquired data samples lie in the Fourier domain, known as k-space in MRI. In two dimensions, k-space is a coordinate system constructed from the k_x and k_y variables, which correspond to spatial frequency in the x- and y-dimensions of the image, respectively.

Depending on the acquisition, an additional third k-space dimension, k_z , may also be considered, corresponding to the third dimension, z, of a three-dimensional (3D) image. Acquiring data in this manner enables the image to be recovered via an inverse Fourier transform, as described in [Appendix A.1](#).

In practice, only a finite number of discrete k-space samples can be obtained, and we only wish to represent the image with a finite number of pixels or voxels (image-space samples). Hence, the discrete inverse Fourier transform is instead utilized. This also allows fast image reconstruction, since for Cartesian sampling it can be implemented very efficiently using the fast Fourier transform (FFT) algorithm [1].

The effect of the discretization can be understood via the convolution theorem, which states that the multiplication of two functions in k-space is equivalent to the image-space convolution of the inverse Fourier transforms of those functions. We consider the discretization to

Table 1
Summary of key selected reconstruction approaches discussed in this review.

Method	Acceleration strategy	Details
Full sequential sampling	None	Each frame acquired as a fully sampled image.
Key-hole	Temporal averaging	Initial acquisition of outer k-space reused throughout scan.
Restricted FOV	Temporal averaging	Motion limited to small image-space FOV.
SENSE	Parallel imaging	Image-space unaliasing algorithm. Requires coil profiles. Auto-calibrating and non-Cartesian variants.
GRAPPA	Parallel imaging	k-Space synthesis algorithm. Auto-calibrated with ACS. Also non-Cartesian variants.
SPIRiT	Parallel imaging	Self-consistent GRAPPA-like k-space synthesis for arbitrary trajectories. Auto-calibrated with ACS.
UNFOLD	Temporal filtering	Staggered k-space sampling induces aliasing, resolved through temporal frequency filtering.
k-t BLAST/k-t SENSE	Spatio-temporal	SENSE-like unfolding algorithm in high-dimensional x-f domain. Single coil/multi-coil.
k-t GRAPPA	Spatio-temporal	GRAPPA-like k-space synthesis in high-dimensional k-t domain.
k-t SPARSE/k-t SPARSE-SENSE	Compressed sensing/CS and parallel imaging	Sparsity imposed in wavelet domain (spatial) and Fourier domain (temporal). Incoherent artifacts from random sampling. Single coil/multi-coil.
k-t FOCUS/MASTeR	Compressed sensing and motion correction	Sparsity imposed between separate, motion-corrected temporal frames.
XD-GRASP/5D Free-running	Compressed sensing and multiple temporal dimensions	Sparsity imposed along multiple temporal dimensions (e.g. cardiac and respiratory).
k-t SLR/k-t PCA	Temporal low-rank	Low-rank approximation imposed on pixel time series/temporal frequency profiles.
L + S	Compressed sensing and low-rank	Models dynamic image as slowly varying low-rank component and sparse dynamic component.
Image-/k-space- to-image networks	Deep learning	Deep neural network learns correlations between undersampled image/k-space and fully sampled image.
Unrolled networks	Deep learning	Iterative reconstruction alternating between denoising network and data consistency.
SSDU	Deep learning	Splits training data for network input and loss calculation. No fully sampled training data required.
Subject-specific networks	Deep learning	Network trained specifically for every scan.

FOV field-of-view, SENSE sensitivity encoding, GRAPPA generalized autocalibrating partially parallel acquisitions, SPIRiT iterative self-consistent parallel imaging reconstruction, ACS auto-calibration signal, UNFOLD unaliasing by Fourier-encoding the overlaps using the temporal dimension, CS compressed sensing, MASTeR motion-adaptive spatio-temporal regularization, XD-GRASP extra-dimensional iGRASP, 5D five-dimensional, SLR sparsity and low-rank, PCA principal component analysis, SSDU self-supervised via data undersampling, iGRASP iterative golden-angle radial sparse parallel MRI, MRI magnetic resonance imaging, FOCUS focal underdetermined system solver




Fig. 1. k-Space sampling trajectories (top row) and corresponding reconstructions of the Shepp-Logan phantom (bottom row) for (a) a fully sampled Cartesian acquisition, (b) a two-fold undersampled Cartesian acquisition, (c) an undersampled radial acquisition, and (d) an undersampled spiral acquisition. In the k-space diagrams, white circles represent unacquired Cartesian samples and green circles represent acquired samples

be a multiplication by the sampling function $A(k_x, k_y)$, which consists of a Dirac delta function centered at each sample location and has zero value elsewhere. The inverse Fourier transform of $A(k_x, k_y)$ is called a point-spread function (PSF), since it will be convolved with the image and in doing so spreads each pixel of the image.

For k-space sampling on a regular Cartesian grid, the PSF consists of isolated individual points on a regular grid extending infinitely in all image dimensions, with their separation inversely proportional to the spacing between k-space samples. This means each pixel value is replicated an infinite number of times and the overall effect is a tessellation of the image. When the k-space samples $b_{n,m}$ are acquired on a Cartesian grid with spacing consistent with the Nyquist-Shannon sampling theorem [2,3], the separation between replicated pixels is equal to the image field-of-view (FOV), there is no overlap of image tessellations and the discrete inverse Fourier transform allows the image pixel values $\rho_{u,v}$ to be reliably recovered from the k-space data. This is illustrated in Fig. 1a where a Cartesian sampling pattern in k-space returns an example image, the Shepp-Logan phantom, perfectly.

However, when k-space samples are acquired only on a subset of the Cartesian grid points, or on a non-Cartesian trajectory of k-space locations that lie between the grid points, the Nyquist criterion may be breached and image aliasing and the introduction of undersampling artifacts can occur. This can be seen in Fig. 1b-d. Specifically, the Nyquist-Shannon sampling theorem specifies that the grid spacing between k-space samples be

$$\Delta k_i = \frac{2\pi}{\text{FOV}_i}, \quad (1)$$

where i is one of x , y , or z and FOV_i denotes the FOV, or image length, in the i -direction. The size of the k-space that must be filled with samples separated by Δk_i is determined by the desired image resolution; the number of image pixels in the i -direction must be equal to the number of samples in that k_i -direction of k-space. When Cartesian undersampling is applied in one direction with a reduction (or acceleration) factor of R (such that only $1/R$ of the total k-space samples are acquired), the effective FOV in the corresponding image dimension is similarly reduced by a factor of R . This can be seen by rearranging Eq. (1) to make FOV_i the subject and selecting the k-space separation in the

direction of undersampling, $\Delta k_{i,\text{under}}$, to be R times larger than the Nyquist distance $\Delta k_{i,\text{nyq}}$, giving

$$\text{FOV}_{i,\text{under}} = \frac{2\pi}{\Delta k_{i,\text{under}}} = \frac{2\pi}{R\Delta k_{i,\text{nyq}}} = \frac{\text{FOV}_{i,\text{nyq}}}{R}. \quad (2)$$

As an example, in Fig. 1b undersampling by a factor of $R = 2$ has been applied in the phase encoding direction (k_y) such that only every second line of k-space samples is acquired. The resultant image is achieved via an inverse discrete Fourier transform of the undersampled k-space, with unacquired lines zero-filled. Since the FOV in the y -direction has been halved, the tessellations of the image are overlapping and multiple copies of the image in different positions are overlaid, a phenomenon known as image aliasing.

Cartesian undersampling can also be performed in two directions, which would induce horizontal as well as vertical aliasing in Fig. 1b. In two-dimensional (2D) imaging, however, this does not provide practical scan acceleration since each phase-encoding line must be traversed by the gradient coils regardless of how many data samples are taken along it. In 3D imaging, however, 2D Cartesian undersampling may be implemented while the acquisitions remain fully sampled in the frequency-encoding (k_x) direction. When non-Cartesian sampling trajectories, such as radial (Fig. 1c) or spiral (Fig. 1d) trajectories, are used, none of the sampled points lie at the ideal k-space locations determined by the Nyquist criterion. In this case, the extent of the undersampling is different in different directions, and different along different lines in the same direction. For this reason, no coherent tessellation of the image is shifted into our FOV and overlaid on the center image as was the case for Cartesian undersampling. Instead, incoherent artifacts appear across the resultant images, as can be seen in Fig. 1c and d.

The reconstruction of such images can be achieved via first gridding the data (interpolating the acquired samples to the Cartesian grid), and, for well-chosen gridding functions with a sufficient number of non-Cartesian samples, good quality images may be obtained [4,5].

To achieve sufficient acceleration to enable high temporal and spatial resolution in dynamic imaging, however, it is impractical to sample the k-space fully. Hence, a great deal of research and innovation has been directed toward reconstruction techniques that are able to produce high-quality dynamic images from undersampled data.

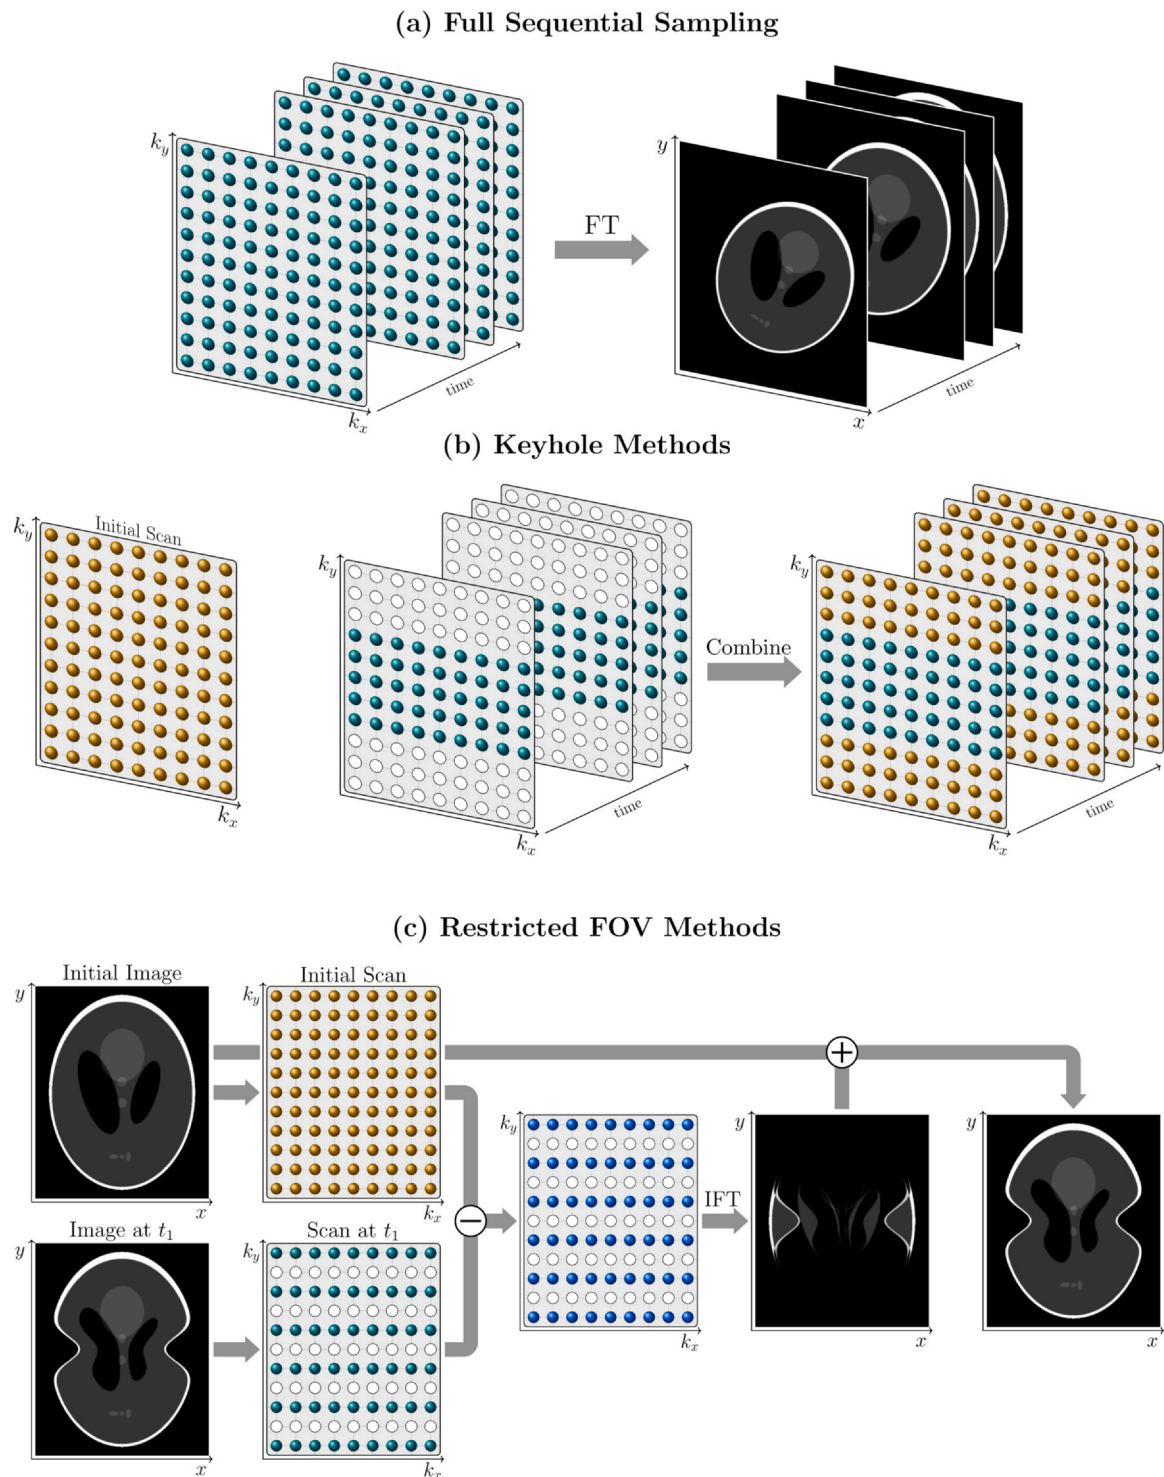


Fig. 2. Early reconstruction methods for dynamic imaging. (a) A fully sampled k-space is acquired at each time point. (b) In keyhole methods, the center of k-space is acquired at each time point and combined with the outer-k-space data from an initial scan. (c) The difference between undersampled data at time t_1 and the same samples from an initial scan gives an unaliased difference image since motion is restricted to the center half of the y-axis. Combining the initial image with the difference image returns an unaliased image for time t_1 . *FT* Fourier transform, *IFT* inverse Fourier transform, *FOV* field of view

3. Early approaches to dynamic image reconstruction

The earliest attempts to implement MRI for dynamic imaging applications aimed to acquire fully sampled 2D k-spaces sequentially [6,7]. Despite using the fastest acquisition sequences available at the time, they were nevertheless severely limited in both spatial and temporal resolution.

Throughout the 1990s, various techniques were proposed to improve the temporal and spatial resolution available in dynamic imaging. In

general, these methods made use of the fact that the multiple time-point images acquired during a dynamic scan are very similar; they differ only in changing contrast and/or physiological motion. Therefore, it is reasonable to suspect that the image series can be reconstructed using less data than would be required for a series of completely unrelated images (i.e. less than consecutive, fully sampled k-space acquisitions as illustrated in Fig. 2a).

Keyhole methods [8,9], for instance, rely on the assumption that most of the dynamic information lies near the center of k-space, thereby

corresponding to large-scale and spatially smooth changes. In these methods, the k-space is fully sampled at the beginning of the scan, and only data in a central “keyhole” region of k-space are reacquired at each time point, as shown in Fig. 2b. In reconstructing the time-point images, the data acquired outside the keyhole region from the initial scan are combined with the updated data within the keyhole.

Partial- and half-Fourier approaches [10–12] similarly use under-sampled k-space acquisitions at each time-point and synthesize the remaining unacquired data. In this instance, undersampling is achieved via a truncation of k-space, rather than an increase in the separation of k-space lines, which, as previously discussed, would correspond to a reduced image FOV. For purely real images (without the complex phase inherent in real coil-weighted MR images), k-space is Hermitian, and thus one half of k-space can be determined as the conjugate transpose of the other. In more realistic imaging scenarios, this real-image assumption does not hold exactly, but does provide the basis of several partial-Fourier methods which tend to acquire just over half of the total k-space. For example, in one approach [11], half of the k-space was acquired at each time point alongside a small number of low-frequency lines in the unacquired half of k-space. The remaining higher-frequency lines were then calculated by conjugating and reflecting the acquired lines, under the Hermitian k-space assumption. The complex phase of the underlying image can also be directly accounted for by utilizing an initial fully sampled scan to calculate magnitude and phase calibration constants which relate the Hermitian-symmetric pairs in the two halves of k-space [12]. By assuming these remain constant throughout the scan, only half the k-space needs to be acquired at each time point, with the calibration constants used to synthesize the other half.

Another class of methods assumes that the motion is restricted to a particular region of the image [13,14], such as the heart in cardiac cine. If the dynamic region occupies only one Rth of the total image FOV, only one Rth of the k-space lines need to be reacquired at each time point to reconstruct the dynamic region. Without further correction, however, time-point images reconstructed from such data would be affected by aliasing of the static image region, since the k-space is now undersampled relative to the overall FOV. Different approaches to alleviating this difficulty were proposed. For example, in one subtraction-based technique [13], the acquired k-space lines at each time point were subtracted from the same lines in a fully sampled initial scan. The inverse Fourier transform of this difference data was taken to provide a reduced-FOV difference image which was correctly positioned within the full FOV using prior knowledge of the dynamic region location. Finally, the difference image was added to the fully sampled static image to yield the full-FOV dynamic image for that time point. This approach is depicted in Fig. 2c. An alternative proposal was to instead construct a k-space data set corresponding only to the static region of the image [14]. This was achieved via the temporal averaging of all k-space lines with pixel values in the dynamic image region set to zero. Subtraction was again used to determine the dynamic and reduced FOV in the dynamic region, with the dynamic and static images combined to yield the full-FOV image at each time point. By averaging over all data to achieve their static image, however, an increase in the signal-to-noise ratio (SNR) was realized.

4. Parallel imaging

A critical development in the history of MRI scan acceleration was the introduction of parallel imaging techniques. In parallel imaging, k-space data are acquired simultaneously in every coil of a multi-coil array [15], as opposed to in a single volume coil, as had been common previously. Since the k-space location of the RF signal received in a coil is determined by the gradient pulse sequence, the same k-space location is acquired in all coils simultaneously, and no increase in scan time is required to acquire N_c copies of the k-space, each corresponding to one of the N_c coils in the array. Together with providing improvements in SNR, an advantage of this approach is that rather than producing

identical (and thereby redundant) copies of k-space, each set of k-space data instead corresponds to the image multiplied by the unique sensitivity profile of the coil that it was acquired with. That is, it consists of discrete samples of the Fourier transform of $S_j(\mathbf{x})\rho(\mathbf{x})$, instead of just the image $\rho(\mathbf{x})$, where $S_j(\mathbf{x})$ is the sensitivity profile of the j th coil. Essentially, each coil only “sees” a limited part of the imaging region close to its position. In this way, parallel imaging arrays provide additional spatial encoding of the k-space data, which can be used to partially replace the Fourier spatial encoding and hence accelerate scan times.

Many algorithms have been proposed to take advantage of the additional information provided by multi-coil arrays [16–26]. Of these, we will consider three of the most popular approaches, SENSE [16], GRAPPA [24], and iterative self-consistent parallel imaging reconstruction (SPIRiT) [25], in detail. These methods have found widespread use in MR image reconstruction. Additionally, they form the basis of many methods specifically designed for dynamic applications, as is discussed later. However, to introduce parallel image reconstruction, we begin instead with the parallel imaging with localized sensitivities (PILS) [17] technique. PILS has strict coil sensitivity requirements which mean it is often not applicable in real-world situations, but it nevertheless provides an intuitive explanation of how multi-coil arrays can provide scan acceleration while avoiding the introduction of aliasing artifacts.

4.1. PILS

In PILS [17], a 2D acquisition is accelerated in the phase-encoding (k_y) direction by a reduction factor of R , such that only every Rth k-space line is acquired and the scan time is similarly reduced by a factor of R . To achieve this acceleration and still obtain an unaliased full-FOV image, the multi-coil array must contain at least R coils ($N_c \geq R$), and each coil sensitivity profile must be localized to a region with a height no greater than FOV_y/R , where FOV_y is the full image FOV in the y -direction. Beyond this region, the coil sensitivity must be zero. Additionally, the position of the center of each sensitive region must be known.

An example of coil sensitivity profiles consistent with this requirement for a four-coil array can be seen in Fig. 3a, and the corresponding k-space sampling pattern for $R = 4$ is illustrated in Fig. 3b. When the inverse Fourier transform is applied to each k-space, without zero-padding, the vertical FOV is also reduced by factor of R in accordance with Eq. (2). As discussed previously, this would normally lead to image aliasing. Now, however, as each coil image is zero outside of the localized region, the height of the non-zero region of the image is equal to or less than the reduced height of the image (FOV_y/R). Therefore, the image wraparound effect induced by k-space undersampling, while still present, does not cause multiple parts of the image to be overlaid at the same position, as can be seen in Fig. 3c. The known y -value can then be used to position each reduced-FOV coil image correctly within the full FOV (Fig. 3d). Finally, the coil images can be combined (using, e.g., a pixel-wise sum-of-squares approach) to yield a full-FOV image covering the entire imaging region of interest free from any aliasing effects (Fig. 3e).

4.2. SENSE

A more flexible framework is provided by the SENSE technique [16], which acts in image space to “unfold” an aliased image. SENSE is compatible with arbitrary coil array configurations, provided that the corresponding coil sensitivity profiles introduce sufficient spatial encoding to the k-space data. This only occurs when there is sufficient variation between the profiles; if the coil sensitivity values are too similar, a SENSE reconstruction may fail.

A successful SENSE reconstruction also requires that coil sensitivity profiles be accurately known. In practice, coil sensitivity maps are

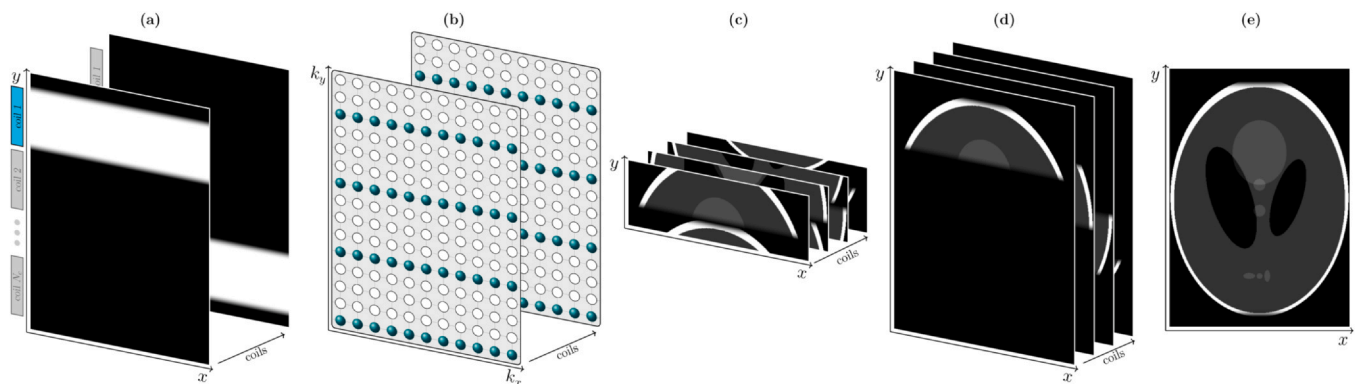


Fig. 3. Stages in a PILS acquisition and reconstruction. (a) Localized coil sensitivity maps consistent with the requirements for a PILS reconstruction. (b) The undersampled k-space acquired with each coil. (c) The inverse Fourier transform of each k-space gives a series of unaliased images with a reduced FOV in the y -direction. (d) The reduced-FOV images shifted to the correct position in the full FOV. (e) The resultant unaliased image formed via a sum-of-squares combination of the coil images in (d). *PILS* parallel imaging with localized sensitivities, *FOV* field-of-view

dependent on the object being imaged and need to be redetermined for every scan. A fast low-resolution calibration scan is often performed prior to the undersampled SENSE, enabling low-resolution images to be reconstructed for each coil in the array. Dividing each individual coil image by an un-weighted image (formed, for instance, using a full body coil, or via sum-of-squares [16,27], or adaptive combination [28] of the individual coil images) effectively removes the underlying object from the images and provides acceptable coil profiles. Potential noise in profiles thus obtained may be reduced through local polynomial fitting [16,27]. Since the profiles tend to be smoothly varying, the low resolution of the calibration scan is sufficient to capture their spatial variation. However, movement of either the patient or the coils between the calibration and the scan can introduce a mismatch between the estimated coil sensitivities and those corresponding to the undersampled data [29], and lead to artifacts in the reconstructed image. Auto-calibrated techniques [29,30] such as modified SENSE [31] have been proposed to mitigate this issue. These techniques utilize additional k-space lines acquired with full sampling at the center of k-space to estimate the sensitivity profiles. Since the calibration data are acquired simultaneously to the undersampled SENSE data, motion occurring between the acquisitions is eliminated.

As depicted in Fig. 4, in SENSE, the inverse Fourier transform of each undersampled coil-weighted k-space is taken to yield a set of N_c images, all of which are affected by aliasing. In the example shown, undersampling with a reduction factor of $R = 2$ has been implemented

in the phase encoding (k_y) direction only. In general, however, SENSE is applicable to Cartesian undersampling in multiple directions and higher reduction factors, provided the overall reduction factor is less than or equal to N_c , the total number of coils in the phase encoding direction. Each pixel in each of the aliased images is formed via the superposition of multiple sensitivity-weighted pixels from the full-FOV coil-weighted images. This is highlighted in Fig. 4; the red boxes in Fig. 4a indicate the positions of two pixels in the full-FOV images, ρ_1 and ρ_2 , which alias to the same position in the reduced-FOV images, labeled a and indicated by the red boxes in Fig. 4c.

The relationships between unaliased pixel values and aliased coil-weighted image pixels can be expressed in sets of linear equations, as detailed in Appendix A.2. Sets of linear equations such as these can be solved provided the number of equations is greater than or equal to the number of unknowns. Here, there is one equation for each coil in the array, and each set of equations contains all the pixels R that are aliased to the same position in the reduced-FOV image (the unknowns), hence the requirement for the number of coils to be equal or greater than the acceleration factor.

When each of the sets of equations is solved, the calculated pixel values ρ_n can be positioned appropriately to recover a full-FOV image without aliasing, as illustrated in Fig. 4d.

Generally, however, and particularly for non-uniform sampling patterns, a large system of equations will be obtained, and thus a more computationally efficient iterative scheme is often used to reconstruct images in SENSE [19]. The problem becomes to iteratively solve

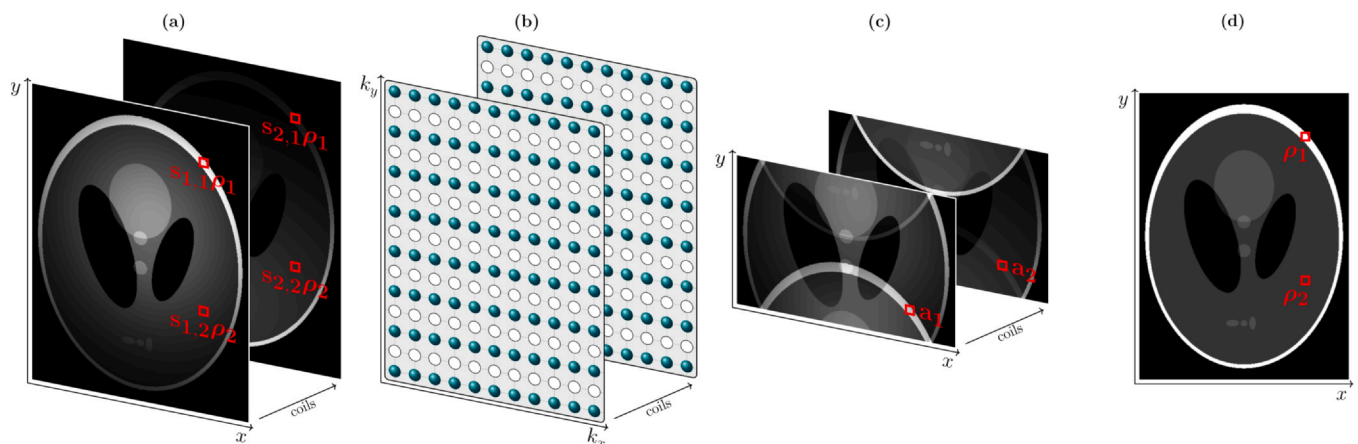


Fig. 4. Stages in a SENSE reconstruction. (a) Coil-weighted images, as “seen” by the coils in the array. (b) An undersampled k-space is acquired with each coil. (c) Taking the inverse Fourier transform of each k-space gives a series of aliased images. (d) SENSE reconstruction “unfolds” the aliased image to give the full-FOV image. *SENSE* sensitivity encoding, *FOV* field-of-view

$$\mathbf{b} = F S \rho \quad (3)$$

for ρ , where ρ is the vector of image pixels, \mathbf{b} is a vector of the data acquired in all coils, S is a matrix of coil sensitivities, and F is a Fourier transform operator. More generally, F and S can be replaced by an encoding operator E which contains any operations required to transform the solution image to the acquired data. Within this framework, non-Cartesian trajectories, such as radial and spiral, can easily be incorporated since the implementation of E can utilize either a non-uniform FFT algorithm or k-space gridding [4] between the Cartesian and acquired grids at each iteration.

Noise amplification, and a corresponding reduction in SNR, is an inevitable result of SENSE and other parallel-imaging-based acceleration methods. The SNR of an undersampled SENSE reconstruction is related to the SNR of a fully sampled acquisition via the equation

$$SNR_{\text{under}}(x, y) = \frac{SNR_{\text{fully}}(x, y)}{g(x, y) \sqrt{R}}. \quad (4)$$

Here g is the g-factor (geometry factor), which is often used to characterize the spatial distribution of the noise amplification and depends on the geometry of the coil array [16,32]. Eq. (4) reveals that the loss in SNR is also dependent on the square root of the acceleration factor R , since with higher acceleration fewer samples are available from which to reconstruct the image.

The SENSE equations can become ill-conditioned when there are areas of poor k-space coverage (such as the corner of a rectangular k-space FOV when a radial or spiral trajectory is used) or if there is insufficient SENSE due to insufficiently distinct coil sensitivity profiles [19,33]. Mathematically, ill-conditioned problems are characterized by the property that small changes to their input (the k-space data, in the present case) may result in large changes to their output (the reconstructed images). In such scenarios, SENSE can become sensitive to, and amplify, noise in the acquired data. Various regularization techniques can be employed to restrict this noise amplification and improve image quality. For instance, k-space can be filtered to reduce the contribution of values found in regions with poor coverage [19] or a singular value decomposition (SVD)-truncation can be applied to the encoding operator E [34]. Alternatively, the iterative formulation allows regularization functions to be added to $\|E\rho - \mathbf{b}\|_2^2$, the SENSE objective, and minimized during the iteration process. A simple Tikhonov objective function, calculated as the norm of the image pixel values, has been demonstrated to reduce image noise [33,35], while more complicated penalty functions are often utilized within a similar framework in CS techniques, as discussed further in Section 6.

While primarily introduced here as background for the more-advanced dynamic reconstruction methods discussed in later sections, SENSE itself has been applied on a frame-by-frame basis to accelerate real-time cardiac imaging [36,37]. Such applications do not take advantage of the temporal redundancies intrinsic to dynamic MRI datasets, and as such a trade-off between spatial and temporal resolution remains a limiting factor.

4.3. GRAPPA

Distinct from SENSE and other methods that operate directly in image space to “unfold” aliased images, are the class of parallel imaging algorithms that attempt to recover unacquired data in k-space [24–26]. Among these, GRAPPA [24] is the most common. Unlike SENSE, GRAPPA is auto-calibrating, meaning that accurate coil sensitivity profiles are not required as an input to the method. This may be particularly advantageous in imaging scenarios where movement occurs between the sensitivity calibration scan and the undersampled scan, although, as mentioned auto-calibrated extensions of SENSE have also been developed. In GRAPPA, the requirement for sensitivity profiles is replaced by the need for an auto-calibration signal (ACS), which generally consists of a fully sampled set of k-space lines around $k_y = 0$

(assuming undersampling in the k_y direction). The idea underlying GRAPPA is that multi-coil k-space data are correlated, since each coil’s k-space corresponds to the same physiological image modulated by a generally smooth coil sensitivity profile. Specifically, GRAPPA proposes that each unacquired k-space sample may be reliably synthesized by a linear combination of the acquired k-space samples in its multi-coil k-space neighborhood, and, furthermore, that the weights used in this linear combination are constant throughout k-space.

The first assumption builds on ideas introduced in the simultaneous acquisition of spatial harmonics (SMASH) [21] method, and its subsequent auto-calibrated extensions [22,23]. SMASH utilizes linear combinations of a single k-space line acquired in different coils to synthesize unacquired neighboring lines. The linear combination weights are chosen such that the combined coil sensitivity profile closely matches a spatial harmonic which modulates the exponential term in the integrand of the Fourier transform (Eq. (A.1)) so that it now corresponds to an adjacent k-space position, as described in Appendix A.3.

In GRAPPA, this idea is extended and made more robust by including k-space samples from multiple locations surrounding the target location, and by synthesizing the missing samples in every coil rather than just in a k-space corresponding to the final image unmodulated by coil sensitivities.

Fig. 5a and b depicts the k-space geometry of a standard GRAPPA reconstruction with two-fold undersampling in the k_y direction. Near the center of k-space, shown in Fig. 5a, the k-space is fully sampled due to the acquisition of ACS lines, whereas in other regions, depicted in Fig. 5b, only every second phase-encoding line is acquired. In these figures, the shaded region represents the multi-coil k-space neighborhood of a k-space sample to be synthesized. As indicated by the arrows in Fig. 5b, all the acquired samples that sit within this neighborhood contribute to the synthesis of the unacquired sample via a linear combination which utilizes a calibrated kernel of weights. A mathematical description of the GRAPPA synthesis is included in Appendix A.4.

To implement GRAPPA and synthesize the missing samples in every coil, the kernel of weights, g , must first be calibrated using the ACS lines. As indicated by the arrows in Fig. 5a, the calibration process utilizes the same multi-coil neighborhood geometry as the synthesis step, but here both the target point being synthesized and the source points contributing to the synthesis are known, since the ACS region is fully sampled. By equating the known value of each target point with the value that would be produced were it to be synthesized from the neighboring source points with GRAPPA, an equation is constructed for every possible kernel position within the ACS region. The resultant system of equations is solved to find the weights, often in a least-squares manner since the number of equations (determined by the number of viable kernel positions) ideally exceeds the number of weights, increasing robustness to noisy data.

Once the kernel of weights has been calculated, the unacquired samples are synthesized across every coil, creating a fully sampled multi-coil k-space, as illustrated in Fig. 5c. From this, coil images can be calculated using the standard inverse Fourier transform operation and combined using, for example, a sum-of-squares technique.

In certain cases, it can be advantageous to undersample in more than one direction. For example, for a 3D acquisition, undersampling can be implemented in both the phase-encoding (k_y) and slice (k_z) directions. GRAPPA is readily applicable to multi-dimensional undersampling. [25] However, as can be seen in Fig. 5d, which depicts the GRAPPA kernels for a 2D undersampling scheme with two-fold undersampling in each of the k_x - and k_y -directions, different neighborhood acquisition patterns arise for different samples being synthesized. In these cases, a separate set of weights must be calibrated and applied for each pattern.

Variants of GRAPPA have also been developed for non-Cartesian trajectories [38–44], including, for instance, radial [38–41], spiral [42,43], and Periodically Rotated Overlapping Parallel Lines

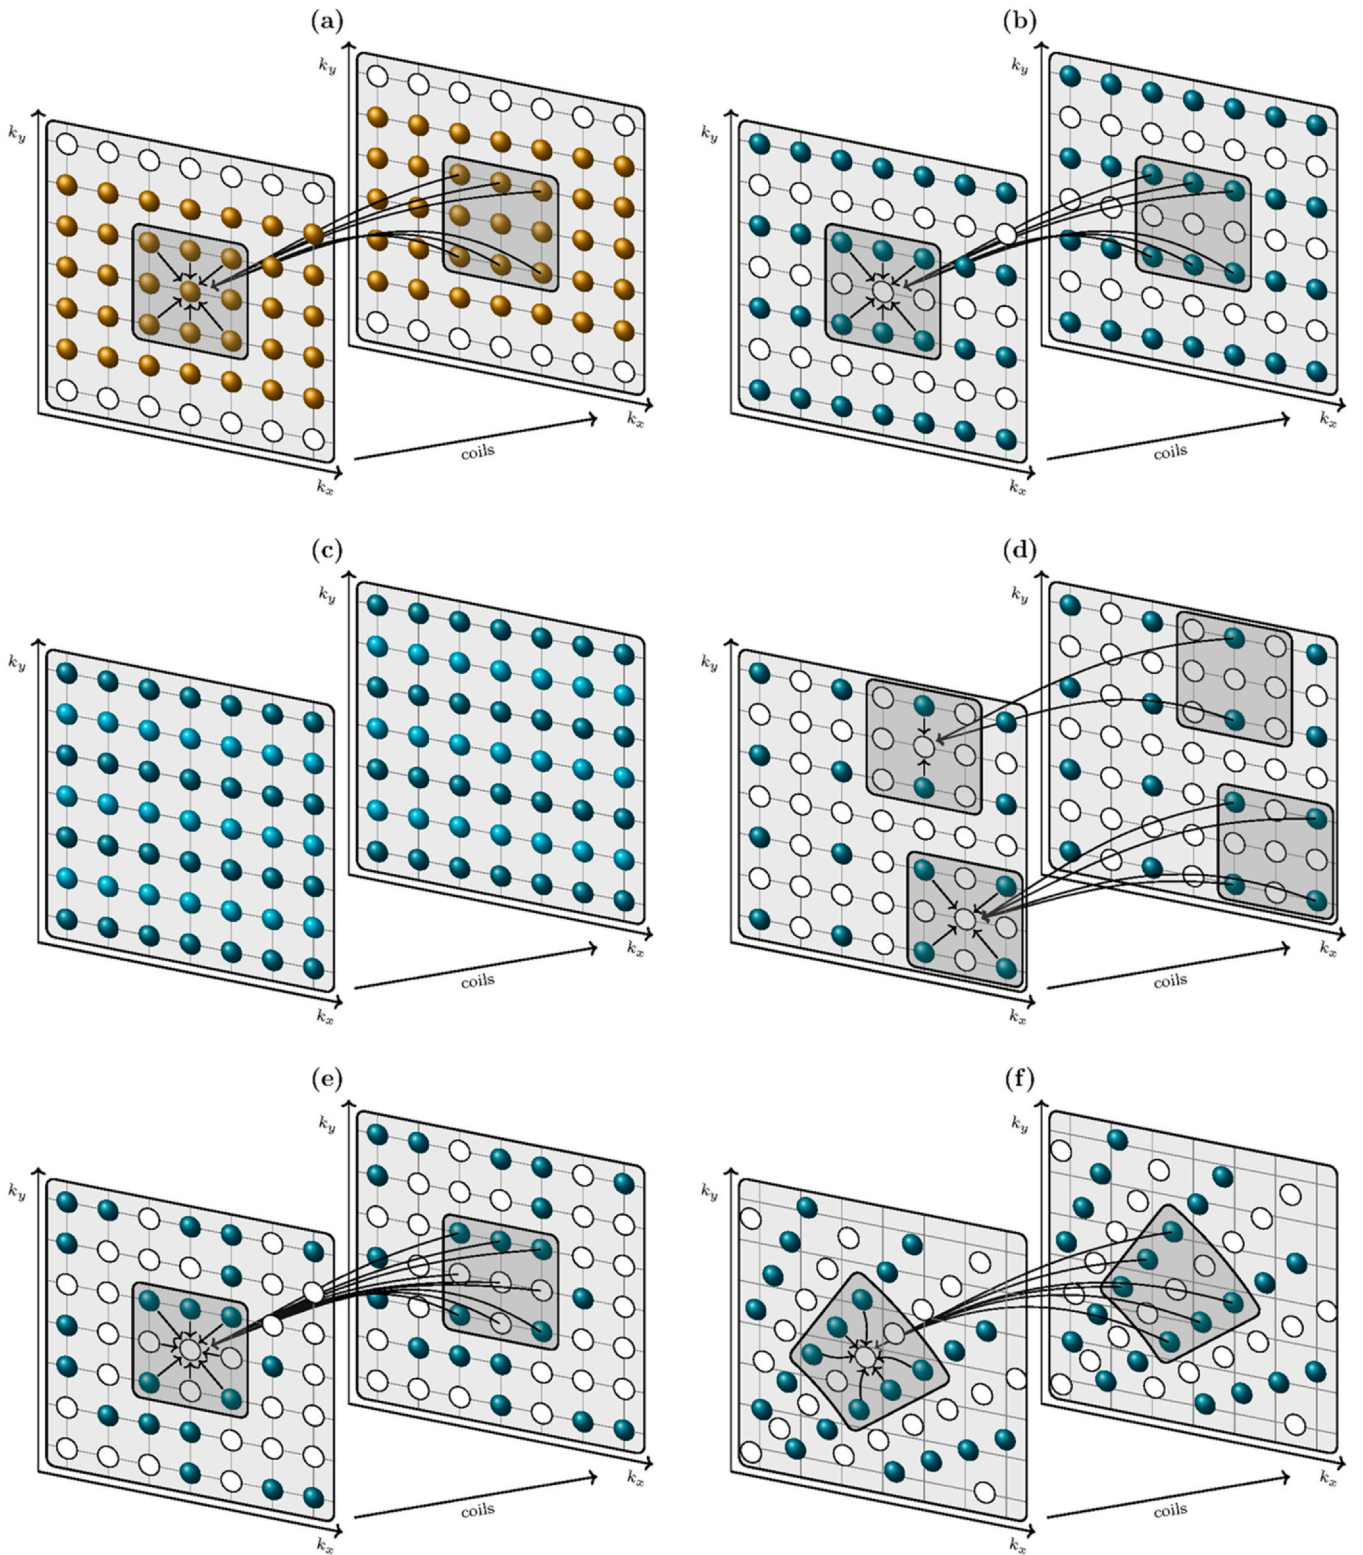


Fig. 5. Kernel calibration and k-space synthesis for GRAPPA and SPIRiT reconstructions. (a) Kernel calibration in the ACS region at the center of k-space. (b) k-Space synthesis with GRAPPA for a two-fold undersampling in the k_y direction. (c) Fully sampled multi-coil k-space with unacquired samples synthesized using GRAPPA. (d) GRAPPA applied to 2D undersampling where different neighborhood acquisition patterns exist for different positions in k-space. (e) A SPIRiT kernel for an arbitrary sampling pattern with both acquired and unacquired samples contributing to the synthesis of an unacquired sample. (f) GRAPPA applied to a radial sampling trajectory where the neighborhood takes the shape of a truncated wedge. *GRAPPA* generalized autocalibrating partially parallel acquisitions, *SPIRiT* iterative self-consistent parallel imaging reconstruction, ACS auto-calibration signal, 5D five-dimensional

with Enhanced Reconstruction (PROPELLER) [44] acquisitions. These applications pose the additional complication that the length and direction of the displacement between neighboring samples vary with position. For example, in a radial trajectory, the rectangular kernel geometry from Cartesian GRAPPA is replaced with a truncated wedge, as is illustrated in Fig. 5f. The displacement between the target point at the center of this kernel and a neighboring source point will increase as the radius of the kernel center increases, while the displacement vector will be rotated as the angular position of the kernel changes. In such non-Cartesian circumstances, Eqs. (A.10)–(A.12) do not apply and the kernel of weights cannot be applied throughout k-space. Initial implementations of radial GRAPPA [38,39] required a fully sampled calibration dataset and then utilized segmentations in both the radial and azimuthal directions. Within each segment, the differences were assumed to be small, allowing a standard calibration where the radial data are treated as Cartesian, and separate weights were calibrated and used in each segment. More accuracy can be obtained by defining a separate segment for each kernel position, with that kernel at its center [40]. A trade-off exists between the benefits of smaller calibration segments (which more closely satisfy the invariant displacement vector assumption and reduce blurring) and the increased number of viable kernel positions allowed by larger segments, which tend to reduce noise. Additional kernel instances can be achieved for smaller calibration segments in dynamic imaging applications by obtaining multiple fully sampled radial k-spaces before the undersampled scan [40,43]. Similarly, calibration data from multiple slices can be used together in a 3D stack-of-stars acquisition [41].

Self-calibration schemes for non-Cartesian GRAPPA have also been proposed [45–49]. For example, the radial self-calibrated approach [45] calculates weights using larger kernels on the undersampled radial acquisitions (with correspondingly increased displacements between target and source points). The required weights for smaller kernels positioned on the full radial grid are then interpolated from the weights calculated at various radii but the same angular position. A conceptually similar approach involves searching the acquired radial k-space for similar sampling patterns to each needed for the synthesis of an unacquired sample [46]. Self-calibration for arbitrary non-Cartesian trajectories has been realized through the use of a Cartesian calibration region (gridded from a highly sampled region of the k-space) [47,48]. This can be then used to grid or interpolate multiple instances of each neighborhood sampling pattern that appear in the undersampled trajectory for calibration.

As was the case for SENSE, GRAPPA can be applied frame-by-frame to dynamic cardiac data [50] without exploiting through-time correlations. However, greater efficiency can be realized from GRAPPA-based reconstructions of dynamic data by constructing calibration regions from neighboring temporal frames (as mentioned above for radial GRAPPA) or by applying GRAPPA-like concepts to a high-dimensional spatio-temporal domain. These approaches are discussed further in Section 5.

4.4. SPIRiT

The SPIRiT technique further generalizes GRAPPA by avoiding the requirement of different calibration kernels for every acquisition pattern, and thereby permits the use of arbitrary Cartesian k-space undersampling. This is achieved by constructing a single kernel with weights corresponding to every sample in the multi-coil neighborhood, regardless of whether or not those samples have been acquired. Fig. 5e illustrates the concept; with an arbitrary Cartesian undersampling pattern the target point at the center of the kernel in the first coil has both acquired and unacquired samples in its multi-coil neighborhood, and the target point is synthesized (as indicated by the arrows) with a linear combination of all these samples. However, since the unacquired samples are unknown, and must also be synthesized using the same procedure, the synthesis step cannot be immediately

implemented following kernel calibration as it was with GRAPPA. Instead, an iterative solution is found to satisfy

$$\arg \min_{\mathbf{b}} \|D\mathbf{b} - \mathbf{b}_{\text{acq}}\|_2^2 + \lambda \|(G - I)\mathbf{b}\|_2^2, \quad (5)$$

where \mathbf{b} is the vector of synthesized fully sampled multi-coil k-space data to solve for, \mathbf{b}_{acq} is a vector of the acquired data, D is an under-sampling operator which selects only the acquired points from the fully sampled k-space, λ is a penalty weighting, I is the identity, and G is a matrix constructed from the kernel elements which applies the kernel convolution to each coil's k-space. Eq. (5) consists of two constraints; the first term requires that the synthesized data matches the acquired data at those locations where data are acquired, while the second term requires the solution to be consistent with the kernel, such that applying the linear-combination synthesis to the solution (where unlike Fig. 5e all neighboring samples now have values) returns the same multi-coil fully sampled k-space.

The SPIRiT formalism can also be adapted for non-Cartesian acquisitions by simply replacing the undersampling operator D with an operator that interpolates the fully sampled data to the off-grid locations of the acquired non-Cartesian samples.

5. Exploiting temporal redundancy

During the period when pioneering techniques in parallel MRI were being introduced, so too were a series of reconstruction methods designed specifically for dynamic imaging applications. Just as correlations in multi-coil k-space data were exploited to accelerate scans with parallel MRI, these methods exploit the correlations present in the temporal dimension of dynamic scans.

As such, these methods often consider a higher dimensional x-t domain, where t represents the temporal dimension of the scan and x represents the spatial dimensions (generally two or three). The goal of image reconstruction becomes to fully reconstruct the image in x-t space such that a full FOV image without artifacts exists for every discrete value of t , as determined by the desired frame-rate of the scan. From this concept follows other useful parameter domains such as k-t space (the k-space at each time point), and x-f and k-f space, where f is the temporal frequency. The latter two can be obtained from x-t and k-t space, respectively, via a one-dimensional (1D) Fourier transform applied in the temporal dimension.

Previous works had explored the theory underlying spatio-temporal sampling [51–54], but the first prominent application of these ideas to MRI was the unaliasing by Fourier-encoding the overlaps using the temporal dimension (UNFOLD) [55] method. UNFOLD allows undersampling by separating aliased components in the temporal frequency dimension and can also be combined with parallel imaging [56–58]. Building on UNFOLD, k-t broad-use linear acquisition speed-up technique (BLAST) and k-t SENSE [59] exploit high dimensional aliasing in x-f space to achieve higher acceleration or reduce remaining artifacts from parallel imaging. In this section, we discuss the UNFOLD and k-t SENSE approaches, and also consider the temporal GRAPPA (TGRAPPA) [60] and k-t GRAPPA [61] methods which extend GRAPPA for dynamic applications.

5.1. UNFOLD

Like SENSE, the UNFOLD [55] method acts in image space to unfold aliased images. The manner in which it does so is also similar to SENSE; both UNFOLD and SENSE are able to separate individual pixel contributions to an aliased pixel since multiple aliased images are formed via linear combinations of the unaliased pixels with different combination weights. In SENSE, each aliased image corresponds to a single coil in the array, and the aliased pixel values are modulated by the coil sensitivities. In UNFOLD, each aliased image instead corresponds to a time point in a dynamic acquisition, and the aliased images correspond to different undersampling trajectories which are used at each time point.

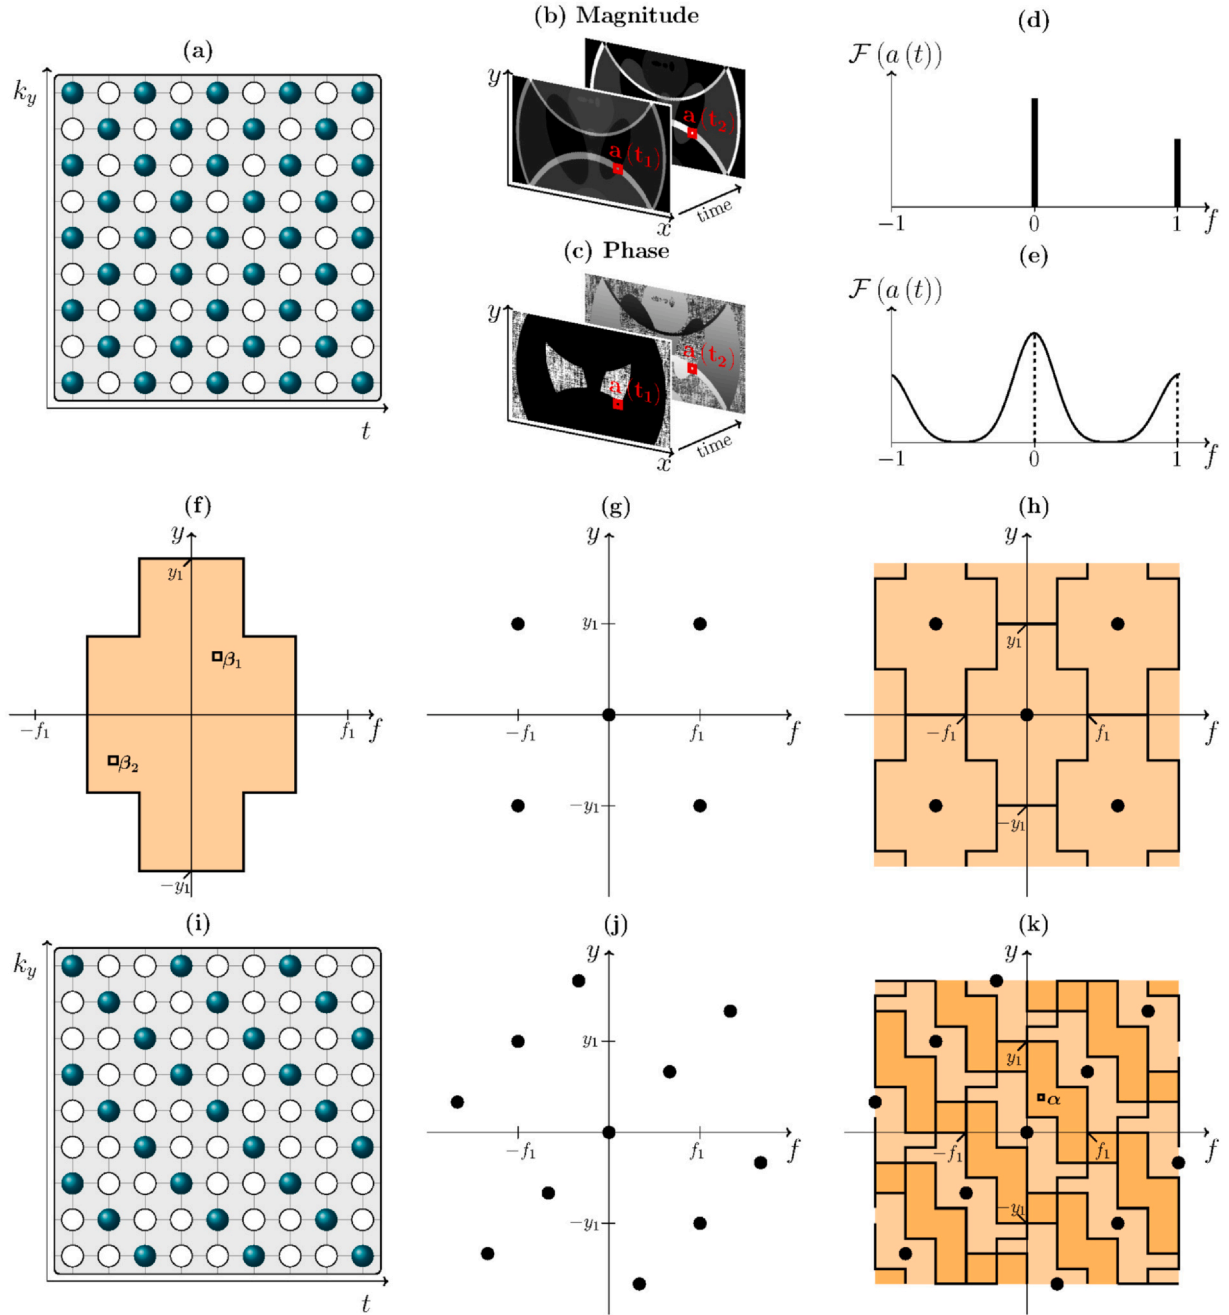


Fig. 6. Diagrams depicting k-t acquisition patterns, x-f support regions, and x-f aliasing for methods exploiting temporal redundancies. (a) A staggered two-fold-undersampled k-t acquisition pattern for UNFOLD. Each line is fully sampled in the k_x -direction (not pictured). (b) and (c) Magnitude and phase of aliased images produced at odd and even time points from the trajectory in (a). (d) and (e) The temporal frequency spectrum for an arbitrary pixel acquired with the trajectory in (a) when motion is (e) and is not (d) present. (f) An x-f support region in UNFOLD, assuming the center half of the y-axis requires twice the temporal frequency information to describe the motion as than outer half does. (g) The PSF in x-f space associated with the k-t acquisition pattern in (a). (h) Tessellation of the x-f support region shown in (f) due to the k-t acquisition pattern in (a). (i) A three-fold-undersampled k-t acquisition pattern for k-t SENSE. (j) The PSF in x-f space associated with the k-t acquisition pattern in (i). (k) Aliasing of the support region shown in (f) due to the k-t acquisition pattern in (i). *UNFOLD* unaliasing by Fourier-encoding the overlaps using the temporal dimension

As an example, consider the two-fold undersampled 2D Cartesian acquisition pattern utilized in the original presentation of UNFOLD [55]. At each time point every second phase-encoding line is acquired, but the acquisition alternates between odd and even lines. This staggered sampling pattern is depicted in Fig. 6a, where the frequency encoding (k_x) direction (not shown for simplicity) is assumed to be fully sampled. As we have seen previously, two-fold undersampling causes image aliasing such that two pixels are superimposed at every location in the reduced FOV image at every time point. However, in

accordance with the Fourier shift theorem, the aliased image from the even k-space lines will not be identical to that from the odd k-space lines. Since the sampling function $A(k_x, k_y)$ has been shifted in the k_y -direction for every second time point, a phase shift is applied to the corresponding PSF. A given pixel in the aliased images thus contains contributions from the same pixels from the full FOV image, but these have been modulated by different phases, as can be seen in Fig. 6b and c where the magnitude and phase of the images at each time point are shown separately.

If the image is static, as is the case in Fig. 6b and c, then each aliased pixel value alternates between the sum and the difference of its two constituent pixels (as the phase modulation experienced by the neighboring image tessellation alternates between $e^{i\pi 0} = 1$ and $e^{i\pi} = -1$):

$$a(t_{\text{even}}) = \rho_1 + \rho_2 \quad a(t_{\text{odd}}) = \rho_1 - \rho_2. \quad (6)$$

Therefore, taking the Fourier transform in the temporal dimension for a given aliased pixel yields exactly two non-zero components: the direct current (DC) component (attributable to ρ_1) and the Nyquist frequency (attributable to ρ_2). This is shown in Fig. 6d and permits a more flexible view of UNFOLD which can accommodate motion; rather than simply unfolding aliased images, UNFOLD separates aliased information in the temporal frequency direction.

When, as is the case in a dynamic imaging scenario, the constituent pixels contain motion, the discrete frequency spikes will be replaced with broader spectra, as shown in Fig. 6e. The bandwidth of these curves is determined by the motion present in the pixel itself; the more rapid and aperiodic the motion, the larger the range of temporal frequencies required to represent it accurately. When the temporal bandwidths are sufficiently narrow such that they do not overlap, as is the case in Fig. 6e, UNFOLD can be used to separate the aliased pixels by filtering in x-f space. Within this framework, it can be seen that UNFOLD can be applied to arbitrary undersampling schemes, provided the sampling pattern is altered each frame in a periodic manner, there is sufficient bandwidth to permit the same number of non-overlapping pixel spectra as there are frames in one period and that the accumulated sampling pattern over a full period is sufficient to provide an image free from aliasing and artifacts.

Additionally, regions of the image where more motion is expected, such as the heart in cardiac imaging, are likely to require wider bandwidths, but these can be stacked in x-f space alongside more-static regions with narrower bandwidths. A simple framework for visualizing the potential acceleration is provided by examining the support regions in x-f space [62]. Fig. 6f depicts a cross-shaped x-f support region, consistent with an imaging scenario where a dynamic region (requiring a larger range of temporal frequencies) is present in the center half of the y-axis. Now, since x-f space is formed via the Fourier transformation of the k-t domain in every dimension, the PSF of the staggered acquisition pattern in Fig. 6a exists in x-f space and is depicted in Fig. 6g. As shown in Fig. 6h, the replications of the cross-shaped support region induced by this PSF tessellate without overlap.

UNFOLD can therefore be implemented as a filtering operation in the x-f domain. Exactly one copy of the support region is retained, and everywhere else is set to zero.

UNFOLD may also be combined with parallel imaging algorithms such as SENSE [56–58] and GRAPPA [57], where it can be used to either increase the acceleration factor, remove residual artifacts remaining after reconstruction, or a combination of the two. Temporal filtering SENSE (TSENSE) [56] and UNFOLD-SENSE [58] are two methods that combine UNFOLD with SENSE. While differing in the details of their implementation, a feature of both is that through UNFOLD-like temporal filtering they allow the estimation of time-varying sensitivity maps which can be utilized in the SENSE reconstruction. Such adaptive coil maps preclude the need for an initial reference scan and can accommodate relatively slow variations (e.g. respiratory motion).

5.2. k-t SENSE and k-t BLAST

The k-t acquisition scheme used in UNFOLD avoids aliasing by ensuring support region replicates do not overlap in the x-f domain, and then applies a filter in that domain to extract just one copy. Effectively, SENSE-like unaliasing concepts are used to unalias x-t images by separating them in the temporal frequency dimension f. Building on these ideas, the k-t BLAST (single coil) and k-t SENSE (multi-coil array) methods go further; by permitting higher acceleration factors that do

lead to aliasing in the x-f domain, k-t SENSE reconstructs the unaliased time series of images by applying a SENSE-like unfolding algorithm in the higher dimensional space.

In k-t SENSE, the staggered acquisition pattern of UNFOLD is utilized with higher acceleration factors. This is depicted in Fig. 6i where an acceleration of $R = 3$ each frame is applied, leading to closer peaks in the PSF (Fig. 6j). In the x-f domain, this leads to the overlapping of the same support region which produced no x-f aliasing for two-fold undersampling (Fig. 6k). The arbitrary point α , indicated in the aliased x-f domain, is seen to be formed from the overlap of the points labeled β_1 and β_2 in Fig. 6f.

In general, an arbitrary number of points may alias to the same position in x-f space, and a set of linear equations describing these points can be constructed (see Appendix A.5). These share the same structure as the SENSE equations and can be solved in a similar way. For k-t BLAST, and k-t SENSE when the number of coils is less than the number of x-f points aliasing to the same position, the equations are underdetermined and no unique solution exists. For this reason, an initial low-spatial-resolution training scan is used to estimate a matrix M^2 containing the expected magnitudes of each β . Alternatively, the training scan can be avoided by utilizing an auxiliary TSENSE [56] reconstruction to generate the training data [63].

A solution is found that ensures each β matches its expected signal magnitude, with the vector of values β found by solving

$$\beta = \underline{\beta} + (S^* \Psi^{-1} S + M^{-2})^{-1} S^* \Psi^{-1} (\beta_{\text{alias}} - S \underline{\beta}), \quad (7)$$

where S is a matrix of coil sensitivities, Ψ is a matrix containing the covariance of noise levels across the coils, $\underline{\beta}$ is a baseline estimate of β , and $*$ denotes the conjugate transpose. When k-t SENSE was first proposed, the baseline estimate was obtained using temporal averaging. However, this results in a number of nulled temporal frequencies, since the DC component (equivalent to the temporal average) is aliased to $R-1$ higher frequencies, and thus these frequencies are also removed when the baseline subtraction is performed [64]. Alternative baselines formed using TGRAPPA [64] (discussed below) or reduced-encoding imaging by generalized-series reconstruction (RIGR) [65,66] techniques have been proposed; these approaches are able to recover the nulled frequency information and result in reduced artifact power. When arbitrary non-Cartesian sampling schemes are used or the time-invariant coil sensitivity assumption cannot be met, the equation becomes

$$\beta = \underline{\beta} + (E^* \Psi_{k,t}^{-1} E + M^{-2})^{-1} E^* \Psi_{k,t}^{-1} (\mathbf{b}_{k,t,\gamma} - E \underline{\beta}), \quad (8)$$

where E is a forwards operator similar to that which models the acquisition in iterative SENSE (Eq. (3)). Here, E explicitly incorporates the sampling pattern and applies coil-sensitivity weights to every time-point, $\Psi_{k,t}$ is the noise covariance matrix in k-t space and $\mathbf{b}_{k,t,\gamma}$ is the vector of acquired k-t samples in the γ th coil.

5.3. TGRAPPA and k-t GRAPPA

TGRAPPA and k-t GRAPPA have both been proposed as extensions to GRAPPA for dynamic imaging applications. In TGRAPPA [60], a staggered sampling pattern, as used in UNFOLD and k-t SENSE, is implemented to avoid the requirement of acquiring ACS lines; the lines near the center of k-space across neighboring temporal frames are collected and used together as the GRAPPA calibration region. Once the kernel is calibrated, however, the GRAPPA k-space synthesis proceeds normally and separately for each time point.

The k-t GRAPPA method [61] explicitly incorporates the higher dimensional k-t domain by applying the GRAPPA kernel in the k-t domain. An unacquired line is synthesized via a linear combination of acquired lines which include both spatial neighbors, acquired at the same time point as the line being synthesized, and temporal neighbors, acquired at the same k-space location before or after the time the line

being synthesized corresponds to. To calibrate the k-t GRAPPA kernel, an ACS region in k-t space, similar to that seen for regular GRAPPA in Fig. 5a, is obtained by fully sampling the center of k-space at each time point.

6. Compressed sensing and low-rank methods

CS was first proposed for MRI reconstruction in the mid-2000s [67–69] and has become a popular technique for many applications. The success of CS lies in the fact that medical images do not consist of random independent pixels stacked side-by-side. Rather, they exhibit clear spatial correlations and, in particular, have sparse representations upon transformation to certain domains (meaning they can be represented by a number of elements in these domains that is far fewer than the number of pixels or voxels in the image). This fact can be used as prior information to regularize the ill-conditioned inverse problem that is undersampled MR reconstruction, provided the acquisition scheme would generate incoherent aliasing artifacts (as seen for radial and spiral sampling in Fig. 1c and d, but not for uniform Cartesian undersampling as seen in Fig. 1b). The general CS MRI reconstruction problem can be expressed as

$$\arg \min_{\rho} \|E\rho - \mathbf{b}\|_2^2 + \lambda \|\Psi\rho\|_1, \quad (9)$$

where Ψ is the sparsifying transform, λ is a penalty weighting, and the first term, which involves the encoding operator E , the acquired data \mathbf{b} and the image ρ , forces the sparse solution to be consistent with the acquired data. The L_1 -norm is applied to the sparse domain since its minimization promotes sparsity. If desired, multiple sparsity terms can be included in Eq. (9) with different penalty weightings.

Eq. (9) is solved iteratively, using, for example, a conjugate-gradient descent algorithm or projection onto convex sets, which can often lead to long reconstruction times making real-time applications of CS challenging.

Total variation (TV) and wavelet transformations are two examples of transforms often used in CS MRI reconstruction. These are depicted in Fig. 7, where the sparse representations of a reconstructed cardiac image in each of the domains are apparent. The TV is a first-order approximation of the spatial derivative; regions of relatively constant pixel value produce small values whereas tissue boundaries result in larger values. Minimizing the number of non-zero TV coefficients thus promotes large smooth regions in the image. The wavelet domain contains information about the presence of structure in the image at different scales, and is well known to provide a sparse representation of natural images, as is exemplified, for instance, in its use as the basis for JPEG 2000 image compression.

The k-t SPARSE [70] method, proposed for cardiac MRI, was the first application of CS for dynamic imaging. It utilizes random-order phase encoding in the k-t domain to ensure incoherent artifacts and incorporates sparsifying transforms in both the spatial and temporal dimensions. A wavelet transform, as previously discussed, was applied in the spatial dimensions. In the temporal dimension, a Fourier transform was applied, since it sparsifies smooth and periodic motion. In this manner, similarities between CS in the temporal dimension and UNFOLD, which applied filtering to the temporal frequency f , become apparent. Indeed, CS has also been applied to cardiac cine imaging with the x-f domain used directly as the sparse domain [71], and in that study it was shown to outperform k-t BLAST. Increased artifact incoherence can be achieved by replacing the random order but Cartesian phase encodes of k-t SPARSE with radial sampling. The iterative golden-angle radial sparse parallel MRI (iGRASP) technique [72], for instance, acquires radial spokes with golden angle (111.25°) separation between successive spokes, ensuring relatively uniform k-space coverage for any arbitrary series of consecutive acquisitions.

The CS framework is also readily compatible with parallel imaging techniques; when \mathbf{b} consists of multi-coil data and the encoding operator E incorporates coil sensitivity profiles, Eq. (9) is simply a regularized form of the iterative SENSE equation (Eq. (3)). As an example, the k-t SPARSE-SENSE approach, which utilized the sparsifying transforms from k-t SPARSE and a multi-coil array, was able to achieve eight-fold acceleration ($R = 8$) with comparable image quality to a two-fold accelerated GRAPPA reconstruction for first-pass cardiac perfusion [73]. The same approach applied to real-time cine, again with an acceleration factor of $R = 8$, was able to produce diagnostic images with a temporal resolution of 40–50 ms [74,75].

As the potential of CS reconstruction for improving dynamic MRI became clear, more advanced methods incorporating motion estimation (ME) and motion correction (MC) were proposed [76–82]. One such method, k-t focal underdetermined system solver [76,77], imposes sparsity not on a transform of the image itself but on the residual between the image and an initial image prediction. The prediction may be obtained using a RIGR-like method, or a more sophisticated ME/MC technique which estimates motion vector field between frames and applies these fields to fully sampled reference images.

The motion-adaptive spatio-temporal regularization [78] method similarly utilizes inter-frame motion, but without relying on a reference frame. Instead, forwards and backwards motion fields are estimated between every pair of adjacent frames from an initial reconstruction, and the sparsity constraint is applied to the difference between an estimated frame and its neighboring frame warped by the motion field to be in the same motion state. Motion can alternatively be estimated across groups of frames, reducing the impact artifacts in the initial reconstruction can have on motion fields in pairwise estimation [79,80].




Fig. 7. Sparse domains for CS reconstruction. The reconstructed MR image in (a) has sparse representations in the total variation (b) and wavelet (c) domains as only a small fraction of the total coefficients in these domains are non-zero. CS compressed sensing, MR magnetic resonance

The CS framework is also applicable to non-real-time imaging applications where retrospective data sorting can be applied. The extra-dimensional iGRASP (XD-GRASP) [83] method arranges data in multiple temporal dimensions (depending on the application) and then applies different sparsity transforms along each. For instance, in an XD-GRASP reconstruction of 2D free-breathing cardiac imaging, the position in the cardiac cycle and the position in the respiratory cycle were used as the temporal dimensions [83]. The concept has also been applied for five-dimensional free-running cardiac imaging [84,85] (three spatial dimensions, plus respiratory, and cardiac temporal dimensions) with a 3D golden-angle radial trajectory, where both the respiratory and cardiac signals are able to be extracted from the acquired k-space center, alleviating any requirement for a synchronized electrocardiogram signal [85].

An important consideration in the implementation of CS-based reconstructions is the choice of the penalty weighting(s) λ , as seen in Eq. (9). In general, no universally optimal value of λ exists, and the empirically determined values often used depend on multiple factors such as the acceleration factor, the sampling scheme, the image type, the data scaling, and the sparsifying transform [86]. Selecting a too-low value results in insufficient regularization, leading to an increased level of noise and artifacts in the image. Overly large values instead prioritize sparsity over data fidelity and tend to excessively smooth or blur the image in both the spatial and temporal dimensions [83]. More recently, techniques that optimize these weighting(s) over databases, as occurs in deep learning (discussed later), have been proposed [165].

Another class of methods closely related to CS are low-rank approaches [87–90]. These are similarly constructed as a regularized minimization problem, but the sparsity penalty of CS is replaced with a penalty based on the rank of a Casorati matrix M , which is constructed from the dynamic image ρ in various ways (discussed below). The problem to solve becomes

$$\arg \min_{\rho} \|E\rho - \mathbf{b}\|_2^2 + \lambda R(M), \quad (10)$$

where $R(M)$ is the rank of M [91]. Rank is a matrix property given by the number of linearly independent vectors that are required such that any column of the matrix can be constructed via a linear combination of the vectors. That is, if the matrix M consists of K columns and has rank 3, then there exist vectors \mathbf{a}_1 , \mathbf{a}_2 , and \mathbf{a}_3 such that every column \mathbf{m}_k can be written as

$$\mathbf{m}_k = c_1 \mathbf{a}_1 + c_2 \mathbf{a}_2 + c_3 \mathbf{a}_3, \quad (11)$$

a linear combination of the vectors. The idea is depicted in Fig. 8 where similar patches in the initial image are located and their pixels stacked in the Casorati matrix [92–94]. By applying a low-rank constraint to the matrix, incoherent noise is removed from the patch and aggregating the patches returns a denoised image.




Fig. 8. Low-rank patch-based denoising. Similar patches in the noisy image are collected and a Casorati matrix is formed from their pixel values. By reducing the rank of the matrix, incoherent noise is removed from the patches. Finally, the patches are aggregated to form a denoised image

Reconstruction methods that enforce the low rankness of Casorati matrices bear similarities to CS reconstruction, in that sparsity is enforced in the coefficient domain. That is, low-rank matrices require few non-zero coefficients c_1 , c_2 , etc. The key difference is that the sparsifying transformation is in this case dependent on the matrix (or image).

For dynamic imaging applications, low-rank methods can be used under the assumption that temporal variations can be well-described with only a few basis functions [91,95–97]. In k-t sparsity and low-rank (SLR) [91], the matrix is arranged such that each row contained the time series for a given pixel, and a SVD is applied to enforce low rankness by retaining only the highest singular values. In k-t principal component analysis (PCA) [95], each column of M instead contains the temporal frequency profile of a given pixel, and a PCA is applied to the columns. Using only the first 10 principal components (enforcing $R(M) = 10$), improvements can be seen over k-t BLAST.

Low-rank and sparse techniques have also been combined [82,98–100]. For example, the $L + S$ method [99] represents the high-dimensional image as the sum of a low-rank matrix L , which models the slowly varying background component, and a sparse matrix S , which captures the more-dynamic motion. This approach was able to provide high acceleration factors in a range of dynamic applications, including cardiac cine and cardiac perfusion.

7. Machine learning

In recent years, much attention has been focused on emerging machine-learning methods for image reconstruction. These approaches demonstrate promising potential for the next generation of reconstruction techniques, enabling increased acceleration factors, high reconstruction quality, and reduced computational time [101–106].

The fundamental idea behind machine learning is that, with a sufficient amount of training data, correlations in the data can be iteratively learnt. For the deep-learning networks often employed in MRI, multiple layers exist between the input and output of the network, and parameters that determine how each layer is calculated from the previous one are learnt during the training stage. Since the prior information learnt by the network is derived from the data itself, rather than being a human-designed constraint imposed on the reconstruction as in CS, it is reasonable to expect deep learning may lead to improved reconstruction quality. Additionally, deep-learning networks have the advantage that, while the training stage may be slow and computationally expensive, the actual reconstruction stage is generally very fast. Nevertheless, optimal design of the network remains an open question.

As the technology matures, increased clinical adoption of deep-learning reconstruction is expected and present limitations will need to be addressed [107,108]. These include the availability of large and unbiased datasets for training; the generalizability of methods beyond the scanner vendors, field strengths, and imaging sequences they were

trained with; the interpretability of the networks, particularly important in establishing trust in the methods by clinicians; and the need for widespread clinical validation.

In this section, we consider proposed techniques as separated into three categories: supervised techniques, which rely on the existence of large databases of high-quality training data, self-supervised techniques, which do not require fully sampled training data, and subject-specific techniques, in which a distinct set of network parameters are learnt every time as part of the reconstruction process. Deep-learning-based image reconstruction is overviewed generally, with specific applications and extensions for the dynamic case highlighted.

7.1. Supervised techniques

Perhaps the most widely researched class of deep-learning techniques proposed for MRI reconstruction are supervised techniques, which rely on the existence of a fully sampled ground-truth image for every undersampled dataset used during the training stage, in order to calculate a loss between the known ground-truth and the deep-learning reconstruction. Alternatively, networks can be trained under a supervised regime to “mimic” much slower iterative methods such as CS by treating the reconstructions from those methods as the ground-truth objective.

One of the first MRI-reconstruction applications of supervised deep learning was a proposal to use a convolutional neural network (CNN) to obtain fully sampled images from the input of an undersampled image [101]. Here, convolutions with learnt kernels are applied between layers, and the network is trained using many fully sampled MR datasets that are retrospectively undersampled to provide zero-filled images. The network learns the implicit relationship between the zero-filled and fully sampled images, since the parameters are tuned iteratively to reduce the difference between the network output and the fully sampled images.

Many deep-learning MRI methods similarly tackle the problem as one of denoising a zero-filled image, as depicted in Fig. 9a, while others directly take the acquired k-space as an input (Fig. 9b) and thereby perform the dual roles of recovering missing data and performing an inverse Fourier transform [102,109]. Alternatively, data consistency constraints may be imposed with an unrolled network architecture [110–114], as seen in Fig. 9c. In this scenario, the network effectively takes the place of the regularization term as seen in CS and low-rank formulations.

Various supervised techniques have been proposed specifically for dynamic imaging applications [111–121]. For example, a deep cascade of CNNs with data consistency constraints was proposed and applied to real-time 2D cardiac cine images [114]. The network was able to reconstruct a sequence of 30 frames, each 11-fold undersampled, within 10 s and outperformed existing state-of-the-art approaches including k-t SLR [91] and the $L + S$ method [99]. Like UNFOLD, k-t SENSE and other methods that build on spatio-temporal sampling results, deep-learning networks are able to exploit the spatio-temporal correlations present in dynamic MR data. The deep cascade of CNNs is among several methods that utilize 3D (2D + time) convolutions [114–118]. Such spatio-temporal convolutions may be truly 3D, or can be constructed as a separable operation with a 2D spatial convolution followed by a 1D temporal convolution [122]. Separable convolutions decrease the number of trainable parameters to be learnt, allowing more efficient training [117,122], with the potential cost of a reduction in the network’s ability to exploit correlations. In myocardial perfusion, it has been shown that true 3D spatio-temporal convolutions offer improved performance in the reduction of aliasing artifacts relative to separable convolutions [118]. Other examples of leveraging spatio-temporal correlations with deep-learning include a method combining 3D (2D + time) convolutions in x-t space with 3D (2D + temporal frequency) convolutions in x-f space [113], a method utilizing recurrent

CNNs [119], and CINENet, which performs separable four-dimensional (4D) (3D + time) convolutions to enable 3D cine in a single breath-hold with a reconstruction time of just ~ 5 s [120].

As was the case for CS methods, explicit ME and correction can be incorporated into deep-learning dynamic reconstruction techniques. Building on successful networks such as FlowNet [123] and recurrent all-pairs field transforms [124], which have been proposed for optical flow applications in computer vision, various networks have been proposed for ME in MRI [110–112,125–129]. These networks take, as input, either a pair [126–129] or group [111,112,125] of images in different motion states, and output non-rigid motion field(s) between these states. ME has been demonstrated for both respiratory [110,128] and cardiac [111,112,125,126] motion. By estimating both forwards and backwards motion fields, or estimating diffeomorphic fields which can be inverted [129,130], the motion can be incorporated into the encoding operator of a data consistency term [131]. When incorporated as part of an unrolled iterative reconstruction, the ME network parameters can be learnt to optimize objectives based not only on the correctness of their estimated motion fields, but also on the quality of the final reconstructed images [110–112]. Recent work leveraging these concepts for cardiac cine was able to demonstrate excellent image quality from a 20-fold undersampled Cartesian acquisition [111].

7.2. Self-supervised techniques

The availability of training data of sufficient quality and quantity can pose a significant challenge to the feasibility of supervised techniques. This is particularly true in dynamic applications where it may be impossible to sample k-t space fully at the desired spatial and temporal resolution. To address this problem, various self-supervised methods, where training is implemented without fully sampled ground-truth images (or without ground-truth images paired with undersampled data), have been proposed [132–142].

The self-supervised learning via data undersampling (SSDU) [132] method, for instance, uses a training set comprised entirely of undersampled k-space acquisitions. The approach by which the framework is trained to recover high-quality images despite not being exposed to ground-truth images is depicted in Fig. 9d. Each undersampled k-space is separated into two disjoint subsets of k-space samples with a higher level of undersampling. An unrolled network incorporating data consistency is used to reconstruct images given the input of the zero-filled image corresponding to one subset of samples, with the loss calculated in k-space against the second subset.

The SSDU technique bears conceptual similarities with Noise2Noise [143], which was proposed in computer vision for denoising images and trained only with pairs of distinct noisy images of the same scene, and is one of a series of methods utilizing similar ideas for MRI reconstruction [133–140,144]. Among these is Phase2-Phase [134], where the distinct undersampled k-spaces corresponded to data acquired in adjacent respiratory phases, DC-SiamNet [137], which utilizes contrastive learning to account for the overall image structure, and NLIN-Net [139] which jointly estimates coil sensitivity maps alongside the images. The latter was trained and applied to reconstruct real-time 2D cine images from ~ 19 -fold undersampled radial acquisitions, achieving comparable image quality to a CS implementation which could not be implemented online due to its computational demand [139]. SSDU itself has also been applied to dynamic CMR; a version utilizing multiple k-space mask pairs for each training sample [141] was applied to a simultaneous multi-slice acquisition of myocardial perfusion, with four-fold undersampling in each slice and a temporal resolution of 116 ms, where it was demonstrated to provided improved image quality and sharpness relative to CS-based techniques [142].

Generative and diffusion-based techniques have also been proposed for MRI and are trained to generate samples from a known image

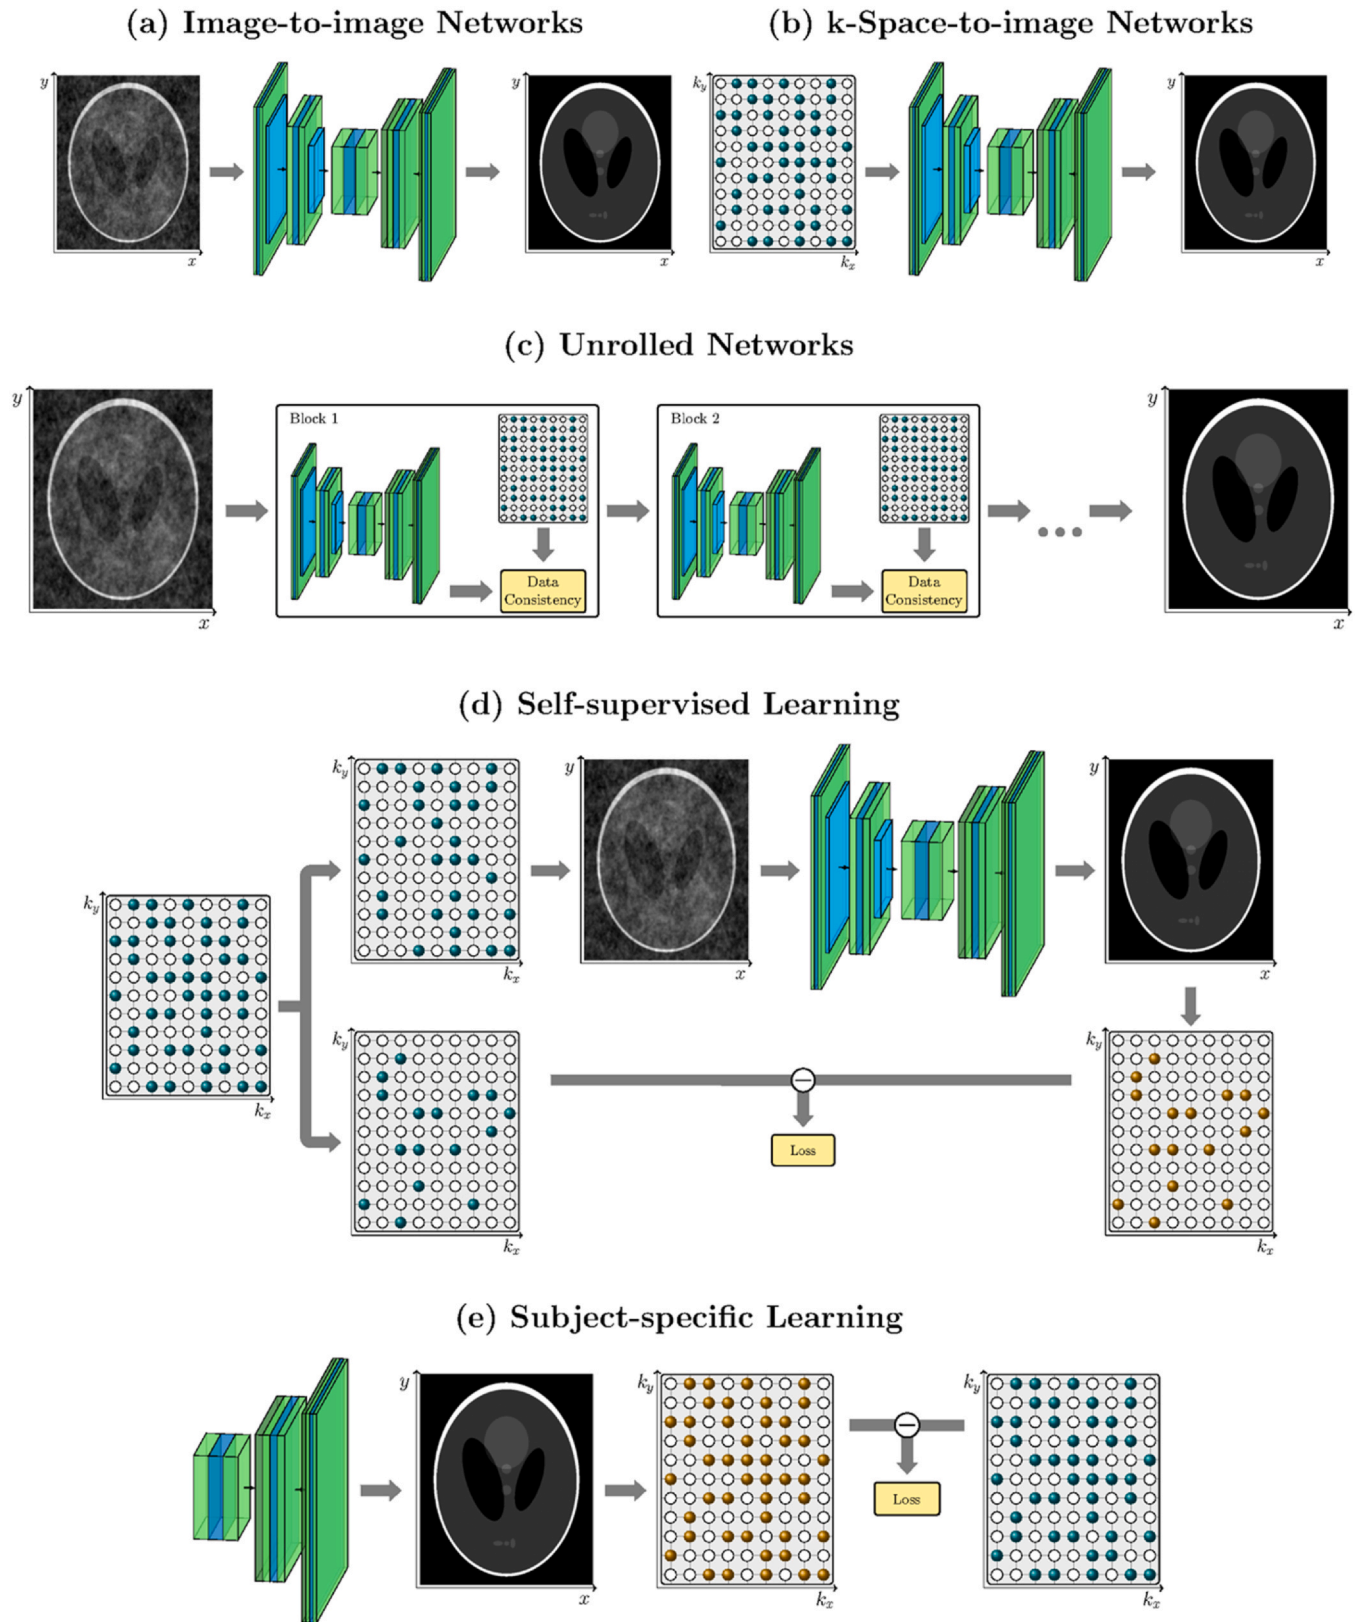


Fig. 9. Deep-learning neural networks for image reconstruction. Network layers are formed of various operations, such as convolutions, non-linear activations, batch normalizations, max-pooling, and up-sampling. The input of the network can be in the image (a) or k-space (b) domains. Alternatively, networks may be embedded in unrolled iterative reconstructions (c). When fully sampled data are not available for training, networks may be trained from undersampled data in a self-supervised manner (d) or trained to determine network parameters unique to one specific image (e)

distribution [145–148]. Since the training does not require paired samples of ground-truth and undersampled images, the high-quality images used in training need not exactly match the desired output. For instance, fully sampled static 2D images could be used in a framework proposed for cine reconstruction [146]. In another proposed generative-model-based technique, the network was trained using only undersampled data, and, in a similar manner to SSDU, the difference between a separate k-space acquisition and an undersampled k-space synthesized from the reconstruction was compared, here with a discriminator network [149]. While these techniques have yielded promising results for static MR applications, there has currently only been limited investigation into their use for dynamic reconstruction problems [150] and further research is warranted.

7.3. Subject-specific techniques

Subject-specific deep-learning methods form another class of reconstruction techniques [151–163]. These utilize deep neural networks while obviating the need for large amounts of training data. When applied to MRI, an unsupervised network can be trained to output an image (or series of images, for dynamic applications) given the input of either a latent variable [152–154] or image with random pixel values [155–157]. A forward model can then be applied to the output image, yielding corresponding k-space values which are compared against the acquired k-space to determine the network training loss, as is depicted in Fig. 9e. Since the network does not take specific information from the scan input, such as acquired k-space samples or a zero-filled image, and is not trained prior to reconstruction, a unique set of network parameters must be learnt for every scan. Thus, removing the requirement for large amounts of training data comes at the cost of long training/reconstruction times. The random-pixel input image case is known as the deep image prior approach [164] and relies on a CNN architecture. It has been successfully applied to stationary cardiac magnetic resonance fingerprinting (MRF) [155] as well as 2D cine MRF [156,157]. Its success lies in the fact that the CNN architecture is able to learn natural image structure faster than it learns noise, and provided the network training is stopped at an appropriate time an image that well-fits the acquired k-space measurements but has not been over-fitted to noise may be obtained. When the technique was adapted for dynamic cine imaging [156], two networks were utilized in a manner that allowed temporal correlations to be exploited. The first generated a spatial basis consisting of a series of images and the second generated a temporal basis that determined the weights for combining the spatial basis images at each time point.

The SSDU method has also been adapted for subject-specific reconstruction in the scenario that a training database of undersampled scans is not available [158]. In this case, the acquired undersampled data are split into three disjoint subsets, two for input and loss-calculation, as before, with the third deployed for self-validation, to prevent over-fitting. Applied to real-time cine [159], the subject-specific SSDU was able to achieve good quality images in cases where the spatio-temporal correlations were not effectively learned in database-trained SSDU.

Another technique for achieving subject-specific deep-learning-based reconstruction involves utilizing implicit neural representations (INRs) [160–163]. INRs use a neural network, typically a fully connected multi-layer perceptron, to represent the image as a continuous function of the image-space coordinates x , y and, for 3D images, z .

Again, these networks must be trained for every image or image series being reconstructed. At inference, each coordinate corresponding to a pixel position is passed to the network, and the network outputs the value of that pixel. In this way, INRs can also be employed to perform implicit super resolution, since the resolution of the pixel grid can be set arbitrarily. In dynamic cardiac MR, INRs have been proposed for 2D [160,161] and 3D [162] cine. Additionally, a k-space-based INR has been proposed [163], where k-space values corresponding to a particular coil and time-point, rather than pixel values, were produced by the network. The technique was able to achieve 30-frame cine images with $2\text{ mm} \times 2\text{ mm}$ resolution from a single heartbeat with good image quality, significantly outperforming CS- and low-rank-plus-sparse-based reconstructions.

8. Conclusion

The field of dynamic cardiac MR image reconstruction has advanced greatly over the past three decades and remains today an active area of research. From its early days when low-resolution 2D dynamic images were only possible at low temporal resolution, now sophisticated reconstruction techniques exploiting parallel imaging, spatio-temporal redundancies, prior information, and the latest machine learning architectures are able to provide good-quality high-resolution images at high temporal resolution. Without such innovations, the role of modern MRI in dynamic imaging applications such as cardiac cine, speech MRI, and perfusion would not be possible.

Funding

The authors acknowledge financial support from (1) King's BHF Centre for Award Excellence PG/18/59/33955 and RG/20/1/34802, (2) EPSRC EP/V044087/1 (3) Wellcome EPSRC Centre for Medical Engineering (NS/A000049/1), (4) Millennium Institute for Intelligent Healthcare Engineering ICN2021_004, FONDECYT 1210637 and 1210638, (5) IMPACT, Center of Interventional Medicine for Precision and Advanced Cellular Therapy, Santiago, Chile. ANID-Basal funding for Scientific and Technological Center of Excellence, IMPACT, #FB210024 (6) the Department of Health through the National Institute for Health Research (NIHR) comprehensive Biomedical Research Centre award, (7) NIHR Cardiovascular MedTech Co-operative, (8) the Technical University of Munich - Institute for Advanced Study and (9) the Government of Denmark. The views expressed are those of the authors and not necessarily those of the BHF, NHS, the NIHR or the Department of Health.

Author contributions

René M. Botnar: Writing – review & editing, Conceptualization. **Claudia Prieto:** Writing – review & editing, Conceptualization. **Andrew Phair:** Writing – review & editing, Writing – original draft, Conceptualization.

Declaration of competing interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Appendix

A.1. The Fourier transform

Mathematically, the k-space $b(\mathbf{k})$ is related to the image $\rho(\mathbf{x})$ by

$$b(\mathbf{k}) = \int \rho(\mathbf{x}) e^{-i\mathbf{k} \cdot \mathbf{x}} d\mathbf{x}, \quad (\text{A.1})$$

which we recognize as the continuous Fourier transform in two or three dimensions, depending on the number of k-space and image-space variables in the vectors \mathbf{k} and \mathbf{x} , respectively. The image can thus be recovered from the acquired k-space data by applying the inverse Fourier transform,

$$\rho(\mathbf{x}) = \frac{1}{(2\pi)^Q} \int b(\mathbf{k}) e^{i\mathbf{k} \cdot \mathbf{x}} d\mathbf{k}, \quad (\text{A.2})$$

where Q is the number of dimensions. In practice, both the acquired k-space data and the image being reconstructed consist of discrete data samples, and so the discrete inverse Fourier transform is instead considered. For 2D imaging, this may be written as

$$\rho_{u,v} = \frac{1}{N} \frac{1}{M} \sum_{n=1}^N \sum_{m=1}^M b_{n,m} e^{i(x_u(k_x)n + y_v(k_y)m)}, \quad (\text{A.3})$$

which can also be generalized to 3D.

A.2. SENSE equations

In SENSE, a set of linear equations may be written for each pixel in the reduced FOV describing how the aliased pixels are formed as the superposition of multiple coil-weighted pixels. In general, these take the form

$$\begin{aligned} a_1 &= s_{1,1}\rho_1 + s_{1,2}\rho_2 + \dots + s_{1,R}\rho_R \\ a_2 &= s_{2,1}\rho_1 + s_{2,2}\rho_2 + \dots + s_{2,R}\rho_R \\ &\vdots \\ a_{N_c} &= s_{N_c,1}\rho_1 + s_{N_c,2}\rho_2 + \dots + s_{N_c,R}\rho_R \end{aligned} \quad (\text{A.4})$$

and in the particular case illustrated in Fig. 4, where $N_c = 2$ and $R = 2$, they become

$$\begin{aligned} a_1 &= s_{1,1}\rho_1 + s_{1,2}\rho_2 \\ a_2 &= s_{2,1}\rho_1 + s_{2,2}\rho_2. \end{aligned} \quad (\text{A.5})$$

A.3. SMASH k-space synthesis

In SMASH, k-space line measurements acquired in different coils are linearly combined with weights chosen such that the combined coil sensitivity profile closely matches a spatial harmonic which modulates the exponential term in the integrand of the Fourier transform (Eq. (A.1)). Expressing the linear combination in terms of the Fourier transform definition of k-space gives

$$\sum_{j=1}^{N_c} g_j b_j(k_x, k_y) = \iint \sum_{j=1}^{N_c} g_j S_j(x, y) \rho(x, y) e^{-i(k_x x + k_y y)} dx dy \quad (\text{A.6})$$

where g_j are the weights. If the weights are chosen such that

$$\sum_{j=1}^{N_c} g_j S_j(x, y) = e^{im\Delta k_y y} \quad (\text{A.7})$$

then the combination becomes

$$\begin{aligned} \sum_{j=1}^{N_c} g_j b_j(k_x, k_y) &= \iint \rho(x, y) e^{-i(k_x x + (k_y - m\Delta k_y)y)} dx dy \\ &= b(k_x, k_y - m\Delta k_y), \end{aligned} \quad (\text{A.8})$$

which is the k-space shifted by m grid units in the k_y direction. Therefore, applying the linear combination to an acquired line in k-space yields a neighboring, unacquired line.

A.4. GRAPPA k-space synthesis

In GRAPPA, all acquired samples that sit within a k-space neighborhood of an unacquired sample contribute to its synthesis via a linear combination which utilizes a calibrated kernel of weights. This may be expressed as

$$b_\gamma(s, t) = \sum_j \sum_u \sum_v g_{u,v,\gamma,j} b_j(s + u\Delta k_x, t + v\Delta k_y), \quad (\text{A.9})$$

where $b_\gamma(s, t)$ is the k-space value at the location $(k_x, k_y) = (s, t)$ in the γ th coil, $g_{u,v,\gamma,j}$ is the kernel element (linear combination weight) used for the synthesis of the center sample in the γ th coil from the sample in the (u, v) kernel position in the j th coil, and the second and third summations are over all values of u and v that correspond to acquired samples in the neighborhood.

The applicability of weights calibrated in the ACS region to the rest of k-space can be demonstrated by considering Eq. (A.9) with the Fourier transform definition for the k-space samples (Eq. (A.1)) substituted in,

$$\iint \rho_\gamma(x, y) e^{-i(sx+ty)} dx dy = \iint \sum_j \sum_u \sum_v g_{u,v,\gamma,j} \rho_j(x, y) e^{-i((s+u\Delta k_x)x+(t+v\Delta k_y)y)} dx dy. \quad (\text{A.10})$$

Since the kernel elements $g_{u,v,\gamma,j}$ have been determined over the ACS region, rather than at a single position, we can assume this relationship is not a coincidental equality of integrals but instead due to the equality of the underlying integrands, as in the composite sensitivity profile which takes the place of Fourier basis functions in SMASH. Equating the integrands and multiplying both sides by $\exp(i(U\Delta k_x + V\Delta k_y))$, for arbitrary integers U and V , yields

$$\rho_\gamma(x, y) e^{-i((s+U\Delta k_x)x+(t+V\Delta k_y)y)} = \sum_j \sum_u \sum_v g_{u,v,\gamma,j} \rho_j(x, y) e^{-i((s+(u+U)\Delta k_x)x+(t+(v+V)\Delta k_y)y)}. \quad (\text{A.11})$$

Integrating both sides of Eq. (A.11) with respect to x and y gives

$$b_\gamma(s + U\Delta k_x, t + V\Delta k_y) = \sum_j \sum_u \sum_v g_{u,v,\gamma,j} b_j(s + (u + U)\Delta k_x, t + (v + V)\Delta k_y), \quad (\text{A.12})$$

which describes, as required, the synthesis of a sample at the shifted location $(k_x, k_y) = (s + U\Delta k_x, t + V\Delta k_y)$ using the neighboring samples in the same relative positions and the same kernel of weights.

A.5. *k-t BLAST and k-t SENSE equations*

For a *k-t BLAST* acquisition, if a total of H points β in x - f space alias to the same position α , these points must satisfy

$$\alpha = \beta_1 + \beta_2 + \dots + \beta_H, \quad (\text{A.13})$$

or, in the multi-coil case where there is an x - f space corresponding to each coil and the coil sensitivity profiles are assumed to be time-invariant,

$$\begin{aligned} \alpha_1 &= s_{1,1}\beta_1 + s_{1,2} & \beta_2 + \dots + s_{1,H}\beta_H \\ \alpha_2 &= s_{2,1}\beta_1 + s_{2,2} & \beta_2 + \dots + s_{2,H}\beta_H \\ & \vdots & \vdots \\ \alpha_{N_c} &= s_{N_c,1}\beta_1 + s_{N_c,2} & \beta_2 + \dots + s_{N_c,H}\beta_H. \end{aligned} \quad (\text{A.14})$$

References

- [1] Cooley JW, Tukey JW. An algorithm for the machine calculation of complex Fourier series. *Math Comput* 1965;19(90):297–301.
- [2] Nyquist H. Certain topics in telegraph transmission theory. *Trans Am Inst Elect Eng* 1928;47(2):617–44.
- [3] Shannon CE. A mathematical theory of communication. *Bell Syst Tech J* 1948;27(3):379–423.
- [4] Jackson JI, Meyer CH, Nishimura DG, Macovski A. Selection of a convolution function for Fourier inversion using gridding (computerised tomography application). *IEEE Trans Med Imaging* 1991;10(3):473–8.
- [5] Pipe JG, Menon P. Sampling density compensation in MRI: rationale and an iterative numerical solution. *Magn Reson Med* 1999;41(1):179–86.
- [6] Belliveau JW, Rosen BR, Kantor HL, Rzedzian RR, Kennedy DN, McKinstry RC, et al. Functional cerebral imaging by susceptibility-contrast NMR. *Magn Reson Med* 1990;14(3):538–46.
- [7] Belliveau J, Kennedy D, McKinstry R, Buchbinder B, Weisskoff R, Cohen M, et al. Functional mapping of the human visual cortex by magnetic resonance imaging. *Science* 1991;254(5032):716–9.
- [8] Van Vaals JJ, Brummer ME, Thomas Dixon W, Tuithof HH, Engels H, Nelson RC, et al. “Keyhole” method for accelerating imaging of contrast agent uptake. *J Magn Reson Imaging* 1993;3(4):671–5.
- [9] Jones R, Haraldseth O, Müller T, Rinck P, Øksendal A. K-space substitution: a novel dynamic imaging technique. *Magn Reson Med* 1993;29(6):830–4.
- [10] Margosian P. Faster MR imaging: imaging with half the data. *Proceedings of the 4th Annual Meeting of the SMRM, London; 1985, p. 1024.*
- [11] Feinberg DA, Hale JD, Watts JC, Kaufman L, Mark A. Halving MR imaging time by conjugation: demonstration at 3.5 kG. *Radiology* 1986;161(2):527–31.
- [12] Perman WH, El-Ghazzawy O, Gado MH, Larson KB, Perlmutter JS. A half-Fourier gradient echo technique for dynamic MR imaging. *Magn Reson Imaging* 1993;11(3):357–66.
- [13] Hu X, Parrish T. Reduction of field of view for dynamic imaging. *Magn Reson Med* 1994;31(6):691–4.
- [14] Fredrickson JO, Pelc NJ. Temporal resolution improvement in dynamic imaging. *Magn Reson Med* 1996;35(4):621–5.
- [15] Roemer PB, Edelstein WA, Hayes CE, Souza SP, Mueller OM. The NMR phased array. *Magn Reson Med* 1990;16(2):192–225.
- [16] Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: sensitivity encoding for fast MRI. *Magn Reson Med* 1999;42(5):952–62.
- [17] Griswold MA, Jakob PM, Nittka M, Goldfarb JW, Haase A. Partially parallel imaging with localized sensitivities (PILS). *Magn Reson Med* 2000;44(4):602–9.
- [18] Kyriakos WE, Panych LP, Kacher DF, Westin C-F, Bao SM, Mulkern RV, et al. Sensitivity profiles from an array of coils for encoding and reconstruction in parallel (SPACE RIP). *Magn Reson Med* 2000;44(2):301–8.
- [19] Pruessmann KP, Weiger M, Börner P, Boesiger P. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magn Reson Med* 2001;46(4):638–51.
- [20] Breuer FA, Blaimer M, Heidemann RM, Mueller MF, Griswold MA, Jakob PM. Controlled aliasing in parallel imaging results in higher acceleration (CAIPIRINHA) for multi-slice imaging. *Magn Reson Med* 2005;53(3):684–91.
- [21] Sodickson DK, Manning WJ. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magn Reson Med* 1997;38(4):591–603.
- [22] Jakob PM, Griswold MA, Edelman RR, Sodickson DK. AUTO-SMASH: a self-calibrating technique for SMASH imaging. *Magn Reson Mater Phys Biol Med* 1998;7(1):42–54.
- [23] Heidemann RM, Griswold MA, Haase A, Jakob PM. VD-AUTO-SMASH imaging. *Magn Reson Med* 2001;45(6):1066–74.
- [24] Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, et al. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med* 2002;47(6):1202–10.
- [25] Lustig M, Pauly JM. SPIRiT: iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magn Reson Med* 2010;64(2):457–71.
- [26] Uecker M, Lai P, Murphy MJ, Virtue P, Elad M, Pauly JM, et al. ESPIRiT - an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. *Magn Reson Med* 2014;71(3):990–1001.
- [27] Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. Coil sensitivity maps for sensitivity encoding and intensity correction. *Proceedings of the 6th Annual Meeting of the ISMRM, Sydney; 1998, p. 2087.*
- [28] Walsh DO, Gmitro AF, Marcellin MW. Adaptive reconstruction of phased array MR imagery. *Magn Reson Med* 2000;43(5):682–90.
- [29] Griswold MA, Breuer F, Blaimer M, Kannengiesser S, Heidemann RM, Mueller M, et al. Autocalibrated coil sensitivity estimation for parallel imaging. *NMR Biomed* 2006;19(3):316–24.
- [30] McKenzie CA, Yeh EN, Ohliger MA, Price MD, Sodickson DK. Self-calibrating parallel imaging with automatic coil sensitivity extraction. *Magn Reson Med* 2002;47(3):529–38.
- [31] Wang J, Kluge T, Nittka M, Jellus V, Kuehn B, Kiefer B. Parallel acquisition techniques with modified SENSE reconstruction mSENSE. *Proceedings of the First Würzburg Workshop on Parallel Imaging Basics and Clinical Applications; 2001, p. 89.*
- [32] Robson PM, Grant AK, Madhuranthakam AJ, Lattanzi R, Sodickson DK, McKenzie CA. Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-based parallel imaging reconstructions. *Magn Reson Med* 2008;60(4):895–907.
- [33] Lin F-H, Kwong KK, Belliveau JW, Wald LL. Parallel imaging reconstruction using automatic regularization. *Magn Reson Med* 2004;51(3):559–67.
- [34] Hoge WS, Brooks DH, Madore B, Kyriakos WE. A tour of accelerated parallel MR imaging from a linear systems perspective. *Concepts Magn Reson Part A Educ J* 2005;27(1):17–37.
- [35] King K, Angelos L. SENSE image quality improvement using matrix regularization. *Proceedings of the 9th Annual Meeting of the ISMRM, Glasgow; 2001, p. 1771.*
- [36] Weiger M, Pruessmann KP, Boesiger P. Cardiac real-time imaging using SENSE. *Magn Reson Med* 2000;43(2):177–84.
- [37] Pruessmann KP, Weiger M, Boesiger P. Sensitivity encoded cardiac MRI. *J Cardiovasc Magn Reson* 2001;3(1):1–9.
- [38] Griswold MA, Heidemann RM, Jakob PM. Direct parallel imaging reconstruction of radially sampled data using GRAPPA with relative shifts. *Proceedings of the 11th Annual Meeting of the ISMRM, Toronto; 2003, p. 2349.*

- [39] Breuer FA, Seiberlich N, Blaimer M, Jakob PM, Griswold MA. Quantitative g-factors for radial GRAPPA. Proceedings of the 17th Annual Meeting of the ISMRM, Honolulu; 2009, p. 2712.
- [40] Seiberlich N, Ehses P, Duerk J, Gilkeson R, Griswold M. Improved radial GRAPPA calibration for real-time free-breathing cardiac imaging. *Magn Reson Med* 2011;65(2):492–505.
- [41] Wright KL, Lee GR, Ehses P, Griswold MA, Gulani V, Seiberlich N. Three-dimensional through-time radial GRAPPA for renal MR angiography. *J Magn Reson Imaging* 2014;40(4):864–74.
- [42] Heidemann RM, Griswold MA, Seiberlich N, Krüger G, Kannengiesser SA, Kiefer B, et al. Direct parallel image reconstructions for spiral trajectories using GRAPPA. *Magn Reson Med* 2006;56(2):317–26.
- [43] Seiberlich N, Lee G, Ehses P, Duerk JL, Gilkeson R, Griswold M. Improved temporal resolution in cardiac imaging using through-time spiral GRAPPA. *Magn Reson Med* 2011;66(6):1682–8.
- [44] Skare S, Newbould RD, Nordell A, Holdsworth SJ, Bammer R. An auto-calibrated, angularly continuous, two-dimensional GRAPPA kernel for propeller trajectories. *Magn Reson Med* 2008;60(6):1457–65.
- [45] Codella NC, Spincemaille P, Prince M, Wang Y. A radial self-calibrated (RASCAL) generalized autocalibrating partially parallel acquisition (GRAPPA) method using weight interpolation. *NMR Biomed* 2011;24(7):844–54.
- [46] Sayin O, Saybasili H, Zviman MM, Griswold M, Halperin H, Seiberlich N, et al. Real-time free-breathing cardiac imaging with self-calibrated through-time radial GRAPPA. *Magn Reson Med* 2017;77(1):250–64.
- [47] Luo T, Noll DC, Fessler JA, Nielsen J-F. A GRAPPA algorithm for arbitrary 2D/3D non-Cartesian sampling trajectories with rapid calibration. *Magn Reson Med* 2019;82(3):1101–12.
- [48] Chieh S-W, Kaveh M, Akçakaya M, Moeller S. Self-calibrated interpolation of non-Cartesian data with GRAPPA in parallel imaging. *Magn Reson Med* 2020;83(5):1837–50.
- [49] Franson D, Ahad J, Liu Y, Fyrdahl A, Truesdell W, Hamilton J, et al. Self-calibrated through-time spiral GRAPPA for real-time, free-breathing evaluation of left ventricular function. *Magn Reson Med* 2023;89(2):536–49.
- [50] Schwab F, Schwarz F, Dietrich O, Lanz T, Resmer F, Wichmann T, et al. Free breathing real-time cardiac cine imaging with improved spatial resolution at 3 T. *Invest Radiol* 2013;48(3):158–66.
- [51] Willis P, Bresler Y. Optimal scan for time-varying tomography. I. Theoretical analysis and fundamental limitations. *IEEE Trans Image Process* 1995;4(5):642–53.
- [52] Willis NP, Bresler Y. Optimal scan for time-varying tomography. II. Efficient design and experimental validation. *IEEE Trans Image Process* 1995;4(5):654–66.
- [53] Willis NP, Bresler Y. Lattice-theoretic analysis of time-sequential sampling of spatiotemporal signals. Part I. *IEEE Trans Inf Theory* 1997;43(1):190–207.
- [54] Willis NP, Bresler Y. Lattice-theoretic analysis of time-sequential sampling of spatiotemporal signals. Part II. Large space-bandwidth product asymptotics. *IEEE Trans Inf Theory* 1997;43(1):208–20.
- [55] Madore B, Glover GH, Pelc NJ. Unaliasing by Fourier-encoding the overlaps using the temporal dimension (UNFOLD), applied to cardiac imaging and fMRI. *Magn Reson Med* 1999;42(5):813–28.
- [56] Kellman P, Epstein FH, McVeigh ER. Adaptive sensitivity encoding incorporating temporal filtering (TSENSE). *Magn Reson Med* 2001;45(5):846–52.
- [57] Madore B. Using UNFOLD to remove artifacts in parallel imaging and in partial-Fourier imaging. *Magn Reson Med* 2002;48(3):493–501.
- [58] Madore B. UNFOLD-SENSE: a parallel MRI method with self-calibration and artifact suppression. *Magn Reson Med* 2004;52(2):310–20.
- [59] Tsao J, Boesiger P, Pruessmann KP. k-t BLAST and k-t SENSE: dynamic MRI with high frame rate exploiting spatiotemporal correlations. *Magn Reson Med* 2003;50(5):1031–42.
- [60] Breuer FA, Kellman P, Griswold MA, Jakob PM. Dynamic autocalibrated parallel imaging using temporal GRAPPA (TGRAPPA). *Magn Reson Med* 2005;53(4):981–5.
- [61] Huang F, Akao J, Vijayakumar S, Duensing GR, Limkeman M. k-t GRAPPA: a k-space implementation for dynamic MRI with high reduction factor. *Magn Reson Med* 2005;54(5):1172–84.
- [62] Tsao J. On the UNFOLD method. *Magn Reson Med* 2002;47(1):202–7.
- [63] Ponce IP, Blaimer M, Breuer FA, Griswold MA, Jakob PM, Kellman P. Auto-calibration approach for k-t SENSE. *Magn Reson Med* 2014;71(3):1123–9.
- [64] Blaimer M, Ponce IP, Breuer FA, Jakob PM, Griswold MA, Kellman P. Temporal filtering effects in dynamic parallel MRI. *Magn Reson Med* 2011;66(1):192–8.
- [65] Liang Z-P, Lauterbur PC. An efficient method for dynamic magnetic resonance imaging. *IEEE Trans Med Imaging* 1994;13(4):677–86.
- [66] Xu D, King KF, Liang Z-P. Improving k-t SENSE by adaptive regularization. *Magn Reson Med* 2007;57(5):918–30.
- [67] Lustig M, Lee JH, Donoho DL, Pauly JM. Faster imaging with randomly perturbed, under-sampled spirals and $\|L\|_1$ reconstruction. Proceedings of the 13th Annual Meeting of the ISMRM, Miami Beach. Citeseer; 2005, p. 685.
- [68] Lustig M, Donoho D, Pauly JM. Sparse MRI: the application of compressed sensing for rapid MR imaging. *Magn Reson Med* 2007;58(6):1182–95.
- [69] Lustig M, Donoho DL, Santos JM, Pauly JM. Compressed sensing MRI. *IEEE Signal Process Mag* 2008;25(2):72–82.
- [70] Lustig M, Santos JM, Donoho DL, Pauly JM. k-t SPARSE: high frame rate dynamic MRI exploiting spatio-temporal sparsity. Proceedings of the 14th Annual Meeting of the ISMRM, Seattle; 2006, p. 2420.
- [71] Gamper L, Boesiger P, Kozerke S. Compressed sensing in dynamic MRI. *Magn Reson Med* 2008;59(2):365–73.
- [72] Feng L, Grimm R, Block KT, Chandarana H, Kim S, Xu J, et al. Golden-angle radial sparse parallel MRI: combination of compressed sensing, parallel imaging, and golden-angle radial sampling for fast and flexible dynamic volumetric MRI. *Magn Reson Med* 2014;72(3):707–17.
- [73] Otazo R, Kim D, Axel L, Sodickson DK. Combination of compressed sensing and parallel imaging for highly accelerated first-pass cardiac perfusion MRI. *Magn Reson Med* 2010;64(3):767–76.
- [74] Feng L, Srichai MB, Lim RP, Harrison A, King W, Adluru G, et al. Highly accelerated real-time cardiac cine MRI using k-t SPARSE-SENSE. *Magn Reson Med* 2013;70(1):64–74.
- [75] Bassett EC, Kholmovski EG, Wilson BD, DiBella EV, Dossall DJ, Ranjan R, et al. Evaluation of highly accelerated real-time cardiac cine MRI in tachycardia. *NMR Biomed* 2014;27(2):175–82.
- [76] Jung H, Ye JC, Kim EY. Improved k-t BLAST and k-t SENSE using FOCUSS. *Phys Med Biol* 2007;52(11):3201.
- [77] Jung H, Sung K, Nayak KS, Kim EY, Ye JC. k-t FOCUSS: a general compressed sensing framework for high resolution dynamic MRI. *Magn Reson Med* 2009;61(1):103–16.
- [78] Asif MS, Hamilton L, Brummer M, Romberg J. Motion-adaptive spatio-temporal regularization for accelerated dynamic MRI. *Magn Reson Med* 2013;70(3):800–12.
- [79] Royuela-delVal J, Cordero-Grande L, Simmross-Wattenberg F, Martín-Fernández M, Alberola-López C. Nonrigid groupwise registration for motion estimation and compensation in compressed sensing reconstruction of breath-hold cardiac cine MRI. *Magn Reson Med* 2016;75(4):1525–36.
- [80] Royuela-delVal J, Cordero-Grande L, Simmross-Wattenberg F, Martín-Fernández M, Alberola-López C. A Jacobian weighted temporal total variation for motion compensated compressed sensing reconstruction of dynamic MRI. *Magn Reson Med* 2017;77(3):1208–15.
- [81] Mohsin YQ, Lingala SG, DiBella E, Jacob M. Accelerated dynamic MRI using patch regularization for implicit motion compensation. *Magn Reson Med* 2017;77(3):1238–48.
- [82] Tolouee A, Alirezaie J, Babyn P. Nonrigid motion compensation in compressed sensing reconstruction of cardiac cine MRI. *Magn Reson Imaging* 2018;46:114–20.
- [83] Feng L, Axel L, Chandarana H, Block KT, Sodickson DK, Otazo R. XD-GRASP: golden-angle radial MRI with reconstruction of extra motion-state dimensions using compressed sensing. *Magn Reson Med* 2016;75(2):775–88.
- [84] Feng L, Coppo S, Piccini D, Yerly J, Lim RP, Masci PG, et al. 5D whole-heart sparse MRI. *Magn Reson Med* 2018;79(2):826–38.
- [85] Di Sopra L, Piccini D, Coppo S, Stuber M, Yerly J. An automated approach to fully self-gated free-running cardiac and respiratory motion-resolved 5D whole-heart MRI. *Magn Reson Med* 2019;82(6):2118–32.
- [86] Jaspán ON, Fleysher R, Lipton ML. Compressed sensing MRI: a review of the clinical literature. *Br J Radiol* 2015;88(1056):20150487.
- [87] Liang Z-P. Spatiotemporal imaging with partially separable functions. In: 2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro. Rotterdam, Netherlands: IEEE; 2007, p. 988–91.
- [88] Haldar JP, Liang Z-P. Spatiotemporal imaging with partially separable functions: a matrix recovery approach. In: 2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro. Rotterdam, Netherlands: IEEE; 2010, p. 716–9.
- [89] Zhao B, Haldar JP, Brinegar C, Liang Z-P. Low rank matrix recovery for real-time cardiac MRI. In: 2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro. Rotterdam, Netherlands: IEEE; 2010, p. 996–9.
- [90] Haldar JP, Liang Z-P. Low-rank approximations for dynamic imaging. In: 2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro. Rotterdam, Netherlands: IEEE; 2011, p. 1052–5.
- [91] Lingala SG, Hu Y, DiBella E, Jacob M. Accelerated dynamic MRI exploiting sparsity and low-rank structure: kt SLR. *IEEE Trans Med Imaging* 2011;30(5):1042–54.
- [92] Akçakaya M, Basha TA, Goddu B, Goepfert LA, Kissinger KV, Tarokh V, et al. Low-dimensional-structure self-learning and thresholding: regularization beyond compressed sensing for MRI reconstruction. *Magn Reson Med* 2011;66(3):756–67.
- [93] Qu X, Hou Y, Lam F, Guo D, Zhong J, Chen Z. Magnetic resonance image reconstruction from undersampled measurements using a patch-based nonlocal operator. *Med Image Anal* 2014;18(6):843–56.
- [94] Bustin A, Ginami G, Cruz G, Correia T, Ismail TF, Rashid I, et al. Five-minute whole-heart coronary MRA with sub-millimeter isotropic resolution, 100% respiratory scan efficiency, and 3D-PROST reconstruction. *Magn Reson Med* 2019;81(1):102–15.
- [95] Pedersen H, Kozerke S, Ringgaard S, Nehrke K, Kim WY. k-t PCA: temporally constrained k-t BLAST reconstruction using principal component analysis. *Magn Reson Med* 2009;62(3):706–16.
- [96] Trémouilhac B, Dikaos N, Atkinson D, Arridge SR. Dynamic MR image reconstruction - separation from undersampled (k, t)-space via low-rank plus sparse prior. *IEEE Trans Med Imaging* 2014;33(8):1689–701.
- [97] Liu F, Li D, Jin X, Qiu W, Xia Q, Sun B. Dynamic cardiac MRI reconstruction using motion aligned locally low rank tensor (MALLRT). *Magn Reson Imaging* 2020;66:104–15.
- [98] Gao H, Rapacchi S, Wang D, Moriarty J, Meehan C, Sayre J, et al. Compressed sensing using prior rank, intensity and sparsity model (PRISM): applications in cardiac cine MRI. Proceedings of the 20th Annual Meeting of the ISMRM; Melbourne; 2012, p. 2242.
- [99] Otazo R, Candes E, Sodickson DK. Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components. *Magn Reson Med* 2015;73(3):1125–36.
- [100] Roohi SF, Zonoobi D, Kassim AA, Jaremko JL. Multi-dimensional low rank plus sparse decomposition for reconstruction of under-sampled dynamic MRI. *Pattern Recognit* 2017;63:667–79.
- [101] Wang S, Su Z, Ying L, Peng X, Zhu S, Liang F, et al. Accelerating magnetic resonance imaging via deep learning. In: 2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI). Prague, Czech Republic: IEEE; 2016, p. 514–7.
- [102] Sun J, Li H, Xu Z, et al. Deep ADMM-Net for compressive sensing MRI. *Adv Neural Inf Process Syst* 2016;29:10–8.
- [103] Hammernik K, Knoll F, Sodickson D, Pock T. Learning a variational model for compressed sensing MRI reconstruction. Proceedings of the 24th Annual Meeting of the ISMRM; Singapore; 2016, p. 1088.

- [104] Hammernik K, Klatzer T, Kobler E, Recht MP, Sodickson DK, Pock T, et al. Learning a variational network for reconstruction of accelerated MRI data. *Magn Reson Med* 2018;79(6):3055–71.
- [105] Hyun CM, Kim HP, Lee SM, Lee S, Seo JK. Deep learning for undersampled MRI reconstruction. *Phys Med Biol* 2018;63(13):135007.
- [106] Lin DJ, Johnson PM, Knoll F, Lui YW. Artificial intelligence for MR image reconstruction: an overview for clinicians. *J Magn Reson Imaging* 2021;53(4):1015–28.
- [107] Fotaki A, Puyol-Antón E, Chiribiri A, Botnar R, Pushparajah K, Prieto C. Artificial intelligence in cardiac MRI: is clinical adoption forthcoming? *Front Cardiovasc Med* 2022;8:818765.
- [108] Zhang Q, Fotaki A, Ghadimi S, Wang Y, Doneva M, Wetzl J, et al. Improving the efficiency and accuracy of cardiovascular magnetic resonance with artificial intelligence - review of evidence and proposition of a roadmap to clinical translation. *J Cardiovasc Magn Reson* 2024;26(2):101051.
- [109] Han Y, Sunwoo L, Ye JC. k-Space deep learning for accelerated MRI. *IEEE Trans Med Imaging* 2019;39(2):377–86.
- [110] Qi H, Hajhosseini R, Cruz G, Küstner T, Kunze K, Neji R, et al. End-to-end deep learning nonrigid motion-corrected reconstruction for highly accelerated free-breathing coronary MRA. *Magn Reson Med* 2021;86(4):1983–96.
- [111] Pan J, Hamdi M, Huang W, Hammernik K, Kuestner T, Rueckert D. Unrolled and rapid motion-compensated reconstruction for cardiac CINE MRI. *Med Image Anal* 2024;91:103017.
- [112] Yang J, Küstner T, Hu P, Liò P, Qi H. End-to-end deep learning of non-rigid groupwise registration and reconstruction of dynamic MRI. *Front Cardiovasc Med* 2022;9:880186.
- [113] Qin C, Duan J, Hammernik K, Schlemper J, Küstner T, Botnar R, et al. Complementary time-frequency domain networks for dynamic parallel MR image reconstruction. *Magn Reson Med* 2021;86(6):3274–91.
- [114] Schlemper J, Caballero J, Hajnal JV, Price A, Rueckert D. A deep cascade of convolutional neural networks for MR image reconstruction. In: *International Conference on Information Processing in Medical Imaging*. Boone, NC, USA: Springer; 2017, p. 647–58.
- [115] Hauptmann A, Arridge S, Lucka F, Muthurangu V, Steeden JA. Real-time cardiovascular MR with spatio-temporal artifact suppression using deep learning - proof of concept in congenital heart disease. *Magn Reson Med* 2019;81(2):1143–56.
- [116] Shen D, Ghosh S, Haji-Valizadeh H, Pathrose A, Schiffrs F, Lee DC, et al. Rapid reconstruction of highly undersampled, non-Cartesian real-time cine k-space data using a perceptual complex neural network (PCNN). *NMR Biomed* 2021;34(1):e4405.
- [117] Sandino CM, Lai P, Vasanawala SS, Cheng JY. Accelerating cardiac cine MRI using a deep learning-based ESPIRiT reconstruction. *Magn Reson Med* 2021;85(1):152–67.
- [118] Wang J, Weller DS, Kramer CM, Salerno M. DEep learning-based rapid Spiral Image REconstruction (DESIRE) for high-resolution spiral first-pass myocardial perfusion imaging. *NMR Biomed* 2022;35(5):e4661.
- [119] Qin C, Schlemper J, Caballero J, Price AN, Hajnal JV, Rueckert D. Convolutional recurrent neural networks for dynamic MR image reconstruction. *IEEE Trans Med Imaging* 2018;38(1):280–90.
- [120] Küstner T, Fuin N, Hammernik K, Bustin A, Qi H, Hajhosseini R, et al. CINENet: deep learning-based 3D cardiac CINE MRI reconstruction with multi-coil complex-valued 4D spatio-temporal convolutions. *Sci Rep* 2020;10(1):1–13.
- [121] Kofler A, Haltmeier M, Schaeffter T, Kolbitsch C. An end-to-end-trainable iterative network architecture for accelerated radial multi-coil 2D cine MR image reconstruction. *Med Phys* 2021;48(5):2412–25.
- [122] Tran D, Wang H, Torresani L, Ray J, LeCun Y, Paluri M. A closer look at spatio-temporal convolutions for action recognition. *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*; 2018, p. 6450–9.
- [123] Dosovitskiy A, Fischer P, Ilg E, Hausser P, Hazirbas C, Golkov V, et al. FlowNet: learning optical flow with convolutional networks. *Proceedings of the IEEE International Conference on Computer Vision*; 2015, p. 2758–66.
- [124] Teed Z, Deng J. Raft: recurrent all-pairs field transforms for optical flow. In: *Computer Vision-ECCV 2020: 16th European Conference*; August 23–28, 2020; Glasgow, UK; *Proceedings, Part II* 16. Cham, Switzerland: Springer; 2020, p. 402–19.
- [125] Hammernik K, Pan J, Rueckert D, Küstner T. Motion-guided physics-based learning for cardiac MRI reconstruction. In: *2021 55th Asilomar Conference on Signals, Systems, and Computers*. Pacific Grove, CA, USA: IEEE; 2021, p. 900–7.
- [126] Qin C, Bai W, Schlemper J, Petersen SE, Piechnik SK, Neubauer S, et al. Joint motion estimation and segmentation from undersampled cardiac MR image. In: *Machine Learning for Medical Image Reconstruction: First International Workshop, MLMI 2018, Held in Conjunction with MICCAI 2018, September 16, 2018; Granada, Spain, Proceedings 1*. Cham, Switzerland: Springer; 2018, p. 55–63.
- [127] Balakrishnan G, Zhao A, Sabuncu MR, Guttat J, Dalca AV. Voxelmorph: a learning framework for deformable medical image registration. *IEEE Trans Med Imaging* 2019;38(8):1788–800.
- [128] Qi H, Fuin N, Cruz G, Pan J, Kuestner T, Bustin A, et al. Non-rigid respiratory motion estimation of whole-heart coronary MR images using unsupervised deep learning. *IEEE Trans Med Imaging* 2020;40(1):444–54.
- [129] Munoz C, Qi H, Cruz G, Küstner T, Botnar RM, Prieto C. Self-supervised learning-based diffeomorphic non-rigid motion estimation for fast motion-compensated coronary MR angiography. *Magn Reson Imaging* 2022;85:10–8.
- [130] Dalca AV, Balakrishnan G, Guttat J, Sabuncu MR. Unsupervised learning for fast probabilistic diffeomorphic registration. In: *Medical Image Computing and Computer Assisted Intervention-MICCAI 2018: 21st International Conference*; September 16–20, 2018; Granada, Spain, *Proceedings, Part I*. Cham, Switzerland: Springer; 2018, p. 729–38.
- [131] Batchelor P, Atkinson D, Irarrazaval P, Hill D, Hajnal J, Larkman D. Matrix description of general motion correction applied to multishot images. *Magn Reson Med* 2005;54(5):1273–80.
- [132] Yaman B, Hosseini SAH, Moeller S, Ellermann J, Uğurbil K, Akçakaya M. Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data. *Magn Reson Med* 2020;84(6):3172–91.
- [133] Liu J, Sun Y, Eldeniz C, Gan W, An H, Kamilov US. RARE: image reconstruction using deep priors learned without groundtruth. *IEEE J Sel Top Signal Process* 2020;14(6):1088–99.
- [134] Eldeniz C, Gan W, Chen S, Fraum TJ, Ludwig DR, Yan Y, et al. Phase2Phase: respiratory motion-resolved reconstruction of free-breathing magnetic resonance imaging using deep learning without a ground truth for improved liver imaging. *Invest Radiol* 2021;56(12):809–19.
- [135] Gan W, Sun Y, Eldeniz C, Liu J, An H, Kamilov US. Deep image reconstruction using unregistered measurements without groundtruth. In: *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*. Nice, France: IEEE; 2021, p. 1531–4.
- [136] Millard C, Chiew M. A theoretical framework for self-supervised MR image reconstruction using sub-sampling via variable density Noisier2Noise. *IEEE Trans Comput Imaging* 2023;9:707–20.
- [137] Yan Y, Yang T, Zhao X, Jiao C, Yang A, Miao J. DC-SiamNet: deep contrastive Siamese network for self-supervised MRI reconstruction. *Comput Biol Med* 2023;167:107619.
- [138] Desai AD, Öztürkler BM, Sandino CM, Boutin R, Willis M, Vasanawala S, et al. Noise2Recon: enabling SNR-robust MRI reconstruction with semi-supervised and self-supervised learning. *Magn Reson Med* 2023;90(5):2052–70.
- [139] Blumenthal M, Fantinato C, Unterberg-Buchwald C, Haltmeier M, Wang X, Uecker M. Self-supervised learning for improved calibrationless radial MRI with NLINNet. *Magn Reson Med* 2024;92(6):2447–63.
- [140] Hu Y, Gan W, Ying C, Wang T, Eldeniz C, Liu J, et al. SPICER: self-supervised learning for MRI with automatic coil sensitivity estimation and reconstruction. *Magn Reson Med* 2024;92(3):1048–63.
- [141] Yaman B, Gu H, Hosseini SAH, Demirel OB, Moeller S, Ellermann J, et al. Multi-mask self-supervised learning for physics-guided neural networks in highly accelerated magnetic resonance imaging. *NMR Biomed* 2022;35(12):e4798.
- [142] Demirel OB, Yaman B, Shenoy C, Moeller S, Weingärtner S, Akçakaya M. Signal intensity informed multi-coil encoding operator for physics-guided deep learning reconstruction of highly accelerated myocardial perfusion CMR. *Magn Reson Med* 2023;89(1):308–21.
- [143] Lehtinen J, Munkberg J, Hasselgren J, Laine S, Karras T, Aittala M, et al. Noise2Noise: learning image restoration without clean data. In: *Proceedings of the 35th International Conference on Machine Learning*, 80 of *Proceedings of the 35th International Conference on Machine Learning*. Stockholm, Sweden: PMLR; 2018, p. 2965–74.
- [144] Huang P, Zhang C, Zhang X, Li X, Dong L, Ying L. Self-supervised deep unrolled reconstruction using regularization by denoising. *IEEE Trans Med Imaging* 2023;43:1203–13.
- [145] Oh G, Sim B, Chung H, Sunwoo L, Ye JC. Unpaired deep learning for accelerated MRI using optimal transport driven CycleGAN. *IEEE Trans Comput Imaging* 2020;6:1285–96.
- [146] Lei K, Mardani M, Pauly JM, Vasanawala SS. Wasserstein GANs for MR imaging: from paired to unpaired training. *IEEE Trans Med Imaging* 2020;40(1):105–15.
- [147] Chung H, Ye JC. Score-based diffusion models for accelerated MRI. *Med Image Anal* 2022;80:102479.
- [148] Güngör A, Dar SU, Öztürk Ş, Korkmaz Y, Bedel HA, Elmas G, et al. Adaptive diffusion priors for accelerated MRI reconstruction. *Med Image Anal* 2023;88:102872.
- [149] Cole EK, Ong F, Vasanawala SS, Pauly JM. Fast unsupervised MRI reconstruction without fully-sampled ground truth data using generative adversarial networks. *Proceedings of the IEEE/CVF International Conference on Computer Vision*; 2021, p. 3988–97.
- [150] Qiu S, Pan S, Liu Y, Zhao L, Xu J, Liu Q, et al. Spatiotemporal diffusion model with paired sampling for accelerated cardiac cine MRI. *Proceedings of the 32nd Annual Meeting of the ISMRM*. Singapore; 2024.
- [151] Yazdanpanah AP, Afacan O, Warfield SK. Non-learning based deep parallel MRI reconstruction (NLDPMRI). In: *Medical Imaging 2019: Image Processing*. California, USA: SPIE. 2019;10949, p. 21–7.
- [152] Zou Q, Ahmed AH, Nagpal P, Kruger S, Jacob M. Dynamic imaging using a deep generative STORM (Gen-STORM) model. *IEEE Trans Med Imaging* 2021;40(11):3102–12.
- [153] Zou Q, Torres LA, Fain SB, Higano NS, Bates AJ, Jacob M. Dynamic imaging using motion-compensated smoothness regularization on manifolds (MoCo-STORM). *Phys Med Biol* 2022;67(14):144001.
- [154] Yoo J, Jin KH, Gupta H, Yerly J, Stuber M, Unser M. Time-dependent deep image prior for dynamic MRI. *IEEE Trans Med Imaging* 2021;40(12):3337–48.
- [155] Hamilton JI. A self-supervised deep learning reconstruction for shortening the breathhold and acquisition window in cardiac magnetic resonance fingerprinting. *Front Cardiovasc Med* 2022;9:928546.
- [156] Hamilton JI, Truesdell W, Galizia M, Burris N, Agarwal P, Seiberlich N. A low-rank deep image prior reconstruction for free-breathing ungated spiral functional CMR at 0.55 T and 1.5 T. *Magn Reson Mater Phys Biol Med* 2023;36(3):451–64.
- [157] Hamilton JI, da Cruz GL, Rashid I, Walker J, Rajagopalan S, Seiberlich N. Deep image prior cine MR fingerprinting with B1 + spin history correction. *Magn Reson Med* 2024;91(5):2010–27.
- [158] Yaman B, Hosseini SAH, Akçakaya M. Zero-shot self-supervised learning for MRI reconstruction. *arXiv preprint arXiv:2102.07737*; 2021.
- [159] Demirel OB, Zhang C, Yaman B, Gulle M, Shenoy C, Leiner T, et al. High-fidelity database-free deep learning reconstruction for real-time cine cardiac MRI. In: *2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. Sydney, Australia: IEEE; 2023, p. 1–4.
- [160] Al-Haj Hemidi Z, Vogt N, Quillen L, Wehsbach C, Heinrich MP, Oster J. Cine JENSE: simultaneous cine MRI image reconstruction and sensitivity map estimation using neural representations. In: *International Workshop on Statistical Atlases and Computational Models of the Heart*. Cham, Switzerland: Springer; 2023, p. 467–78.

- [161] Catalán T, Courdurier M, Osses A, Botnar R, Sahli-Costabal F, Prieto C. One-heartbeat cine MRI with implicit neural representations reconstruction. Proceedings of the 32nd Annual Meeting of the ISMRM. Singapore; 2024.
- [162] Shao H-C, Mengke T, Deng J, Zhang Y. 3D cine-magnetic resonance imaging using spatial and temporal implicit neural representation learning (STINR-MR). *Phys Med Biol* 2024;69(9):095007.
- [163] Huang W, Li HB, Pan J, Cruz G, Rueckert D, Hammernik K. Neural implicit k-space for binning-free non-cartesian cardiac MR imaging. In: International Conference on Information Processing in Medical Imaging. Cham, Switzerland: Springer; 2023, p. 548–60.
- [164] Ulyanov D, Vedaldi A, Lempitsky V. Deep image prior. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2018, p. 9446–54.
- [165] Gu J, Yaman B, Moeller S, Ellermann J, Ugurbil K, Akçakaya, M. Revisiting ℓ_1 -wavelet compressed-sensing MRI in the era of deep learning. Proceedings of the National Academy of Sciences. National Academy of Sciences. 119, 33, e220106211, [10.1073/pnas.2201062119](https://doi.org/10.1073/pnas.2201062119).