

Parallel imaging and reconstruction techniques

9

Berkin Bilgic^{a,b} and Tolga Cukur^{c,d}

^a*Athinoula A. Martinos Center for Biomedical Imaging, Charlestown, MA, United States* ^b*Department of Radiology, Harvard Medical School, Boston, MA, United States* ^c*Department of Electrical and Electronics Engineering, Bilkent University, Ankara, Turkey* ^d*National Magnetic Resonance Research Center (UMRAM), Bilkent University, Ankara, Turkey*

Highlights

- The encoding power of receive arrays has significantly improved over the last two decades, and acceleration factors of an order of magnitude became attainable through parallel imaging. This implies that receive arrays now perform most of the image encoding, whereas gradient coils are contributing a much smaller portion.
- Spatial variation in coil sensitivities increases at ultra-high fields, which permit yet higher acceleration factors. Combined with the SNR advantage of higher fields, this renders submillimeter-resolution imaging feasible.
- While controlled aliasing strategies have enabled optimal utilization of the degrees of freedom in modern dense receive arrays, compressed sensing, low-rank constraints, and deep learning priors have provided a yet better trade-off between image SNR and contrast versus scan time.

9.1 Introduction

Image encoding in MRI has been mainly performed using gradient coils, which create spatially varying magnetic fields that encode spins' position into their resonant frequencies. The spatial encoding power of radiofrequency receive arrays has significantly improved over the last two decades, and acceleration factors of an order of magnitude and higher have become attainable through parallel imaging reconstruction. This implies that receive arrays are now performing most of the image encoding, whereas gradient coils are contributing a much smaller portion at high accelerations. As such, parallel imaging and advanced reconstruction algorithms have emerged as important tools in our MR physics arsenal, especially at ultra-high fields (UHF) where high acceleration rates permit acquisitions that probe brain structure and function at mesoscale resolutions.

Spatial variations in coil sensitivity profiles increase at higher fields, and these increased degrees of freedom allow higher acceleration rates to be achieved compared to lower field strength. Unfortunately, undersampling in k -space leads to an intrinsic \sqrt{R} SNR penalty (where R is the acceleration factor) since less data points are acquired, and thus, there is less "noise averaging" being performed. Indeed, at high acceleration factors, this intrinsic SNR penalty can prohibit high-resolution imaging at 3T or lower field strengths as the resulting images may become too noisy for practical use. The SNR boost provided by going to UHF directly counteracts this penalty and thus renders high acceleration

factors feasible. Coupled with the increased degrees of freedom in coil sensitivity profiles, high acceleration factors are well suited for UHF imaging, which enables unprecedented spatial and temporal resolutions to be achieved.

This chapter will start by focusing on the basics of parallel imaging and controlled aliasing that aim to fully utilize the degrees of freedom in dense receive arrays for highly undersampled acquisitions, calibration scans, and coil sensitivity estimation approaches that permit robust utilization of such techniques, and advanced model-based reconstructions and deep learning approaches that can utilize a priori or learned regularizers to further push the acceleration. The chapter will conclude by looking at the effect of going to higher field strengths on parallel imaging capability and examples that illustrate the power and potential of image encoding at UHF.

9.2 Undersampled acquisitions and fundamental parallel imaging approaches

Parallel imaging (PI) methods use the extra degrees of freedom in multichannel receive arrays to reconstruct images from undersampled acquisitions, thereby reducing the scan time. The reduction factor (R) denotes the amount of subsampling performed during the acquisition and is also called the acceleration factor. PI relies on the fundamental relation between the imaging field of view (FOV) and the spacing between the successive k -space lines, Δk , where $\text{FOV} = 1/\Delta k$. This implies that a large FOV requires dense sampling of k -space, and if the prescribed FOV is smaller than the object that is being imaged, aliasing occurs. PI reduces the FOV deliberately by skipping lines in k -space and then tries to resolve the aliasing in image space or estimate the missing lines by interpolating the acquired data in k -space. In Fig. 9.1A, $R = 2$ -fold acceleration is obtained by skipping every other k -space line, thereby

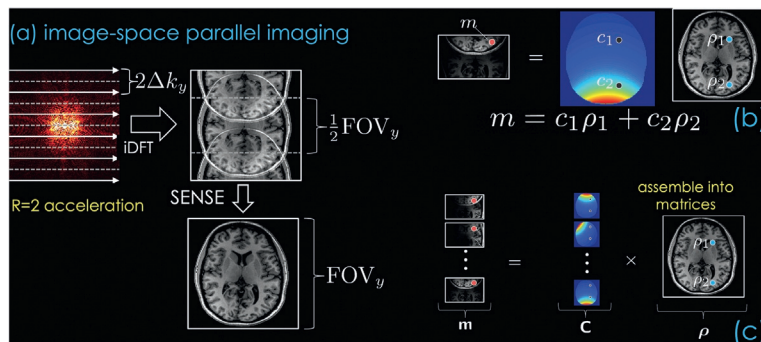


FIG. 9.1

(A) SENSE is a fundamental parallel imaging approach that operates in image space, where aliased signal replicas are unfolded using the coil sensitivity profiles explicitly. This is demonstrated in (B), where each coil contributes a linear equation that relates the aliasing voxels to the acquired coil signal. In (C), equations coming from all coils are assembled into a matrix, which leads to an overdetermined linear system that can be inverted to unalias the signal.

increasing the spacing between successive lines to $2\Delta k$. This, in turn, reduces the imaging FOV by half and causes the voxels that are FOV/2 apart to alias on top of each other.

Sensitivity encoding (SENSE) (Pruessmann et al., 1999) is a fundamental image-space approach where coil sensitivity profiles are used explicitly to unalias voxels by solving a set of linear equations. Generalized auto-calibrating partially parallel acquisition (GRAPPA) (Griswold et al., 2002) uses the spatial variations in coil sensitivities implicitly to synthesize missing k -space lines using a linear combination of the acquired data. Both image- and k -space approaches rely on auto-calibration signal (ACS) data to estimate either coil sensitivity maps or the linear combination weights required for k -space data interpolation.

The simultaneous acquisition of spatial harmonics (SMASH) technique (Sodickson and Manning, 1997) precedes both SENSE and GRAPPA and contains underlying principles of both of these influential approaches. SMASH seeks a linear combination of coil sensitivities that can explicitly form Fourier harmonics (resembling SENSE) and then uses these linear weights to synthesize missing k -space data (akin to GRAPPA). In the following, we will examine these image- and k -space PI algorithms in detail.

9.2.1 PI in image space

As shown in Fig. 9.1B, two voxels that are FOV/2 apart, ρ_1 and ρ_2 , will alias on top of each other in an acquisition at $R = 2$ -fold acceleration. Focusing on a particular coil with sensitivities denoted by c_1 and c_2 in these voxel positions, SENSE models the relation between the acquired aliased signal, m , and the unknown voxel intensities ρ_1 and ρ_2 as

$$m = c_1\rho_1 + c_2\rho_2 \quad (9.1)$$

indicating that the unknown magnetization is weighted by the coil sensitivity, then aliased in image space. Given that we have multiple coils, each providing one additional observation, we can assemble these linear equations into matrix format (Fig. 9.1C)

$$\mathbf{m} = C\rho. \quad (9.2)$$

Here \mathbf{m} is an $N_c \times 1$ vector of aliased signals in each coil, ρ are the unknown $R \times 1$ voxel intensities ($R = 2$ in this example), C is the $N_c \times R$ coil sensitivity matrix, and N_c denotes the number of coils.

The linear system in Eq. 9.2 can be solved in the least squares sense by minimizing

$$\begin{aligned} \rho^* &= \operatorname{argmin}_\rho \|C\rho - \mathbf{m}\|_2^2, \\ &= (C^H C)^{-1} C^H \mathbf{m}. \end{aligned} \quad (9.3)$$

If present, an $N_c \times N_c$ channel noise covariance matrix, Ψ , which describes the levels and correlation of noise in receiver coils, can be incorporated to obtain an SNR-optimal solution (Pruessmann et al., 1999)

$$\rho^*_{optSNR} = (C^H \Psi^{-1} C)^{-1} C^H \Psi^{-1} \mathbf{m}. \quad (9.4)$$

9.2.2 Geometry factor (g-factor)

The matrix inversion in Eq. 9.4 becomes more ill-conditioned at high acceleration rates, which leads to mathematical noise amplification during image reconstruction. This noise amplification is captured by the geometry factor (or g-factor) map and is given by (Pruessmann et al., 1999)

$$g_{\rho} = \left(\left[(C^H \Psi^{-1} C)^{-1} \right]_{\rho, \rho} [C^H \Psi^{-1} C]_{\rho, \rho} \right)^{1/2} \geq 1, \quad (9.5)$$

where the indices $[\]_{\rho, \rho}$ denote the diagonal elements corresponding to voxel ρ . Since the g-factor is greater than or equal to 1, it is common to report the inverse of the g-factor to ensure that it remains in the interval $[0, 1]$. $1/g$ -factor reports the retained SNR after parallel imaging reconstruction as shown in Fig. 9.2. Since coil sensitivity profiles are less spatially varying in the middle of the FOV, g-factor is often worse in these regions. This can be better visualized in the $R = 4$ case in Fig. 9.2, where the zoomed-in region of interest (ROI) exhibits higher noise in the middle (white and deep gray matter) and less noise in the periphery (cortex).

Although an $N_c = 8$ channel array was used in this simulation experiment, acceleration factors beyond $R = 4$ lead to prohibitive noise amplification. We still have more observations than unknowns in Eq. 9.2 since $N_c > R$, but not all these observations are linearly independent because coils that are in close proximity around the head contribute similar encoding information. Further, the acceleration in this experiment is performed only along the phase encoding direction (vertical axis), whereas the eight coils are distributed around the head in 2 dimensions. As such, variations in these eight coils are shared among the vertical and horizontal axes, which makes $R > 4$ -fold acceleration in only one direction difficult.

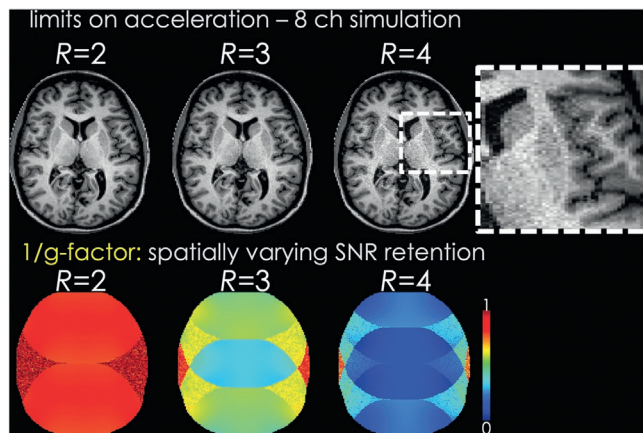


FIG. 9.2

An 8ch SENSE reconstruction simulation is conducted to show the limits of acceleration. While $R=2$ - and 3 -fold accelerations lead to successful reconstructions, noise amplification is visible especially in the $R=4$ case. The noise amplification is spatially varying and more severe in the middle of the FOV. This can also be seen in the $1/g$ -factor maps in the bottom row, where up to $\sim 80\%$ of SNR is lost during parallel imaging reconstruction in the middle of the FOV. The shape of the $1/g$ -factor maps also reflects the acceleration factor, e.g., 4 bands of aliased replicas are visible at $R=4$.

9.2.3 Generalized SENSE reconstruction

The encoding matrix in Eq. 9.4, $(C^H\Psi^{-1}C)^{-1}C^H\Psi^{-1}$, is very small (of size $R \times N_c$) and the equation can be computed rapidly, yet this needs to be performed for each of the collapsed voxels in the image, which may be computationally intensive. Apart from computing these small problems in parallel, another way to facilitate rapid computation is to represent the entire unknown 2D image as a vector \mathbf{p} , which is now of size $(N_x N_y) \times 1$ where N_x and N_y are the number of voxels in readout and phase-encoding axes. This leads to the more flexible relation

$$\mathbf{p}^* = \operatorname{argmin}_{\mathbf{p}} \sum_i \|DFC_i \cdot \mathbf{p} - y_i\|_2^2 \text{ with } i = 1, \dots, N_c, \quad (9.6)$$

where C_i now denotes a matrix whose diagonal entries are the sensitivities of the i th coil for the entire slice, F is a 2D discrete Fourier transform (DFT) operator, D is a k -space undersampling mask, and y_i denotes the acquired k -space data for the i th coil. Though represented as matrix-vector operations, this formulation permits rapid computations through element-wise multiplications that implement $(C_i \cdot \mathbf{p})$, whereas F is conveniently evaluated using the 2D fast Fourier transform (FFT). While it would still be possible to operate in image space by replacing the DFT operator with a “point spread function” matrix that explains image aliasing, the transition to k -space allows us to use nonuniform sampling patterns where the spacing between successive lines is not held constant. Such sampling strategies can be easily represented by the diagonal sampling mask D that has binary elements. Eq. 9.6 can be solved using least squares solvers such as conjugate gradients (CG) and provides the additional flexibility of incorporating regularization, which will be detailed in Section 9.4.

9.2.4 PI in k -space using GRAPPA

Unlike SENSE where coil sensitivity profiles are explicitly used in the reconstruction, GRAPPA uses them implicitly to estimate missing k -space data. This estimation is performed across coils and within a small neighborhood in k -space as described in the following:

$$s_j(k_x, k_y) = \sum_i \sum_{\mu \in U} \sum_{\nu \in V} w_j(i, \mu, \nu) s_i(k_x + \mu \Delta k_x, k_y + \nu \Delta k_y) \text{ with } i = 1, \dots, N_c. \quad (9.7)$$

Here, $s_j(k_x, k_y)$ is the target missing k -space sample in the j th coil (black point in Fig. 9.3A), and w_j is the k -space “kernel” used for interpolating the data for the j th coil. This kernel computes a linear combination of the acquired points $s_i(k_x + \mu \Delta k_x, k_y + \nu \Delta k_y)$ not only inside small neighborhoods in k -space (denoted with U and V for the k_x and k_y axes) but also across all the coils. The underlying assumption in GRAPPA is that there are linear dependencies across coils that permit such estimation, which is akin to SMASH.

As demonstrated in Fig. 9.3B, GRAPPA suffers from similar limitations as SENSE when the acceleration factor is pushed beyond $R > 4$ along a single dimension. At $R = 5$, this reconstruction using a 32ch head array begins to suffer from aliasing artifacts, which become severe at $R = 6$. At this latter acceleration factor, noise amplification especially in the middle of the FOV becomes apparent.

An advantage of GRAPPA over SENSE is in its kernel calibration. Until recently, estimation of coil sensitivities for SENSE has suffered from robustness issues, as they need to be smoothed and carefully masked to exclude nontissue signals (Section 9.5 will detail new techniques that boosted the robustness of sensitivity estimation). GRAPPA, on the other hand, does not need coil sensitivities to be explicitly known, but only requires the estimation of the kernels w_j . To make this possible, fully sampled ACS

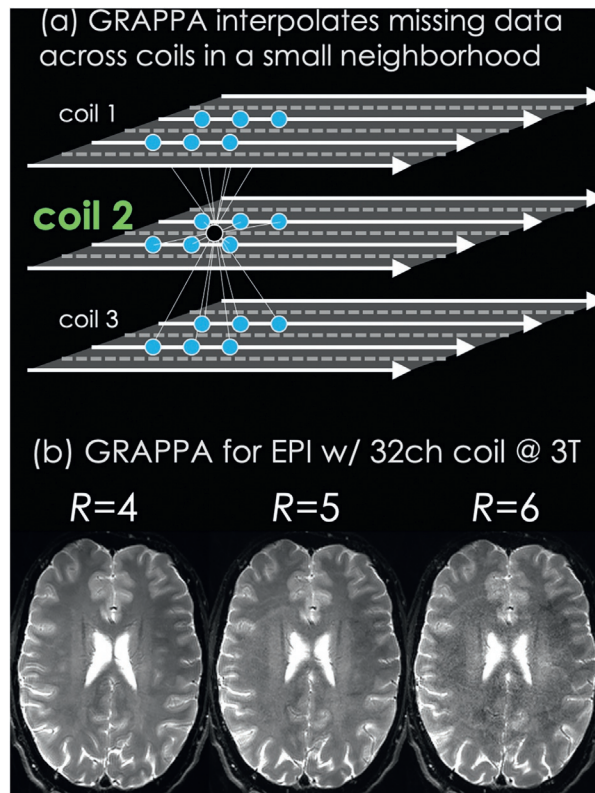


FIG. 9.3

(A) GRAPPA uses a linear combination of acquired k -space data in a small neighborhood across all coils to estimate a missing point. (B) Similar to SENSE, GRAPPA suffers from residual aliasing and noise amplification especially at higher acceleration factors of $R=5$ and 6 . Data were acquired using a single-shot echo planar imaging (EPI) readout with 32ch reception at 3T.

data are used. Here, both the target $s_j(k_x, k_y)$ and the neighboring k -space data $s_i(k_x + \mu\Delta k_x, k_y + \nu\Delta k_y)$ in Eq. 9.7 are known. As such, a large “calibration matrix” can be built by creating one row for each (k_x, k_y) point inside the ACS region. This then permits the estimation of the kernels in the least squares sense due to $W_j = S^\dagger S_j$, where W_j is the vectorized version of the k -space kernel w_j , $S^\dagger = (S^H S)^{-1} S^H$ is the pseudo-inverse of the calibration matrix, and S_j is the vector made from all the target points inside the ACS. The kernel calibration step also admits Tikhonov regularization, where the kernel estimate becomes $W_j = (S^H S + \lambda I)^{-1} S^H S_j$, with λ being the regularization parameter.

Another parameter that needs to be selected in GRAPPA reconstruction is the kernel size. Popular choices include 3×3 and 5×5 in $k_x \times k_y$ axes, but larger kernel sizes can be used as well. The trade-off is that smaller kernel sizes provide better g-factor performance, and larger kernels better mitigate potential reconstruction artifacts. This can be understood by thinking of GRAPPA kernels as k -space filters, where a smaller kernel will have more of a “low-pass” effect on the data, thus denoising it more

effectively. On the other hand, larger kernels can help synthesize higher order Fourier harmonics and thus reduce potential aliasing artifacts especially at higher accelerations. Finally, it is important to note that the ACS data that allow for kernel estimation do not have to have the same contrast as the actual undersampled imaging data. This point is valid for both sensitivity estimation for SENSE and kernel estimation for GRAPPA. As such, low-resolution, short echo time (TE) and short repetition time (TR) gradient echo, ACS data can be used for calibration, which usually require a couple of seconds to acquire. By this way, it becomes unnecessary to incorporate a fully sampled ACS region inside the actual imaging scan, which often has longer TE/TR combinations to provide the desired contrast. Significant time savings can be achieved with such “external” calibration scans.

Similar to SENSE, an analytical expression for g-factor for the j th coil in GRAPPA reconstructions can be defined using the matrix W made from *image-space* kernels (Breuer et al., 2009)

$$g_j = \left([W^H \Psi W]_{jj} / [\Psi]_{jj} \right)^{1/2} \geq 1. \quad (9.8)$$

Alternatively, Monte Carlo simulations can be performed by injecting noise to k -space and running GRAPPA reconstruction several times, then taking ratio between the standard deviation across these reconstructions and the standard deviation of the added noise to obtain an empirical g-factor estimate (Robson et al., 2008). This approach is powerful since it provides a general way to estimate g-factor maps for arbitrary trajectories and advanced reconstructions. A caveat is that, with CS regularization (e.g., L1 or low-rank), using different levels of standard deviation for the added noise may lead to different g-factor maps, since these regularizers may threshold the noise entirely depending on its level. In such cases, it might be helpful to acquire a noise-only reference scan to be able to synthesize noise with appropriate power.

9.2.5 PI in k -space using SPIRiT

Inspired by GRAPPA, SPIRiT takes a different approach to provide a more general solution that uses data more efficiently and also formulates the reconstruction as an optimization problem that can admit additional, e.g., compressed sensing (CS) regularizers (Lustig and Pauly, 2010). The main difference is that GRAPPA enforces consistency between the synthesized points and the neighboring acquired data, whereas SPIRiT enforces consistency between every point and its neighborhood across all coils. In other words, the acquired data should be able to generate the missing points, but also the missing points need to be able to synthesize the acquired points as well as other missing data. This also leads to a different definition of the reconstruction kernel. In GRAPPA, a kernel spans a large extent in k -space since only the acquired data are counted in the kernel size. For example, at $R_y = 3$ -fold acceleration, a 3×3 kernel would have a k -space extent of 3×7 in (k_x, k_y) axes. In SPIRiT, a 3×3 kernel implies a literal 3×3 neighborhood, where all 9 points contribute to each other. This is expressed via

$$\begin{aligned} s_j &= \sum_i g_{ij} \otimes s_i \text{ with } i = 1, \dots, N_c \\ \text{or} & \\ s &= Gs, \end{aligned} \quad (9.9)$$

where g_{ij} s denote the SPIRiT kernels, which are convolved with the entire k -space data s_i in the i th coil, then summed over all coils to yield the j th coil's reconstructed k -space, s_j . This is written more

succinctly as $s = Gs$, where s now denotes entire k -space concatenated across all coils, and G is a matrix containing $g_{ij}s$ in appropriate positions. Finally, SPIRiT enforces data consistency via

$$y = D \cdot s, \quad (9.10)$$

where D is a linear operator that selects only the acquired k -space locations out of the entire k -space grid, and y is a vector of the acquired k -space data concatenated together. The self- and data-consistency constraints in Eqs. 9.9 and 9.10 can be combined to yield the SPIRiT loss function

$$s^* = \operatorname{argmin}_s \|Ds - y\|_2^2 + \lambda \|(G - I) s\|_2^2, \quad (9.11)$$

where λ is a regularization parameter. We can eliminate the need for selecting this parameter by solving only for the missing k -space data, \hat{s} , through the formulation $s = D^T y + D_c^T \hat{s}$, where D^T and D_c^T select the acquired and nonacquired data and put them back in the full k -space grid (Lustig and Pauly, 2010). This leads to

$$\begin{aligned} \operatorname{argmin}_s & \|D(D^T y + D_c^T \hat{s}) - y\|_2^2 + \lambda \|(G - I)(D^T y + D_c^T \hat{s})\|_2^2 \\ & = \|y + 0 - y\|_2^2 + \lambda \|(G - I)(D^T y + D_c^T \hat{s})\|_2^2, \\ & = \operatorname{argmin}_{\hat{s}} \|(G - I)D^T y + (G - I)D_c^T \hat{s}\|_2^2, \end{aligned} \quad (9.12)$$

which can be solved using a standard CG solver.

SPIRiT is also auto-calibrating, where the kernel estimation is performed similar to GRAPPA and can admit Tikhonov regularization. Unlike GRAPPA which functions as a k -space filter, the iterative nature of Eq. 9.11 allows SPIRiT to use more involved regularizers during the reconstruction stage as well (Lustig and Pauly, 2010; Murphy et al., 2012), which will be detailed in Section 9.4.

9.3 Controlled aliasing in parallel imaging (CAIPI) and non-Cartesian trajectories

Modern head coils have large channel counts (32ch or higher) which are often distributed uniformly around the head. This makes it possible to accelerate in more than one phase encoding axis. 2D acceleration is a powerful concept and easily permits $R = 4$ -fold acceleration to be distributed as $R = 2 \times 2$ between two-phase encoding axes in 3D-encoded acquisitions (Weiger et al., 2002; Blaimer et al., 2006). Importantly, the concept of 2D acceleration can be generalized to multislice imaging through simultaneous multislice (SMS) encoding (Feinberg et al., 2011; Setsompop et al., 2012; Moeller et al., 2010; Nunes et al., 2006; Larkman et al., 2001). In SMS, multiple slices are excited simultaneously, which causes the acquired signal to be the superposition of all of these excited slices. Either image- or k -space based PI reconstruction is then used to unalias the collapsed slices (Zahneisen et al., 2015). SMS encoding admits undersampling in the phase encoding axis within each slice as well, which is denoted as “ $R_{inplane}$ ” acceleration. Simultaneous excitation of multiple slices is performed using tailored MultiBand (MB) RF pulses, thus leading to the nomenclature $R_{total} = R_{inplane} \times MB$, where R_{total} is the combined acceleration factor and MB is the number of simultaneously excited slices.

Acceleration in 2D allows the utilization of Controlled Aliasing in Parallel Imaging (CAIPI), where aliasing voxels can be pushed further apart in image space to improve g-factor performance. CAIPI is

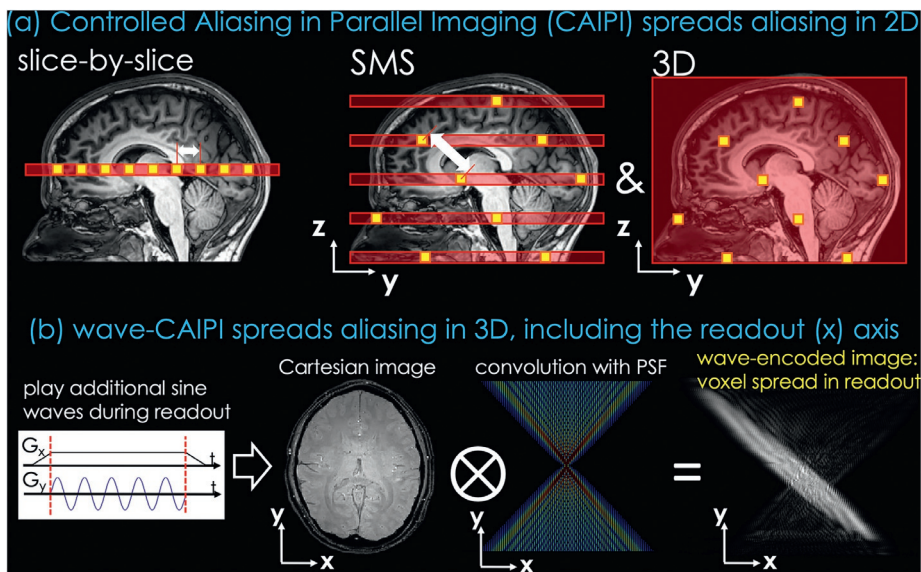


FIG. 9.4

(A) Standard 2D (slice-by-slice) imaging does not lend itself well to high acceleration, as aliasing voxels are very close to each other. When Simultaneous MultiSlice (SMS) or 3D encoding are used, two axes, phase encoding (y) and slice/partition (z) directions, become available to spread aliasing. Controlled aliasing in parallel imaging (CAIPI) improves on this by introducing slice-shifting in SMS and staggered sampling in 3D encoding to further push the distance between aliasing voxels. (B) In wave-CAIPI, the conventionally “fully-sampled” readout (x) direction can be used to spread aliasing in addition to the y - and z -axes, thus better harness the degrees of freedom in receive arrays in all three dimensions.

applicable to both 3D (Breuer et al., 2006) and SMS encoding (Breuer et al., 2005). In 3D imaging, this is achieved by altering the sampling pattern in the 2D phase encoding plane, which changes the aliasing pattern of the collapsing voxels in image space (Fig. 9.4A, right). This is particularly helpful for uniformly distributing the total acceleration factor across the 2D phase encoding axes. For instance, at $R_{total} = 8$, standard 2D sampling would be constrained to distribute aliasing either as $R_y \times R_z = 4 \times 2$ or 2×4 . The high $R = 4$ -fold acceleration in one of the axes would lead to a high g-factor penalty. Instead, with CAIPI, this can be distributed more evenly (i.e., $\sim\sqrt{8}$) by staggering the sampling pattern using Δk_y or Δk_z gradient blips.

Application of the CAIPI concept to SMS encoding is slightly more involved. In spin-warp imaging where k -space is acquired line-by-line, controlled aliasing can be achieved by modulating the phase of the MB RF pulse (Breuer et al., 2005). This permits shifting the aliasing slices with respect to each other in the phase encoding direction, thus increasing the distance between aliasing voxels (Fig. 9.4A, middle). In an $MB = 2$ experiment, the phase of the RF pulse exciting one of the slices can be alternated between 0 and π radians across successive phase encoding lines. This creates a phase ramp in k -space and causes an $FOV/2$ shift in image space in the y -direction. Such phase modulation ideas can be generalized to higher MB cases. For EPI readouts, it is not possible to alternate the phase of the RF pulse

between k -space lines since there is only a single RF per entire k -space. In these cases, blipped-CAIPI (Setsompop et al., 2012) can be utilized. Instead of RF phase modulation, FOV shifting is made possible by playing G_z gradient blips across k -space lines. G_z blips create a linearly varying phase deposition in the slice axis, whose amplitude can be adjusted to provide the desired FOV shift in the target slice positions. Application of SMS in EPI has been impactful, as it has allowed rapid diffusion and functional imaging acquisitions and has been popularized by the Human Connectome Project (Setsompop et al., 2013; Sotiropoulos et al., 2013).

It is possible to push the CAIPI idea further by also utilizing the conventionally fully sampled readout axis to spread aliasing. Zigzag GRAPPA achieves this by playing rapidly oscillating gradients on the G_y gradient during the readout (Breuer et al., 2008). As the name implies, this leads to a zigzag trajectory in k -space, which spreads the aliasing in x - as well as y -axes and improves acceleration capability. The reconstruction is performed using multiple GRAPPA kernels that conform to the shape of the non-Cartesian trajectory. In Bunched Phase Encoding (BPE), signal processing concepts were introduced to reconstruct such zigzag sampled data without the need for sensitivity encoding (Moriguchi and Duerk, 2006). Wave-CAIPI is a generalization of zigzag GRAPPA and BPE, where oscillating gradients are played on both G_y and G_z gradients (Bilgic et al., 2015). Using a sine on, e.g., G_y , and cosine waveform on G_z gradients creates a “corkscrew” trajectory in k -space. Different from the preceding methods, wave-CAIPI represents this non-Cartesian trajectory as a convolution with a point spread function (PSF) so that a fully Cartesian reconstruction becomes possible (Fig. 9.4B). This PSF formalism also helps explain the g-factor benefit of playing sinusoidal waveforms: taking an inverse DFT of the non-Cartesian data leads to image-space picture (Fig. 9.4B, right) where voxels are spread out in the readout direction. The amount of this spreading is a function of the y -position, e.g., there is little spreading in the middle of the FOV, but this increases toward the edge of the object. This way, aliasing voxels are pushed further apart from each other in the x -direction, in addition to the y - and z -axes that are exploited through standard controlled aliasing.

Controlled aliasing can be applied in other non-Cartesian trajectories as well. In radial-CAIPI, RF phase modulation is performed across radial k -space lines. In an $MB = 2$ experiment, alternating the phase between 0 and π radians in one of the slices largely cancels out the aliasing signal contributions from this slice (Yutzy et al., 2011). This permits separating out the two slices even without the aid of sensitivity encoding and renders total acceleration rates of $R_{total} = 10$ or higher feasible using a 12ch array. Blipped-spiral trajectory is another powerful way to spread aliasing across three spatial axes (Zahneisen et al., 2014). Generalizing blipped-CAIPI to spiral imaging permits distributing MB acceleration across slice, as well as between k_x and k_y axes. Further, it is possible to rotate the interleaves in the k_x - k_y plane across different k_z partitions during 3D stack-of-spirals acquisitions, thus exploiting another degree of freedom to introduce complementary sampling for rapid imaging (Deng et al., 2016).

9.4 Model-based reconstruction for parallel imaging

Image reconstruction in PI involves inverting the forward system model expressed in Eq. 9.2. For relatively limited R , traditional PI methods such as SENSE or GRAPPA can offer satisfactory reconstruction performance. However, the inverse problem becomes heavily ill-conditioned due to the reduced number of measurements at higher R (Pruessmann et al., 1999). This results in heavy residual artifacts

and/or noise amplification in reconstructed images. To improve conditioning, prior knowledge on the distribution of MRI data can be incorporated to constrain the solution set of the reconstruction problem. This is commonly achieved by augmenting the inverse problem based on the system model with a regularization term. For image-domain reconstructions such as SENSE, the modified optimization problem can be expressed as

$$\boldsymbol{\rho}^* = \operatorname{argmin}_{\boldsymbol{\rho}} \sum_i \|A_i \boldsymbol{\rho} - y_i\|_2^2 + \lambda_r R(\boldsymbol{\rho}), \quad (9.13)$$

where $\boldsymbol{\rho}$ is the MR image, y_i are k -space data for the i th coil, $A_i = \text{DFC}_i$ is the system matrix component for the i th coil, $R(\cdot)$ is the regularization function, and λ_r is the regularization weight. For k -space reconstructions such as SPIRiT, the regularized reconstruction is instead expressed as follows:

$$s^* = \operatorname{argmin}_s \sum_i \|Ds_i - y_i\|_2^2 + \lambda_1 \|(G - I)s\|_2^2 + \lambda_r R(s), \quad (9.14)$$

where y_i are acquired k -space data for the i th coil and s is the aggregated k -space data across all coils. Differing in terms of their assumptions regarding the data distribution, main-stream regularization functions used in parallel MRI reconstruction include smoothness (Lustig and Pauly, 2010), sparsity (Murphy et al., 2012), low-rank (Haldar and Zhuo, 2016), and recently deep priors (Hammernik et al., 2017). These fundamental approaches to regularized reconstruction are detailed below.

9.4.1 Smoothness priors

The spectrum of MR images rapidly decay from low- toward high-spatial frequencies in k -space. A corollary is that tissue signals show relatively gradual spatial variation in MR images. Furthermore, bodily organs typically contain tissue blocks with relatively uniform signal levels. As such, a traditional approach to separate tissue signals from white noise in MRI has been incorporation of smoothness priors either in image domain or k -space (Lustig et al., 2007). Common regularization functions include total-variation (TV) norm (i.e., l1-norm of the gradient) and l2-norm:

$$R(\rho) = \|\nabla \rho\|_1 \text{ or } R(s) = \|s\|_2^2. \quad (9.15)$$

Smoothness priors allow separation of additive white noise from gradually varying tissue signals in MR images. However, excessive regularization leads to spatial blurring or block artifacts in reconstructions.

9.4.2 Sparsity priors

MR images are considered to have sparse representations in known linear transform domains such as wavelet transform, where they can be represented with much fewer coefficients than would be required in image domain (Murphy et al., 2012). In turn, compressed sensing theory dictates that it should be possible to recover MR images from randomly undersampled acquisitions by enforcing a sparsity prior (Lustig et al., 2007). Although theoretically motivated, random undersampling in k -space yields low SNR efficiency due to the inhomogeneous distribution of energy across the spectrum of MR images. In practice, random undersampling with a variable density across k -space to mimic the image spectrum enhances measurement efficiency while still generating spatially incoherent artifacts in the image domain. Afterward, the sparsity prior can be incorporated to the reconstructions via an l1-norm term:

$$R(\rho) = \|\Phi\{\rho\}\|_1 \text{ or } R(s) = \|\Phi\{F^{-1}s\}\|_1, \quad (9.16)$$

where Φ stands for the transform domain where the images have sparse representations, e.g., wavelet transform. Sparsity priors coupled with random sampling can allow separation of incoherent aliasing artifacts from tissue signals. However, excessive regularization typically leads to loss of detailed, small image features in reconstructions.

9.4.3 Low-rank priors

Multicoil MR images reflect multiplicative modulation of the underlying MR image with sensitivities of individual coils. Multiplication via broad coil sensitivities in the image domain corresponds to convolution with relatively compact kernels in k -space. As such, k -space samples from multiple coils can be cast in the form of a matrix with block Hankel structure, inherently possessing low rank (Haldar and Zhuo, 2016; Shin et al., 2014). Recovery of missing k -space samples in undersampled acquisitions is then equivalent to recovering missing entries in the matrix via structured low-rank matrix completion. This completion can be achieved by enforcing the matrix to have low rank. Thus, low-rank priors can be incorporated via a regularization term

$$R(s) = \sum_i \sigma_i^2(H(s)) \text{ such that } \text{rank}(H) \leq r, \quad (9.17)$$

where H is the structured low-rank matrix, σ_i is the i th singular value of H , and r denotes the upper bound for matrix rank. Structured low-rank matrix completion has dualities to smoothness or sparsity priors, so prescribing excessively low r values during recovery can cause loss of image features during reconstruction.

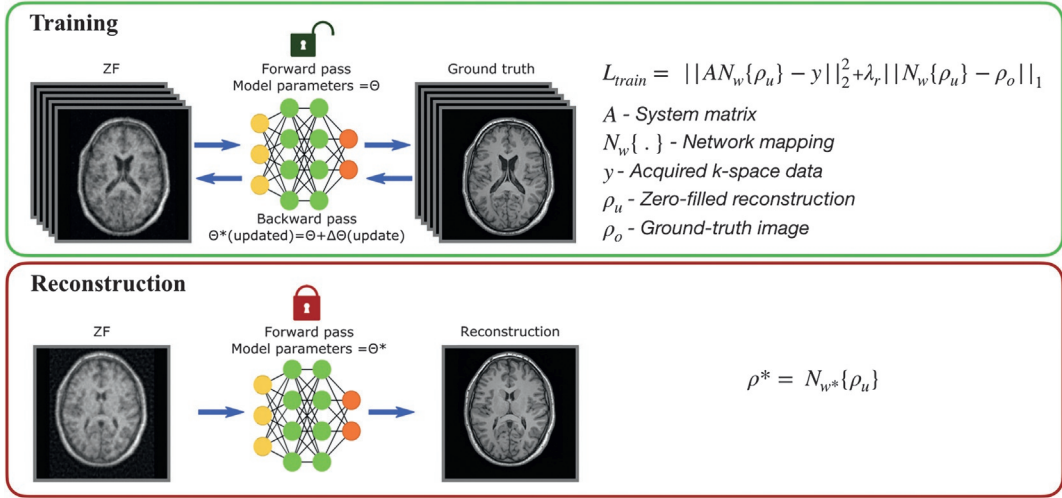
9.4.4 Deep priors

Smoothness, sparsity, or low-rank priors commonly involve hand-constructed regularization terms that rely on certain assumptions regarding the MR image distribution. When these assumptions diverge from the properties of the actual data distribution, they can introduce undesirable reconstruction biases and suboptimal performance. To surmount this difficult challenge, recent studies have adopted deep learning models for constructing priors for MRI reconstruction. Given a training set of undersampled and corresponding fully sampled acquisitions, a neural network can learn an indirect prior to suppress aliasing artifacts encountered in examples of training data (Dar et al., 2020). To do this, the network can be trained via solving a regularized optimization problem

$$w^* = \operatorname{argmin}_w \sum_i \|A_i N_w\{\rho_u\} - y_i\|_2^2 + \lambda_r \|N_w\{\rho_u\} - \rho_o\|_1, \quad (9.18)$$

where N_w is the network mapping with parameters w , ρ_u are MR images obtained via Fourier reconstruction of undersampled acquisitions, and ρ_o are ground-truth MR images based on fully sampled acquisitions. The data-consistency term based on the forward system model can also be incorporated into the network architecture via unrolling based on estimated coil sensitivities (Hammernik et al., 2017), or end-to-end estimation of coil sensitivity and the MR image can be performed (Shin et al., 2014). Afterward, the trained network can unalias Fourier reconstructions of undersampled acquisitions (Fig. 9.5).

Alternatively, a network can learn a direct prior that captures the distribution of MR data to constrain the set of reconstructions to high-quality images. The high-quality MRI prior can be trained using either fully sampled auto-calibration data in undersampled acquisitions, or else trained using MR images derived from fully sampled acquisitions (Korkmaz et al., 2022). Reconstructions based on direct


FIG. 9.5

Given a training set of undersampled and corresponding fully sampled acquisitions, a neural network can learn an indirect prior to suppress aliasing artifacts encountered in examples of training data. To do this, the network can be trained via solving a regularized optimization problem that weights a data-consistency loss based on the forward system model against a pixel-wise loss between the network output and the ground truth image. Afterward, the trained network can de-alias Fourier reconstructions of undersampled acquisitions.

MRI priors involve an iterative inference procedure to minimize the data-consistency loss on acquired k -space samples:

$$w^*, \rho^* = \operatorname{argmin}_w \sum_i \|A_i N_w\{\rho_u\} - y_i\|_2^2, \text{ where } \rho^* = N_{w^*}\{\rho_u\}. \quad (9.19)$$

Building a network model requires prolonged training procedures on a large array of training data. Yet, reconstructions based on deep priors typically offer higher performance than those based on hand-constructed priors as they can learn from and adapt to data flexibly, and they often offer faster inference.

9.5 Estimation of coil sensitivities

Coil sensitivity profiles determine the spatial encoding capabilities introduced by coil arrays over gradient-based encoding. Accordingly, coil sensitivities must be known to set up the forward system model in image-domain methods such as SENSE. While coil sensitivities are not explicitly derived in k -space methods, they are still implicitly embedded in k -space interpolation kernels due to the duality between image and Fourier domains. Therefore, regardless of the reconstruction domain, PI reconstructions must be informed regarding coil sensitivities to leverage coil-driven spatial encoding. In this section, we overview common approaches for the estimation of coil sensitivities including external calibration methods, auto-calibration methods, joint sensitivity estimation and reconstruction methods, and calibrationless methods.

9.5.1 External calibration

A straightforward approach for coil sensitivity estimation is to perform external calibration measurements (Pruessmann et al., 1999). This method requires access to a body coil that offers a homogeneous sensitivity profile over the imaged volume. UHF systems often lack a body coil. For such magnets, an alternative is to operate the coil array of interest in birdcage mode to attain relatively uniform sensitivity across the field-of-view. Images collected with spatially homogeneous sensitivity are taken as a reference, and coil sensitivity profiles can then be estimated by normalizing images collected with the coil array by the reference image (Fig. 9.6A)

$$C_i = \rho_i / \rho_{ref}, \quad (9.20)$$

where ρ_{ref} denotes the reference image and ρ_i is the image for the i th element in the coil array. External calibration is a powerful approach that allows estimation of absolute sensitivity information. However, inadvertent movement in between measurements with the body coil and coil array, or any drifts in coil sensitivities between the calibration measurements and the actual scans, can cause errors in sensitivity estimates.

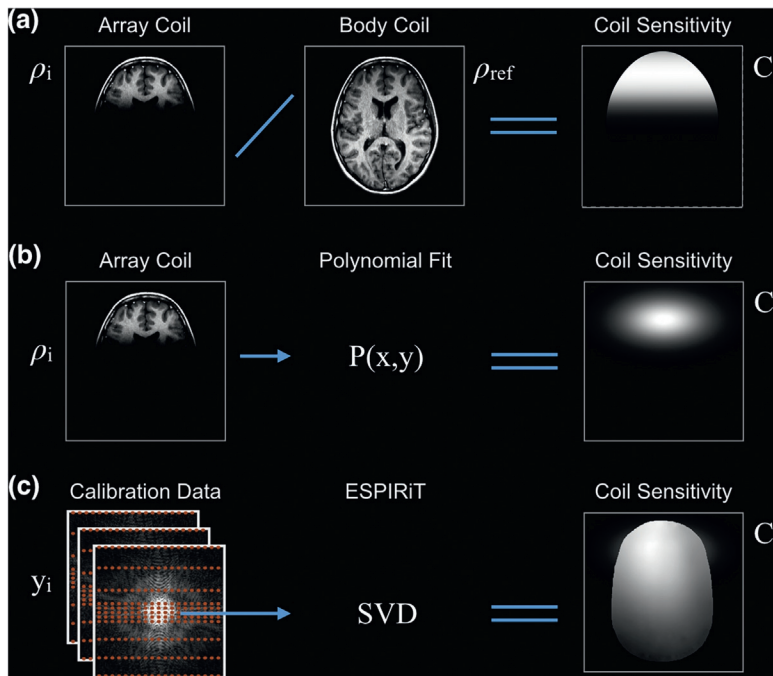


FIG. 9.6

Common methods for coil sensitivity estimation. (A) External calibration method. Images acquired with the coil array of interest are normalized by a reference image acquired separately via a body coil with homogeneous sensitivity. (B) Assuming that coil sensitivity should show gradual spatial variation, a low-order polynomial is fit to individual coil images. (C) Calibration data aggregated across coils are subjected to a singular value decomposition to estimate coil sensitivity profiles.

Indeed, the acquisition of high-quality external calibration data plays an important role for robust echo planar imaging (EPI) reconstruction. For high-resolution whole-brain EPI data often acquired at UHF, reference scans can be lengthy (10 s or longer), which increases the vulnerability to subject motion as well as B_0 changes due to respiration. Since B_0 fluctuations scale with the field strength, it is crucial to perform calibration scans robust to physiologic variations. Standard calibration scans rely on multishot EPI readouts, which obtain “fully-sampled” reference data when their k -space is averaged across multiple shots. These standard reference scans run through all the slices to cover the imaging FOV for the first shot, then proceed to acquire all the slices for the second shot. As such, when the multishot data are combined, each slice receives contributions from several seconds of scan time. An alternative, motion-robust calibration approach is the fast low-angle excitation echo-planar technique (FLEET) method (Polimeni et al., 2016). In FLEET, all shots belonging to slice 1 are acquired first, then all shots for the next slice position are sampled. By this way, each slice’s k -space data are subject to a much shorter acquisition time frame of 100–200 ms, thereby significantly boosting the robustness of such external calibration scans.

9.5.2 Auto-calibration

To mitigate the need to perform separate calibration measurements, coil sensitivities can instead be estimated based on the very data subject to PI reconstruction. For image-domain methods that require explicit sensitivity estimates, polynomial fitting can be used. Coil sensitivity profiles reflect the receive B_1 field expected to vary gradually across space, so they predominantly contain low-spatial-frequency components. In turn, a limited degree polynomial can be fit to individual coil images to capture gradual intensity modulations due to coil sensitivities (Fig. 9.6B)

$$P(x, y) = \sum_{k=0}^n \alpha_k \cdot x^q \cdot y^r, \text{ such that } (q + r) \leq k, \quad (9.21)$$

where P denotes a polynomial over two spatial dimensions (x, y) of degree n and α_k denotes the scaling factor of the k th term in the polynomial fit. Since polynomial fitting can be cast as a linearized regression problem, the solution can be obtained via least-squares minimization (Pruessmann et al., 1999)

$$C_i = \operatorname{argmin}_x \| P(x, y) - \rho_i(x, y) \|_2^2. \quad (9.22)$$

An alternative approach is the ESPIRiT method that estimates coil sensitivity via an eigenvalue problem expressed using fully sampled calibration data at the center of k -space (Uecker et al., 2014). Aggregating auto-calibration data across coils into a calibration matrix, ESPIRiT observes that this matrix has a nonempty null space and calibration data that lie entirely in the row space of the calibration matrix. Since these fundamental observations should also hold for other k -space regions outside the auto-calibration region, aggregated k -space data in remaining regions should be consistent with calibration data. Assuming that linear relationships between k -space patches across coils are characterized with a reconstruction operator W , consistency to calibration data can be expressed as

$$Ws = s, \quad (9.23)$$

where s denotes k -space data. Substituting the forward system model in Eq. 9.23, we get

$$WFC\rho = FC\rho, \quad (9.24A)$$

$$(F^{-1}WF)C\rho = C\rho. \quad (9.24B)$$

It follows from Eq. 9.24B that coil sensitivities are an eigenvector of $(F^{-1}WF)$ and thereby W with an eigenvalue of 1. Thus, coil sensitivities that form a basis for the row space can be derived based on eigendecomposition of W (Fig. 9.6C).

9.5.3 Joint sensitivity estimation and reconstruction

Both external and auto-calibration methods perform a priori estimation of coil sensitivities, and the derived estimates are then used as fixed variables in the forward system model during PI reconstruction. Therefore, any estimation errors for coil sensitivities will translate into reconstruction errors due to inaccuracies in the forward system model. To help mitigate errors in sensitivity estimation, a joint optimization problem can be solved instead that simultaneously estimates coil sensitivities along with the reconstructed image. For image-domain methods, this can be achieved by modifying the problem formulation in Eq. 9.13 as

$$\rho^*, C^* = \operatorname{argmin}_{\rho, C} \sum_i \|\operatorname{DFC}_i \rho - y_i\|_2^2 + \lambda_r R(\rho). \quad (9.25)$$

Meanwhile, for k -space methods, the modified problem formulation is

$$s^*, G^* = \operatorname{argmin}_{s, G} \sum_i \|Ds_i - y_i\|_2^2 + \lambda_1 \|(G - I)s\|_2^2 + \lambda_r R(s), \quad (9.26)$$

where G is the k -space interpolation operator that is to be estimated. Optimization over two distinct sets of variables can be performed via alternating minimization approaches.

9.5.4 Calibrationless reconstruction

Acquiring fully sampled k -space lines is common in PI applications as it enables either explicit estimation of coil sensitivities or implicit sensitivity estimation in the form of interpolation kernels in k -space. However, in certain cases, it may be impractical to perform Nyquist sampling within a dedicated k -space region such as time-interleaved k -space sampling, or imaging at very high acceleration rates. In such cases, performing PI reconstruction in the absence of calibration data can be attempted via calibrationless approaches such as P-LORAKS or SAKE (Haldar and Zhuo, 2016; Shin et al., 2014). Observing that k -space samples from multiple coils should form a system matrix with block Hankel structure, reconstruction is formulated as a structured low-rank matrix completion problem. In this framework, enforcing low rank in the Hankel matrix implicitly enables estimation of missing k -space entries in the matrix, without the need for explicit estimates of coil sensitivities or interpolation kernels. For instance, a k -space reconstruction can be implemented as

$$s^* = \operatorname{argmin}_s \sum_i \|Ds_i - y_i\|_2^2 + \lambda_r \sum_n \sigma_n^2(H(s)) \text{ such that } \operatorname{rank}(H) \leq r, \quad (9.27)$$

where H is the structured low-rank matrix, σ_n is the n th singular value of H , and r denotes the upper bound for matrix rank.

9.6 Coil sensitivity profiles vary more rapidly across space at UHF, thereby improving g-factor performance

As coil sensitivity profiles become more localized and spatially varying at higher field strengths, it is expected that similar coil geometries will have better g-factor performance at UHF. This expectation

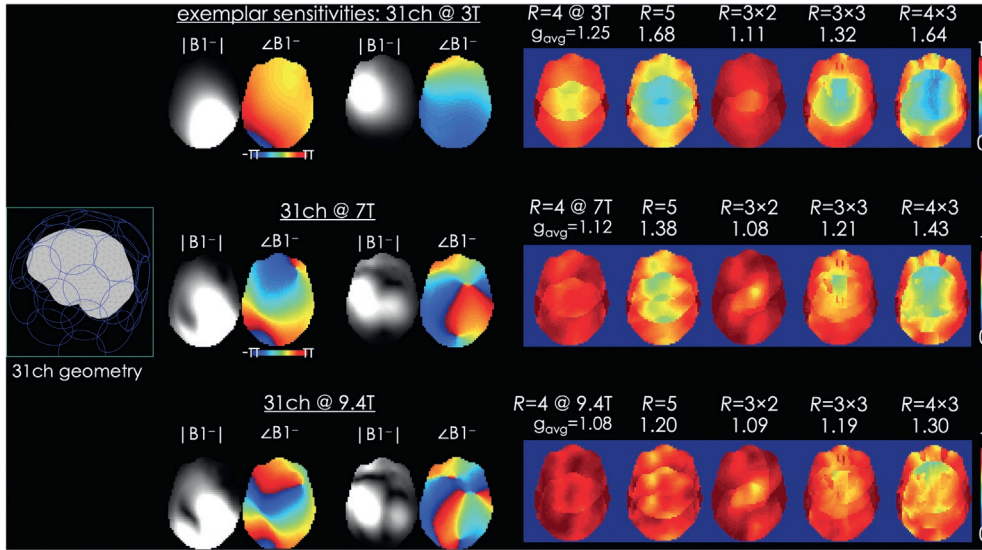


FIG 9.7

(Top) Exemplar magnitude and phase coil sensitivities for two channels out of a 31ch numerical simulation are depicted. $1/g$ -factor maps report results from 1-dimensional undersampling (at $R=4$ and 5) as well as 2-dimensional CAIPI reconstructions (at $R=3 \times 2$, 3×3 and 4×3). (Middle) Going to 7T, coil sensitivity profiles become more spatially varying, and “singularities/fringe lines” become visible in the phase of the sensitivity maps. G-factor performance improves over 3T. (Bottom) At 9.4T, coil profiles are yet more rapidly varying and this leads to further gains in g-factor. For instance, $R=4 \times 3$ acceleration at 9.4T has slightly better performance as the $R=3 \times 3$ case at 3T. This is also true for $R=5$ at 9.4T, compared to the reconstruction at $R=4$ -fold acceleration at 3T.

has been demonstrated through numerical simulations (Wiesinger et al., 2004a) and experimental data (Wiesinger et al., 2004b). We have used the MARIE electromagnetic simulation toolbox (Villena et al., 2016) to explore the field strength dependence of g-factor performance across 3T, 7T, and 9.4T using the same 31ch coil geometry in Fig. 9.7. Noise covariance was taken to be identity ($\Psi=I$). As demonstrated in the exemplar profiles, both phase and magnitude of coil sensitivities become more spatially varying toward higher fields. This is reflected in the g-factor performance, where SENSE at UHF clearly outperformed 3T reconstructions. For example, the average g-factor at 9.4T using $R=4 \times 3$ acceleration was better than that of the 3T g-factor at the lower $R=3 \times 3$ -fold rate. The ability to support higher undersampling factors is complemented by the SNR increase at higher field strengths, which should more than compensate for the increased \sqrt{R} intrinsic SNR penalty and render such acceleration factors feasible at submillimeter resolutions. An active research field is the development of very dense coil arrays at UHF. The construction of 64ch (Uğurbil et al., 2019) and 128ch head arrays (Gruber et al., 2021) increased the PI capability well beyond the commercially available 32ch UHF coils, whereas simulations of 256ch designs helped underline the potential of massively parallel arrays in pushing the acceleration factors even further (Hendriks et al., 2019).

9.7 Exemplar applications enabled by the increased encoding power of UHF systems

Going to higher fields provides two complementary benefits for efficient image encoding. As detailed in the previous section, coil sensitivity profiles are more localized at UHF, thus permitting higher acceleration factors to be achieved. Additionally, the SNR boost provided by higher field strengths renders such high acceleration factors practical. For instance, wave-CAIPI readily permits $R = 3 \times 3$ -fold acceleration with negligible g-factor penalties (Polak et al., 2018), yet the impact of the intrinsic \sqrt{R} penalty on SNR causes the images to be relatively noisy for routine practical use at 3T. As such, with the advent of advanced parallel imaging strategies, we have become SNR limited, rather than encoding limited, especially for high-resolution (~ 1 mm isotropic) data at 3T or lower field strengths. While regularization techniques can help boost the SNR, going to UHF directly makes such acceleration factors practical and impactful. An example to this is the wave-encoded MPRAGE acquisition at 1 mm isotropic resolution in Fig. 9.8A, which was completed in under a minute with near-perfect g-factor

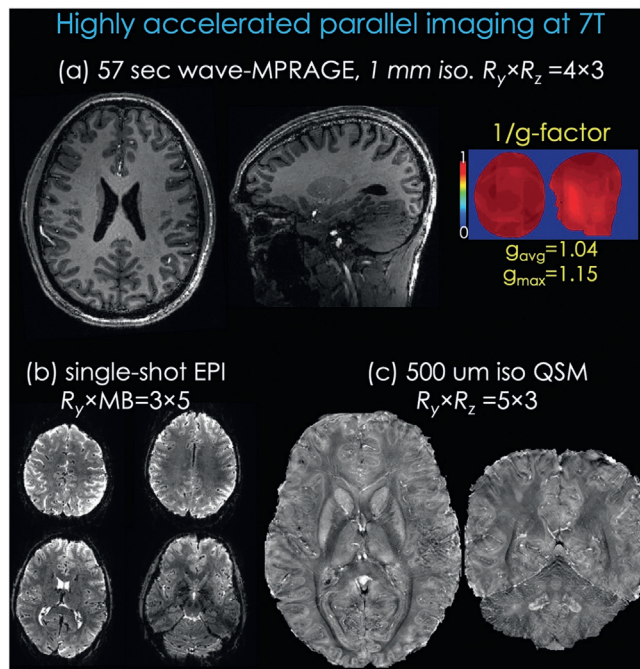


FIG. 9.8

(A) A 57-s, 1-mm isotropic MPRAGE acquisition with adequate SNR became possible at $R = 4 \times 3$ -fold acceleration using wave encoding by capitalizing on the SNR gain of 7T imaging. (B) Single-shot EPI with SMS encoding at $R_y \times MB = 3 \times 5$ -fold acceleration yielded high-quality gradient echo data for functional imaging studies. (C) The superlinear boost of UHF gradient echo imaging permits trading off SNR for resolution, in this case permitting a $500\mu\text{m}$ isotropic QSM acquisition at $R = 5 \times 3$ -fold acceleration using wave encoding. All acquisitions were made using a 31ch head receiver array.

performance. Such a rapid acquisition at $R = 4 \times 3$ -fold acceleration would have been too noisy at 3T when using a standard wave-CAIPI reconstruction.

There is a superlinear relation between the SNR and the field strength for gradient echo (GRE) imaging, e.g., going from 3T to 7T provides a ~ 3.1 -fold SNR boost instead of the expected ~ 2.3 -fold (Pohmann et al., 2016). This makes highly accelerated, high-resolution GRE acquisitions extremely favorable at UHF. Fig 9.8 provides two important example applications. In Fig. 9.8B, whole-brain single-shot EPI with high-quality reconstructions became possible at $R_y \times MB = 3 \times 5$ acceleration, thus providing high temporal resolution and high geometric fidelity simultaneously for functional imaging applications. In Fig. 9.8C, the superlinear SNR gain is traded off for spatial resolution in an $R_y \times R_z = 5 \times 3$ -fold accelerated wave-GRE acquisition at 500 μm isotropic resolution. This permitted quantitative susceptibility mapping (QSM) reconstruction from a ~ 20 min scan where three acquisitions were made with different head orientations relative to the main magnetic field. Combination of the mesoscale resolution and susceptibility contrast mechanism revealed exquisite contrast in the cortex and deep gray matter (axial) as well as in the cerebellum (coronal view).

Acknowledgments

We would like to thank Drs Jason Stockmann and Jonathan Polimeni for their insight and help with figures and coil sensitivity simulations across field strengths.

References

- Bilgic, B., Gagoski, B.A., Cauley, S.F., et al., 2015. Wave-CAIPI for highly accelerated 3D imaging. *Magn. Reson. Med.* 73, 2152–2162.
- Blaimer, M., Breuer, F.A., Mueller, M., et al., 2006. 2D-GRAPPA-operator for faster 3D parallel MRI. *Magn. Reson. Med.* 56, 1359–1364. <https://doi.org/10.1002/mrm.21071>.
- Breuer, F.A., Blaimer, M., Heidemann, R.M., Mueller, M.F., Griswold, M.A., Jakob, P.M., 2005. Controlled aliasing in parallel imaging results in higher acceleration (CAIPIRINHA) for multi-slice imaging. *Magn. Reson. Med.* 53, 684–691.
- Breuer, F.A., Blaimer, M., Mueller, M.F., et al., 2006. Controlled aliasing in volumetric parallel imaging (2D CAIPIRINHA). *Magn. Reson. Med.* 55, 549–556.
- Breuer, F.A., Kannengiesser, S.A.R., Blaimer, M., Seiberlich, N., Jakob, P.M., Griswold, M.A., 2009. General formulation for quantitative G-factor calculation in GRAPPA reconstructions. *Magn. Reson. Med.* 62, 739–746. <https://doi.org/10.1002/mrm.22066>.
- Breuer, F.A., Moriguchi, H., Seiberlich, N., et al., 2008. Zigzag sampling for improved parallel imaging. *Magn. Reson. Med.* 60, 474–478.
- Dar, S.U.H., Özbey, M., Çatlı, A.B., Çukur, T., 2020. A transfer-learning approach for accelerated MRI using deep neural networks. *Magn. Reson. Med.* 84, 663–685.
- Deng, W., Zahneisen, B., Andrew, S.V., 2016. Rotated stack-of-spirals partial acquisition for rapid volumetric parallel MRI. *Magn. Reson. Med.* 76, 127–135. <https://doi.org/10.1002/mrm.25863>.
- Feinberg, D.A., Moeller, S., Smith, S.M., et al., 2011. Correction: multiplexed echo planar imaging for sub-second whole brain fMRI and fast diffusion imaging. *PLoS One* 6. <https://doi.org/10.1371/annotation/d9496d01-8c5d-4d24-8287-94449ada5064>.

- Griswold, M.A., Jakob, P.M., Heidemann, R.M., et al., 2002. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn. Reson. Med.* 47, 1202–1210.
- Gruber, B., Stockmann, J., Mareyam, A., et al., 2021. A 128-channel head coil array for cortical imaging at 7 tesla. In: *Proceedings of the 2021 ISMRM & SMRT Virtual Conference & Exhibition*, p. 176.
- Haldar, J.P., Zhuo, J., 2016. P-LORAKS: low-rank modeling of local k-space neighborhoods with parallel imaging data. *Magn. Reson. Med.* 2016 (75), 1499.
- Hammernik, K., Klatzer, T., Kobler, E., et al., 2017. Learning a Variational network for reconstruction of accelerated MRI data. *Magn. Reson. Med.* 79, 3055–3071.
- Hendriks, A.D., Luijten, P.R., Klomp, D.W.J., Petridou, N., 2019. Potential acceleration performance of a 256-channel whole-brain receive array at 7 T. *Magn. Reson. Med.* 81, 1659–1670.
- Korkmaz, Y., Dar, S.U.H., Yurt, M., Özbey, M., Çukur, T., 2022. Unsupervised MRI reconstruction via zero-shot learned adversarial transformers. *IEEE Trans. Med. Imaging* 41, 1747–1763.
- Larkman, D.J., Hajnal, J.V., Herlihy, A.H., Coutts, G.A., Young, I.R., Ehnholm, G., 2001. Use of multicoil arrays for separation of signal from multiple slices simultaneously excited. *J. Magn. Reson. Imaging* 13, 313–317.
- Lustig, M., Donoho, D., Pauly, J.M., 2007. Sparse MRI: the application of compressed sensing for rapid MR imaging. *Magn. Reson. Med.* 58, 1182–1195.
- Lustig, M., Pauly, J.M., 2010. SPIRiT: iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magn. Reson. Med.* 64, 457–471.
- Moeller, S., Yacoub, E., Olman, C.A., et al., 2010. Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. *Magn. Reson. Med.* 63, 1144–1153. <https://doi.org/10.1002/mrm.22361>.
- Moriguchi, H., Duerk, J.L., 2006. Bunched phase encoding (BPE): a new fast data acquisition method in MRI. *Magn. Reson. Med.* 55, 633–648.
- Murphy, M., Alley, M., Demmel, J., Keutzer, K., Vasanaawala, S., Lustig, M., 2012. Fast 11-SPIRiT compressed sensing parallel imaging MRI: scalable parallel implementation and clinically feasible runtime. *IEEE Trans. Med. Imaging* 31, 1250–1262.
- Nunes, R.G., Hajnal, J.V., Golay, X., Larkman, D.J., 2006. Simultaneous slice excitation and reconstruction for single shot EPI. *Proc. Int. Soc. Magn. Reson. Med.* 14, 293.
- Pohmann, R., Speck, O., Scheffler, K., 2016. Signal-to-noise ratio and MR tissue parameters in human brain imaging at 3, 7, and 9.4 tesla using current receive coil arrays. *Magn. Reson. Med.* 75, 801–809.
- Polak, D., Setsompop, K., Cauley, S.F., et al., 2018. Wave-CAIPI for highly accelerated MP-RAGE imaging. *Magn. Reson. Med.* 79, 401–406.
- Polimeni, J.R., Bhat, H., Witzel, T., et al., 2016. Reducing sensitivity losses due to respiration and motion in accelerated echo planar imaging by reordering the autocalibration data acquisition. *Magn. Reson. Med.* 75 (2), 665–679.
- Pruessmann, K.P., Weiger, M., Scheidegger, M.B., Boesiger, P., 1999. SENSE: sensitivity encoding for fast MRI. *Magn. Reson. Med.* 42, 952–962.
- Robson, P.M., Grant, A.K., Madhuranthakam, A.J., Lattanzi, R., Sodickson, D.K., McKenzie, C.A., 2008. Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-based parallel imaging reconstructions. *Magn. Reson. Med.* 60, 895–907.
- Setsompop, K., Gagoski, B.A., Polimeni, J.R., Witzel, T., Wedeen, V.J., Wald, L.L., 2012. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magn. Reson. Med.* 67, 1210–1224.
- Setsompop, K., Kimmlingen, R., Eberlein, E., et al., 2013. Pushing the limits of in vivo diffusion MRI for the human connectome project. *NeuroImage* 80, 220–233.
- Shin, P.J., Larson, P.E.Z., Ohliger, M.A., et al., 2014. Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion. *Magn. Reson. Med.* 72, 959–970.

- Sodickson, D.K., Manning, W.J., 1997. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magn. Reson. Med.* 38, 591–603. <https://doi.org/10.1002/mrm.1910380414>.
- Sotiropoulos, S.N., Jbabdi, S., Xu, J., et al., 2013. Advances in diffusion MRI acquisition and processing in the human connectome project. *NeuroImage* 80, 125–143.
- Uecker, M., Lai, P., Murphy, M.J., et al., 2014. ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. *Magn. Reson. Med.* 71, 990–1001.
- Uğurbil, K., Auerbach, E., Moeller, S., et al., 2019. Brain imaging with improved acceleration and SNR at 7 tesla obtained with 64-channel receive array. *Magn. Reson. Med.* 82, 495–509.
- Villena, J.F., Polimeridis, A.G., Eryaman, Y., et al., 2016. Fast electromagnetic analysis of MRI transmit RF coils based on accelerated integral equation methods. *IEEE Trans. Biomed. Eng.* 63, 2250–2261.
- Weiger, M., Pruessmann, K.P., Boesiger, P., 2002. 2D SENSE for faster 3D MRI. *MAGMA* 14, 10–19.
- Wiesinger, F., Boesiger, P., Pruessmann, K.P., 2004a. Electrodynamics and ultimate SNR in parallel MR imaging. *Magn. Reson. Med.* 52, 376–390.
- Wiesinger, F., Van de Moortele, P.-F., Adriany, G., De Zanche, N., Ugurbil, K., Pruessmann, K.P., 2004b. Parallel imaging performance as a function of field strength—an experimental investigation using electrodynamic scaling. *Magn. Reson. Med.* 52, 953–964.
- Yutzy, S.R., Seiberlich, N., Duerk, J.L., Griswold, M.A., 2011. Improvements in multislice parallel imaging using radial CAIPIRINHA. *Magn. Reson. Med.* 65, 1630–1637.
- Zahneisen, B., Ernst, T., Poser, B.A., 2015. SENSE and simultaneous multislice imaging. *Magn. Reson. Med.* 74, 1356–1362.
- Zahneisen, B., Poser, B.A., Ernst, T., Stenger, A.V., 2014. Simultaneous Multi-Slice fMRI using spiral trajectories. *NeuroImage* 92, 8–18.