

Dependence of parallel imaging with linear predictability on the undersampling direction

Alex McManus^a, Stephen Becker^a, and Nicholas Dwork^{b,*}

^aUniversity of Colorado Boulder, Department of Applied Mathematics, Boulder, Colorado, United States

^bUniversity of Colorado Anschutz, Department of Biomedical Informatics, Aurora, Colorado, United States

ABSTRACT. Parallel imaging with linear predictability takes advantage of information present in multiple receive coils to accurately reconstruct the image with fewer samples. Commonly used algorithms based on linear predictability include GRAPPA and SPIRiT. We present a sufficient condition for an accurate reconstruction based on the direction of undersampling and the arrangement of the sensing coils. We show, with examples, that the quality of the reconstruction can be high or low for the same data based on whether this condition is met or not met, respectively. We also propose a metric—the acceleration direction metric (ADM)—that uses a fully sampled region centered on the 0 frequency to identify which direction(s) of undersampling would allow for a good-quality image reconstruction prior to full data collection.

© 2025 SPIE and IS&T [DOI: [10.1117/1.JEI.34.2.023031](https://doi.org/10.1117/1.JEI.34.2.023031)]

Keywords: multi-coil magnetic resonance imaging; parallel imaging; linear predictability

Paper 241373G received Nov. 19, 2024; revised Mar. 3, 2025; accepted Mar. 11, 2025; published Mar. 31, 2025.

1 Introduction

A highly effective and ubiquitously used method for accelerating magnetic resonance imaging (MRI) is parallel imaging,¹ where multiple coils simultaneously image the subject from disparate vantage points; the unique information collected with each coil permits an accurate interpolation of missing data in the frequency domain. There are two common approaches to parallel imaging: model-based reconstruction (MBR),^{2,3} where the MRI system is simulated for reconstruction, and linear predictability,⁴ where it is assumed that each Fourier value is a linear combination of nearby Fourier values. MBR assumes knowledge of the coils' sensitivities and generates a single underlying image by solving an optimization problem; algorithms in this vein include SENSE,⁵ ESPIRiT,⁶ PISCO,⁷ and NLINV.⁸ Parallel imaging with linear predictability (PILP) can accelerate MRI without knowledge of the sensitivity maps; instead, PILP uses a single set of coefficients such that missing data can be estimated with a linear combination of nearby known sample values. Examples of PILP algorithms include SMASH,⁹ AUTO-SMASH,¹⁰ P-LORAKS,¹¹ GRAPPA,¹² and SPIRiT.¹³

Although PILP can be extremely effective, it may also fail to yield a high-quality image. Without a method for assuring success, clinicians are left to try the scan and see if the result is of diagnostic quality or not. When the image is poor quality, the clinician must re-scan the patient, which reduces the overall utilization of the MRI machine and may increase the time that a patient is anesthetized (if anesthesia was required for the scanning session). This leaves the clinician

*Address all correspondence to Nicholas Dwork, nicholas.dwork@cuanschutz.edu

weary of its usage. If PILP could be made robust to these errors, clinicians could use this form of accelerated MRI with confidence.

A common method for quantitatively evaluating and optimizing acceleration techniques is the geometry factor (g-factor).¹⁴ In Ref. 15, the sampling pattern for an MBR reconstruction is altered to optimize a function of the geometry factor (g-factor) and image acceleration. G-factor analysis, though, is a retrospective analysis. That is, one can use the g-factor to identify which portions of the image have a high signal-to-noise ratio (SNR) and which portions do not, but only after the data have been acquired. Thus, it cannot be used to ensure success with PILP. Furthermore, g-factor analysis confounds errors in the reconstruction from the sensing array with the signal received from the object to be imaged. Therefore, it does not provide an understanding of what acceleration is appropriate for a given sensing array.

PILP has two sets of parameters: (1) the amount of and direction(s) of undersampling and (2) the size of the interpolation kernel. In Ref. 4, Haldar et al. showed that a sufficient condition for accurate PILP is that the support of the imaged object (the subset of the domain where the image without noise is non-zero) is a strict subset of the field-of-view of the image. That is, if there are regions of the field-of-view that do not image the subject and just image air, then there exists a set of linear coefficients that can accurately interpolate missing data. However, this condition is only sufficient when the size of the interpolation kernel is unlimited. If the support is a large percentage of the field of view, then the kernel must be very large for this condition to be met. The common success of PILP suggests the existence of a less restrictive, sufficient condition. We will show cases where much of the image is air, and yet, the quality of the PILP reconstruction with a reasonably sized kernel is extremely poor.

It has been suggested that PILP yields accurate values for arbitrary displacements in the Fourier domain,¹ permitting an arbitrary acceleration. Presumably, any degradation would be due completely to the reduction in SNR that accompanies a reduced acquisition time.¹⁶ In this study, we will show that this is not the case. We will provide examples where PILP succeeds with one accelerated sampling pattern but fails when the sampling pattern is simply rotated. We will also show examples where the quality is impervious to the direction of undersampling and explain why this is the case. We will provide a novel sufficient condition for when PILP can achieve a high-quality reconstruction, and we will provide examples that show the reconstruction becomes low quality when this condition is not met. Finally, we will present a new method that automatically identifies the directions of undersampling that can be accelerated while maintaining a high-quality image.

2 Background

In this section, we review AUTO-SMASH¹⁰ and show that the generalization of its theoretical basis justifies GRAPPA,¹² SPIRiT,¹³ and general PILP. (A complete review of the theory of AUTO-SMASH using the notation of this paper can be found in Appendix A.) Unless otherwise specified, for the purposes of this discussion, we will assume that processing is performed in two dimensions. These sampling patterns may be generated with a three-dimensional spin-warp trajectory with two dimensions of phase encodes and one dimension of readout; after inverse Fourier transforming along the readout direction, the data are placed in a hybrid (k_x, k_y, z) space and the reconstruction of each slice may be performed independently of every other slice.¹⁷

2.1 MRI Signal Equation

The MRI signal with spin density $\rho: \mathbb{R}^2 \rightarrow [0, \infty)$ for coil j is¹⁶

$$\begin{aligned} S_j(k_x, k_y) &= \iint dx dy C_j(x, y) \rho(x, y) e^{-i(k_x x + k_y y)} \\ &= \mathcal{F}\{C_j \rho\}(k_x, k_y), \end{aligned} \quad (1)$$

where C_j is the sensitivity function for the j 'th coil, $k_x = (\gamma/2\pi) \int_0^{t_x} G_x(\tau) d\tau$, $k_y = (\gamma/2\pi) \int_0^{t_y} G_y(\tau) d\tau$, γ is the gyromagnetic ratio, G_x and G_y are the x and y gradient waveforms,




Fig. 1 In this depiction of parallel imaging with linear prediction, a portion of a sampling pattern is shown: filled-in circles represent a k -space location that was collected and unfilled circles represent a location that was not. A distance threshold of 1 and a metric of $\|\cdot\|_\infty$ are depicted with blue contours. The undersampling rate in each direction leads to different patterns of collected and uncollected points. The arrows represent displacements between collected data that lie within the threshold distance and an uncollected point that must be estimated.

t_x and t_y are the lengths of time that the respective gradient fields are turned on, and $\mathcal{F}\{C_j\rho\}(k_x, k_y)$ denotes the Fourier transform of $C_j\rho$ evaluated at (k_x, k_y) . For simplicity, we are ignoring the effects of relaxation and recovery.

The MRI machine samples the Fourier domain at a set of finite points according to a sampling pattern that is achieved through control of the gradient waveforms. An example sampling pattern for PILP is depicted in Fig. 1. In the example of this figure, the data are undersampled by a factor of 2 in both the horizontal and vertical directions (for a total undersampling factor of 4). Note that with J coils, each sampled point consists of J complex values, one for each coil.

The user supplies a metric and (at least one) threshold, perhaps one per direction. Points that were not collected are synthesized as linear combinations of those values that lie within the threshold. For any uncollected Fourier location, linear coefficients are found to interpolate from all collected points that lie within a threshold's distance. In the example of Fig. 1, the user has specified the $\|\cdot\|_\infty$ as the metric and a threshold of 1.

2.2 AUTO-SMASH

AUTO-SMASH is a seminal work of PILP; the algorithm uses a two-dimensional spin-warp trajectory¹⁸ (e.g., phase encode and readout in k_y and k_x dimensions, respectively), whereas the Nyquist sampling theorem would dictate that the phase encoding lines be separated by Δk_y , AUTO-SMASH separates lines by $M\Delta k_y$ for some integer $M > 1$. Consider the case where there are J coils, each with sensitivity map C_j . AUTO-SMASH defines a composite sensitivity map, C_0^{comp} , created by a linear combination of the individual coil sensitivity maps with linear coefficients $n^{(0)} \in \mathbb{C}^J$. That is, $C_0^{\text{comp}} = \sum_j n_j^{(0)} C_j$. AUTO-SMASH further assumes that for each $m = 1, \dots, M-1$; there exist unknown linear coefficients $n^{(m)} \in \mathbb{C}^J$ such that $\sum_j n_j^{(m)} C_j \approx C_0^{\text{comp}} \exp(im\Delta k_y y)$. Then, by the Fourier shift theorem (as detailed in Appendix A), $\sum_j n_j^{(m)} S_j(k_x, k_y) = \mathcal{F}\{C_0^{\text{comp}}\rho\}(k_x, k_y - m\Delta k_y)$.

This illustrates that the linear coefficients $n^{(m)}$ accurately interpolate unknown values from collected data located $m\Delta k_y$ distance away in the k_y direction. With the approximation specified, these same coefficients are valid across the entire Fourier domain. With SMASH⁹ and AUTO-SMASH,¹⁰ the coils were designed so that $C_0^{\text{comp}} \approx \kappa$, for some constant κ , across the field-of-view of the image. AUTO-SMASH automatically determines the linear coefficients $n^{(m)}$ from a set of lines centered on the 0 frequency [called the fully-sampled region (FSR)] by solving the linear system $\mathcal{S}n^{(m)} = s_{\text{fsr}}$, where s_{fsr} is comprised of the values of the FSR.

3 Methods

We will now show how a generalization of the AUTO-SMASH theoretical basis justifies the more general PILP.

3.1 GRAPPA

GRAPPA, developed heuristically, eliminates the assumption of known $n^{(0)}$ and the existence of an approximately constant composite sensitivity map. Moreover, instead of searching for coefficients that correspond to a single predetermined displacement, GRAPPA attempts to find linear coefficients for multiple displacements. The combination of distance threshold and undersampling rate gives rise to different patterns of collected and uncollected points used for interpolation, as depicted in Fig. 1. Each unique pattern of collected points surrounding an uncollected point is called a kernel. The blue contours in Fig. 1 show two different kernels.

For a given kernel, as with AUTO-SMASH, the linear interpolation coefficients can be determined by solving a linear system. The following is a least-squares problem that simultaneously identifies the interpolation coefficients for all coils

$$\underset{N}{\text{minimize}} \quad \|SN - s_{\text{fsr}}\|_2^2, \quad (2)$$

where S is comprised of the appropriate values from the FSR, N is the matrix of weights, and s_{fsr} is a matrix of Fourier values from the FSR, which has a rectangular size $n_x \times n_y$. For a specific uncollected Fourier location k , let D_k be the number of nonzero points of the relevant kernel. The matrix S will be of size $[\eta_x \eta_y \times JD_k]$, N is a matrix of size $[JD_k \times J]$, s_{fsr} is a matrix of size $[\eta_x \eta_y \times J]$, and η_x and η_y are the number of times that the kernel fits inside the auto-calibration region in the k_x and k_y directions, respectively.

We will now present a novel explicit relationship between Eq. (2) and an assumption similar to that of AUTO-SMASH. An equivalent form of Eq. (2) is

$$\underset{N}{\text{minimize}} \quad \sum_{\mathbf{k}} |\mathcal{S}^{(k)}N - s_{\text{fsr}}^{(k)}|^2, \quad (3)$$

where $k \in \mathbf{k}$ is an individual location inside the auto-calibration region, $\mathcal{S}^{(k)}$ is the row of S that corresponds to location k , and $s_{\text{fsr}}^{(k)}$ is the k 'th row of s_{fsr} . We recognize further that we can write $\mathcal{S}^{(k)}N$ as

$$\mathcal{S}^{(k)}N = \sum_{d \in \mathcal{K}_k} \sum_j S_j(k+d) n_j^d \quad (4)$$

where $S_j(k+d)$ is the signal collected at point $k+d$ from coil j , $n_j^{(d)}$ is the appropriate weight, the inner sum is over the coils, and the outer sum is over the collected sample points that lie within the kernel \mathcal{K}_k .

For a given displacement (d_x, d_y) , we assume

$$\sum_j n_j^{(d)} C_j(x, y) \approx e^{i(d_x x + d_y y)} C_{\ell}(x, y). \quad (5)$$

That is, we assume that there exists a linear combination of sensitivity maps such that the result is a specific sensitivity map multiplied by a complex exponential with frequency that corresponds to the displacement d . Then, Eq. (4) becomes

$$\begin{aligned} \mathcal{S}^{(k)}N &= \sum_{d \in \mathcal{K}_k} \iint dx dy C_{\ell}(x, y) \rho(x, y) e^{-i((k_x - d_x)x + (k_y - d_y)y)} \\ &= \mathcal{F}\{C_{\ell}\rho\}(k_x + d_x, k_y + d_y). \end{aligned} \quad (6)$$

This reformulation shows that, with the assumption of Eq. (5), for a given displacement, GRAPPA seeks a set of linear coefficients $n^{(d)} \in \mathbb{C}^J$ such that the coil sensitivity maps approximate a complex exponential to best take advantage of the Fourier shift theorem. The notable difference is that the linear combination yields a complex exponential weighted by an individual coil's sensitivity rather than a composite sensitivity. Using analogous mathematics as presented

in Appendix A, performing a linear combination of the collected points from all coils linearly interpolates missing values of a specific coil's data. After all missing data are interpolated, the images from multiple coils can be combined into a single image.¹⁹

We can now present a novel sufficient condition for estimation without error when Eq. (5) is perfectly satisfied and without any noise. For a given location k and kernel, if at least one displacement vector to a collected point within the kernel satisfies Eq. (5), then the interpolation will be accurate. If more than one displacement vector satisfies Eq. (5), then GRAPPA finds the set of linear coefficients that interpolate from multiple points in a least-squares optimal sense.

3.2 SPIRiT

SPIRiT¹³ is an extension of GRAPPA. With a fixed kernel size, rather than just interpolating from points that were collected, SPIRiT will use *every point* in the kernel, regardless of whether or not it was collected.

SPIRiT interpolates all values (even those that were collected) by solving the following constrained least-squares problem

$$\text{minimize } \frac{1}{2} \|G\theta - \theta\|_2^2 \quad \text{subject to } \|D\theta - y\|_2^2 \leq \epsilon, \quad (7)$$

where $\|\cdot\|_2$ denotes the ℓ_2 norm, G represents linear interpolation from all values that lie within the kernel (even those that were not collected), D is the linear transformation that isolates the sample points that were collected, y is a vector of the values of the collected data, and ϵ is a bound on the noise power. Equation (7) can be solved with the Fast Iterative Shrinkage-Thresholding algorithm (FISTA),²⁰ as we discuss in Appendix B. In Ref. 13, Lustig et al. set $\epsilon = 0$ and only solve for the values of the uncollected data, which reduces the computational cost of the optimization, but this is not necessary. A solution that yields a smaller value of the objective function can be found when a non-zero value of ϵ is known and used.

Here, again, we present a novel sufficient condition for estimation without error when the approximation of Eq. (5) is perfectly satisfied and without any noise. Consider the set of collected data as a directed graph where the location of each Fourier value is a node and the displacement vectors from each point in the kernel to that location are the directed edges. Accurate interpolation at a location k is possible when there is a path from a collected data point to k such that all edges of that path satisfy Eq. (5). Practically, any error in the approximation and noise in the values are amplified with each edge of the path, so shorter paths lead to more accurate interpolations.

Figure 2 depicts an example that explains this sufficient condition. In this example, we would like to estimate the value of the blue circle. The directions where condition Eq. (5) is satisfied are down and downward-left. We could perfectly estimate the value of the blue circle in two steps: (1) interpolate all values below all sampled values and (2) interpolate all values downward-left of all known or estimated values. In a noiseless situation, if this process were

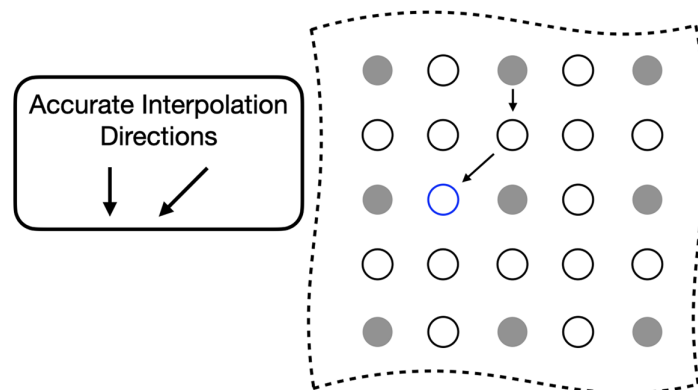


Fig. 2 Depiction of accurate interpolation with SPIRiT. For this hypothetical example, the directions where Eq. (5) is satisfied are down and downward-left. With GRAPPA, the blue circle cannot be interpolated accurately. With SPIRiT, the value of the blue circle can be estimated by two interpolations along the path indicated with the arrows.

iterated until all unknown values were estimated, then it would solve Eq. (7) with an objective function value of 0 (the minimum possible value).

3.3 Acceleration Direction Metric

In Sec. 5, we will provide examples where the quality of the reconstruction is significantly degraded when the sufficient conditions identified are not met. Here, we present a metric that can be calculated from the FSR that identifies any undersampling directions that do not meet Eq. (5). Given the FSR and a kernel size, we solve for the interpolation coefficients using two kernels—one solely in the horizontal direction and one in the vertical direction—of the same size as the kernel used for reconstruction. The kernel has a 0 in the center, and all other values are 1. For example, if the desired kernel is 3×3 , the test kernels look like $\mathbf{k}_h = [1 \ 0 \ 1]$ and $\mathbf{k}_v = \mathbf{k}_h^T$, where \cdot^T denotes transpose. We denote the solutions to Eq. (2) with each of these kernels as N_h^* and N_v^* , respectively. The acceleration direction metric (ADM) is the relative error of the linear system

$$\|SN^* - s_{\text{fsr}}\| / \|s_{\text{fsr}}\|. \quad (8)$$

If the value of ADM is high, then there is not a consistent set of interpolation coefficients for the FSR in the direction specified by the kernel. If its value is low, then there is a consistent set of interpolation coefficients and the data can be undersampled in that direction while retaining a high-quality reconstruction.

4 Experiments

In this paper, we first use Biot-Savart simulations to examine an eight-element birdcage coil (where we ignore coil coupling or high-frequency effects). We then analyze four different datasets: a knee, a brain, an ankle, and a shoulder. All datasets were fully sampled three-dimensional Cartesian data with two dimensions of phase encodes and one dimension of readout. The data on the knee were taken from mriData.²¹ The data of the brain, ankle, and shoulder were released.²² The data were retrospectively downsampled for processing. The data were inverse-transformed along the readout direction and placed into a hybrid space of (k_x, k_y, z) . Then, individual slices of specific z locations were isolated for further processing.

5 Results

Figure 3 shows results from Biot-Savart simulations for an eight-element birdcage coil according to Ref. 23. The top/bottom subfigures show simulations for an axial/sagittal plane that lies at the birdcage coil, respectively. The plots isolate single horizontal and vertical lines that lie at the center of the simulations; the number below each plot is the condition number of the matrix made by concatenating the vectors depicted in the plot. This condition number is a metric that indicates how much variation there is between the sensitivity maps across space. A high condition number implies that the solution to Eq. (2) will be heavily dependent on the noise and will not yield good results when the coefficients are used for interpolation. For both the horizontal and vertical lines of the axial simulations, the condition number is on the order of 10^3 . Although this remains the case for the horizontal line of the sagittal simulation, the condition number for the vertical line is much higher: on the order of 10^8 . By looking at the corresponding plot, it becomes obvious that the sensitivities are approximately scaled versions of each other and that the problem of finding coefficients to linearly combine the sensitivity maps so that they approximate a complex exponential is ill-conditioned.

Figure 4 shows reconstructions for an axial slice of a knee with GRAPPA¹² and SPIRiT¹³ using two separate sampling patterns. Both data were retrospectively undersampled at the same reduction factor of 2; the only difference is the direction of undersampling. This example shows how the quality can depend on the direction of undersampling, which is present with both GRAPPA and SPIRiT reconstructions: the quality of the reconstruction with horizontal (anterior-posterior) undersampling is high, whereas the quality of the reconstruction with vertical (superior-inferior) undersampling is low. Note that the higher values of the g-factor with vertical undersampling correspond with the lower image quality.

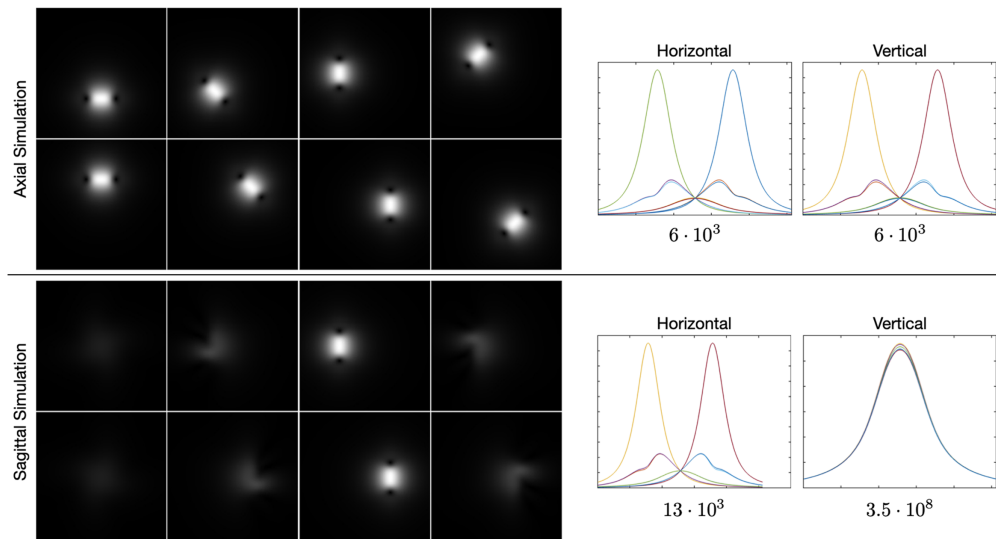


Fig. 3 Biot-Savart simulations of an eight-element birdcage coil. (The top/bottom simulation are the center axial/sagittal slices for a birdcage coil, respectively). The sensitivity maps of each coil are shown on the left. The plots on the right show the sensitivities of each coil for a single horizontal/vertical line through the center of the sensitivity maps. The numbers below each plot show the condition number of a matrix created by concatenating the sensitivities in the plots above.



Fig. 4 Reconstructions of sagittal slices of a knee with different sampling patterns. The columns from left to right are a representation of the sampling indicating the undersampling direction, the GRAPPA reconstruction, the SPIRiT reconstruction, and the g-factor of the SPIRiT reconstructions. The top and bottom rows show undersampling in the horizontal (anterior-posterior) and vertical (superior-inferior) directions, respectively. Both sampling masks used a reduction factor of 2. All reconstructions used a 31×31 FSR and a 3×3 kernel.

The sensitivity maps presented in Fig. 5 illuminate why this happens based on the understanding presented in Sec. 3—for any specific horizontal location (for any point along an anterior-posterior line), the coil sensitivities as a function of vertical location are approximately scaled versions of each other. When undersampled by every other row, GRAPPA and SPIRiT will attempt to interpolate unknown values from collected data that lie above and below in the Fourier domain. With relatively little spatial variation in the coil sensitivities in those directions, there does not exist a linear combination of coil sensitivities such that they approximate a complex exponential and the interpolation coefficients do not yield an accurate estimate.

These data were collected with an eight-channel birdcage coil. Therefore, each coil extends from the most inferior to the most superior portions of the image. Thus, in the superior-inferior




Fig. 5 Approximations of the sensitivity maps for each coil in an eight-coil birdcage for the data of Fig. 4. Note that there are only estimates of the sensitivity in pixels where the magnitude of the corresponding image is sufficiently high for an accurate estimate.




Fig. 6 SPIRiT reconstructions of a sagittal slice of a brain from data collected with an eight-channel birdcage coil using a horizontal (left) and vertical (right) undersampling mask. Top row: 3×3 kernel. Middle row: 5×5 kernel. Bottom row: 7×7 kernel. All reconstructions with an undersampling factor of 2 and a 31×31 FSR.

(SI) (vertical) direction, there is not enough variation to approximate a complex exponential well. This same phenomenon happens when imaging other anatomy with a similar coil arrangement; e.g., imaging the brain with a birdcage coil, as shown in Fig. 6.

Figure 7 shows magnetic resonance (MR) images of an axial slice of the knee; this is the same dataset as Fig. 4. In this case, the quality of the reconstruction is independent of the undersampling direction. Moreover, when simultaneously undersampling in both directions, the

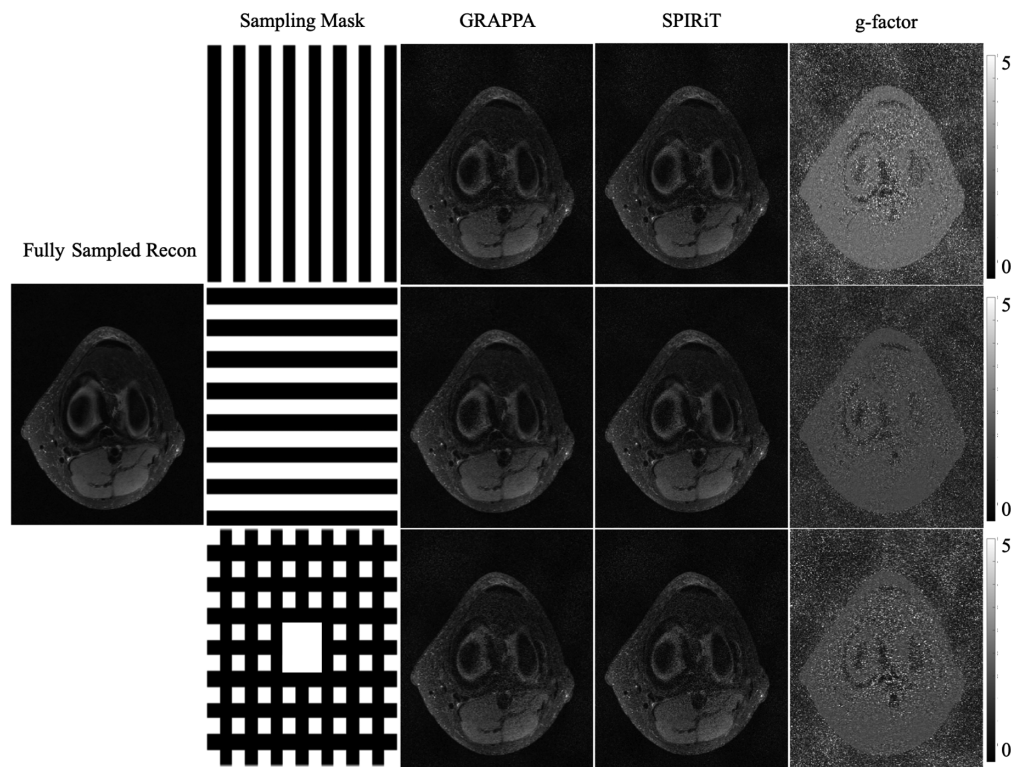


Fig. 7 Axial slices of the knee. The columns, from left to right, are the reconstruction with fully sampled data, representations of the sampling masks used, the GRAPPA reconstructions, the SPIRiT reconstructions, and the g-factors of the SPIRiT reconstructions. All images are size 320×256 . The top and middle rows have undersampling factors of two in a single direction (horizontal and vertical, respectively); the bottom row has a reduction factor of approximately two in both directions (for a total reduction factor of approximately four). All reconstructions from GRAPPA and SPIRiT were made with a 31×31 FSR and a 7×7 kernel.

quality remains high; though the signal-to-noise ratio has been reduced due to the reduction in scan time.^{16,24} Good quality reconstructions after undersampling in both directions (horizontal and vertical—or anterior/posterior and left/right) indicate that the corresponding optimization problems to find the interpolation coefficients for GRAPPA and SPIRiT were solved well. Owing to our understanding that the optimization problem is attempting to find linear coefficients for the coil sensitivities such that they linearly combine into a weighted complex exponential, the high quality indicates that this is true for at least one direction identified with the GRAPPA and SPIRiT kernels.

Figure 8 shows reconstructions of a sagittal slice of an ankle from data collected with an eight-channel dedicated ankle coil arrangement. The reconstructions of retrospectively down-sampled data result in a high-quality reconstruction independent of the direction of undersampling. Figure 9 shows the sensitivities of each coil. In contrast to the sensitivity maps of Fig. 5, the sensitivities exhibit variation in both directions.

We varied reconstruction parameters for GRAPPA and SPIRiT reconstructions of the knee and brain to include kernel sizes of 5×5 and 7×7 and to include an undersampling factor of ~ 3 . In all cases, the same trend was observed: the quality of the reconstruction is highly dependent on the undersampling direction (results are not shown).

Table 1 shows ADM values for the knee, brain, and ankle data studied in this paper. All results conform to the expectations that undersampling directions which yield a high ADM correspond to a poor quality image. The ADM values for those undersampling directions that yielded a poor-quality image are larger than those that yielded a high-quality image. The axial slice of the knee, which yielded a high-quality reconstruction independent of the undersampling direction, has a small ADM for both undersampling directions. For the data we analyzed, a threshold on the relative error of 0.4 would identify directions that would yield a poor-quality




Fig. 8 SPIRiT reconstructions of retrospectively downsampled ankle data. Top: horizontal undersampling pattern. Bottom: vertical undersampling pattern. Left: reduction factor of 2, 3×3 kernel. Center: reduction factor of 3, 3×3 kernel. Bottom: reduction factor of 3, 5×5 kernel. All reconstructions with a 31×31 FSR.




Fig. 9 Approximations to the individual sensitivity maps for the data of Fig. 8. Note that there are only estimates of the sensitivity in pixels where the magnitude of the corresponding image is sufficiently high for an accurate estimate.

image. This, however, is too small a dataset for us to make that a conclusion. Instead, we present this as a preliminary result and hope to pursue it in future work.

Figure 10 shows results for an axial slice of a shoulder from data collected with a 16-channel shoulder array. Note that because the coils cover the three-dimensional structure of the shoulder, many of the coils do not significantly sense the axial slice analyzed. The mean-squared error for the reconstructions with horizontal and vertical acceleration are $4.7 \cdot 10^{-11}$ and $5.3 \cdot 10^{-11}$, respectively. Due to the differences in spatial variance of the sensitivities, the quality of the reconstruction differs when accelerating horizontally or vertically.

Table 1 Examples of the ADM of Eq. (8) applied to specific cases. The metric values are large in undersampling directions that lead to poor-quality images and small in undersampling directions that lead to high-quality images.

Data	Kernel	ADM (vertical)	ADM (horizontal)
Knee—sagittal slice	3×3	55.0% (large)	17.1% (small)
Knee—axial slice	3×3	25.2% (small)	26.6% (small)
Ankle—sagittal slice	3×3	8.1% (small)	5.0% (small)
Brain—sagittal slice	3×3	61.7% (large)	5.0% (small)
Brain—sagittal slice	5×5	52.6% (large)	3.9% (small)
Brain—sagittal slice	7×7	47.3% (large)	3.6% (small)




Fig. 10 Results for data captured with a 16-coil shoulder array. (a) Estimates of the sensitivity maps. (b) The fully sampled reconstruction was compared with reconstructions with 3 \times acceleration in the horizontal and vertical directions. (c) The difference between the accelerated and the fully sampled reconstruction for the area outlined by the white box.

6 Conclusion

We have presented a rigorous physical reasoning that shows how the direction of undersampling can impact the reconstruction quality of PILP. Having the requisite spatial variation in the coil sensitivities is a sufficient condition for a high-quality reconstruction based on linear predictability. Another previously known sufficient condition is that the support of the imaged object be less than the field of view of the image.⁴ Note further that this condition of support is met by the data studied in this paper; however, the quality remains dependent on the direction of undersampling. This indicates that, for the data studied, the more important consideration is the direction in which the coil sensitivities can be linearly combined to approximate a complex exponential.

The linear combination of sensitivities only needs to approximate a complex exponential over the support of the image. Generally, with MRI, the support of the image is a strict subset of the field of view. Consider the reconstruction of the ankle; there is little coil sensitivity in the

superior-anterior quadrant of the image. However, there is plenty of coil variation over the other quadrants, which is where the ankle and leg are located. This is why the reconstruction of the ankle is robust regardless of the direction of undersampling.

The results presented indicate that the ADM can be used to determine undersampling direction(s) that will yield a poor-quality reconstruction so that those undersampling directions could be avoided during the scan. In the future, we hope to test this metric on a much larger dataset to ensure that it is reliable. We will analyze its sensitivity to field strength, patient movement, and size of the FSR. If reliable, we hope to create an adaptive algorithm that uses the FSR to identify good undersampling direction(s) prior to collecting the outer portion of the Fourier domain (that region that exists outside of the FSR). By doing so, PILP would become robust to an inappropriately selected undersampling direction. Although the proposed ADM provides a method for determining appropriate directions of acceleration, it does not indicate which portions of the image experience high or low SNR. G-factor analysis is perfectly suited for this purpose and can be done after the data are collected.

In conclusion, we showed that the quality of PILP algorithms, such as GRAPPA and SPIRiT, depends on the direction of undersampling. This directional dependence is related to the amount of spatial variance in the individual coil sensitivities. The impact on the quality of reconstructions is the difference between having an image of diagnostic quality or not.

7 Appendix A: AUTO-SMASH

In a coil array with J elements, the j 'th coil has a distinct sensitivity function $C_j: \mathbb{R}^2 \rightarrow \mathbb{C}$. A composite sensitivity is generated as a linear combination of individual coil sensitivities with linear coefficients $n_j^{(0)}$ as follows: $C_0^{\text{comp}}(x, y) = \sum_{j=1}^J n_j^{(0)} C_j(x, y)$.

The composite two-dimensional MR signal takes the form

$$\begin{aligned} S^{\text{comp}}(k_x, k_y) &= \iint dx dy \sum_{j=1}^J n_j^{(0)} C_j(x, y) \rho(x, y) e^{-i(k_x x + k_y y)} \\ &= \iint dx dy C_0^{\text{comp}}(x, y) \rho(x, y) e^{-i(k_x x + k_y y)} \\ &= \mathcal{F}\{C_0^{\text{comp}} \rho\}(k_x, k_y). \end{aligned} \quad (9)$$

Suppose that there is another set of complex weights $\{n_j^{(m)}\}$ such that the linear combination of coil sensitivities yields the following composite sensitivity

$$C_m^{\text{comp}}(x, y) = \sum_{j=1}^J n_j^{(m)} C_j(x, y) \approx C_0^{\text{comp}} \exp(im\Delta k_y y). \quad (10)$$

Importantly, this approximation only needs to be valid over the support of the image. With these linear coefficients, the composite MR signal becomes

$$\begin{aligned} \sum_{j=1}^J n_j^{(m)} S_j(k_x, k_y) &= \iint_{\Omega} dx dy \left[\sum_{j=1}^J n_j^{(m)} C_j(x, y) \right] \rho(x, y) e^{-i(k_x x + k_y y)} \\ &\approx \iint_{\Omega} dx dy C_0^{\text{comp}} e^{im\Delta k_y y} \rho(x, y) e^{-i(k_x x + k_y y)} \\ &= \mathcal{F}\{C_0^{\text{comp}} \rho\}(k_x, k_y - m\Delta k_y), \end{aligned} \quad (11)$$

where Ω is the support of ρ . The $n_j^{(m)}$ coefficients serve to interpolate Fourier values at a distance of $m\Delta k_y$. Note that for the special case where $C_0^{\text{comp}} \approx 1$, the Fourier coefficients are those of ρ .

The innovation of AUTO-SMASH is to use lines of an FSR centered on the 0 frequency to estimate the weights $\{n_j^{(m)}\}$ for Eq. (10). These FSR data, S_j^{ACR} , are shifted exactly by the amount $m\Delta k_y$. The composite signal generated using weights $\{n_j^{(0)}\}$ according to Eq. (9) yields

$$S^{\text{comp}}(k_x, k_y - m\Delta k_y) = \sum_{j=1}^J n_j^{(0)} S_j^{\text{ACCS}}(k_x, k_y - m\Delta k_y). \quad (12)$$

Alternatively, following Eq. (11), we can write

$$S^{\text{comp}}(k_x, k_y - m\Delta k_y) = \sum_{j=1}^J n_j^{(m)} S_j(k_x, k_y). \quad (13)$$

Equating (12) and (13) yields

$$\sum_{j=1}^J n_j^{(m)} S_j(k_x, k_y) = \sum_{j=1}^J n_j^{(0)} S_j^{\text{ACCS}}(k_x, k_y - m\Delta k_y). \quad (14)$$

We write the right-hand side of Eq. (14) simply as $S^{\text{comp}}(k_x, k_y - m\Delta k_y)$ to reinforce that is the final, combined image produced using the original weights. For each k_x , this is a (complex) scalar, and the left-hand side is a linear combination of the collected MR signals. To determine $n^{(m)}$, minimize $\|\Sigma n^{(m)} - b\|_2$, where

$$\Sigma = \underbrace{\begin{bmatrix} S_1(k_{x_1}, k_y) & S_2(k_{x_1}, k_y) & \dots & S_J(k_{x_1}, k_y) \\ S_1(k_{x_2}, k_y) & S_2(k_{x_2}, k_y) & \dots & S_J(k_{x_2}, k_y) \\ \vdots & \ddots & \ddots & \vdots \\ S_1(k_{x_{n_x}}, k_y) & S_2(k_{x_{n_x}}, k_y) & \dots & S_J(k_{x_{n_x}}, k_y) \end{bmatrix}}_{n_x \times J},$$

$$n^{(m)} = \underbrace{\begin{bmatrix} n_1^{(m)} \\ n_2^{(m)} \\ \vdots \\ n_J^{(m)} \end{bmatrix}}_{J \times 1}, b = \underbrace{\begin{bmatrix} S^{\text{comp}}(k_{x_1}, k_y - m\Delta k_y) \\ S^{\text{comp}}(k_{x_2}, k_y - m\Delta k_y) \\ \vdots \\ S^{\text{comp}}(k_{x_{n_x}}, k_y - m\Delta k_y) \end{bmatrix}}_{n_x \times 1}.$$

Once the weights $n^{(m)}$ are determined, the matrix-vector multiplication $\Sigma n^{(m)}$ for Σ constructed for a specific (k_x, k_y) will estimate the *composite* Fourier value at $(k_x, k_y - m\Delta k_y)$.

8 Appendix B: FISTA

FISTA solves problems of the form $\text{minimize}_{\theta} \mathcal{F}(\theta) + \mathcal{G}(\theta)$ where \mathcal{F} is differentiable and \mathcal{G} has a simple proximal operator.²⁵ Let $\mathcal{F}(\theta) = \frac{1}{2} \|G\theta - \theta\|_2^2$, which is differentiable with a gradient equal to $(G - I)^T(G - I)\theta$. Let $\mathcal{G}(\theta) = \mathbb{I}_{B[\theta_D - y, \varepsilon]}(\theta)$, where $\theta_D = D\theta$. That is, \mathcal{G} is the indicator function that equals 0 whenever the collected elements of θ are within an ε ball of y , and otherwise is equal to infinity. The proximal operator of \mathcal{G} is a shrinkage function; it brings the elements of θ_D closer to y until the vector lies within the ball. With these definitions, minimizing $\mathcal{F} + \mathcal{G}$ is equivalent to solving Eq. (7). This can be done with FISTA.

Disclosures

No conflicts of interest, financial or otherwise, are declared by the authors.

Code and Data Availability

Data on the knee were acquired from mridata.org; data on the knee, shoulder, brain, and ankle were shared with Ref. 22 and can be found at www.nicholasdwork.com

Ethical Standards

All procedures performed in studies involving human participants were in accordance with the ethical standards of the institutional and/or national research committee and with the 1964

Helsinki Declaration and its later amendments or comparable ethical standards. MR data from humans were gathered with Institutional Review Board (IRB) approval and Health Insurance Portability and Accountability Act (HIPAA) compliance. Informed consent was obtained from all individual participants included in the study.

References

1. M. A. Griswold et al., "Parallel magnetic resonance imaging using the GRAPPA operator formalism," *Magn. Reson. Med.* **54**(6), 1553–1556 (2005).
2. K. P. Pruessmann et al., "Advances in sensitivity encoding with arbitrary k-space trajectories," *Magn. Reson. Med.* **46**(4), 638–651 (2001).
3. J. A. Fessler, "Model-based image reconstruction for MRI," *Signal Process. Mag.* **27**(4), 81–89 (2010).
4. J. P. Haldar and K. Setsompop, "Linear predictability in magnetic resonance imaging reconstruction: leveraging shift-invariant Fourier structure for faster and better imaging," *Signal Process. Mag.* **37**(1), 69–82 (2020).
5. K. P. Pruessmann et al., "Sense: sensitivity encoding for fast MRI," *Magn. Reson. Med.* **42**, 952–962 (1999).
6. M. Uecker et al., "ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA," *Magn. Reson. Med.* **71**(3), 990–1001 (2014).
7. R. A. Lobos, C.-C. Chan, and J. P. Haldar, "New theory and faster computations for subspace-based sensitivity map estimation in multichannel MRI," *IEEE Trans. Med. Imaging* **43**(1), 286–296 (2023).
8. M. Uecker et al., "Image reconstruction by regularized nonlinear inversion—joint estimation of coil sensitivities and image content," *Magn. Reson. Med.* **60**(3), 674–682 (2008).
9. D. K. Sodickson and W. J. Manning, "Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays," *Magn. Reson. Med.* **38**, 591–603 (1997).
10. P. M. Jakob et al., "AUTO-SMASH: a self-calibrating technique for SMASH imaging," *Magn. Reson. Mater. Phys. Biol. Med.* **7**(1), 42–54 (1998).
11. J. P. Haldar and J. Zhuo, "P-LORAKS: low-rank modeling of local k-space neighborhoods with parallel imaging data," *Magn. Reson. Med.* **75**(4), 1499–1514 (2016).
12. M. A. Griswold et al., "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magn. Reson. Med.* **47**(6), 1202–1210 (2002).
13. M. Lustig and J. M. Pauly, "SPIRiT: iterative self-consistent parallel imaging reconstruction from arbitrary k-space," *Magn. Reson. Med.* **64**(2), 457–471 (2010).
14. F. A. Breuer et al., "General formulation for quantitative G-factor calculation in GRAPPA reconstructions," *Magn. Reson. Med.* **62**(3), 739–746 (2009).
15. P. T. Weavers et al., "Acceleration apportionment: a method of improved 2D SENSE acceleration applied to 3D contrast-enhanced MR angiography," *Magn. Reson. Med.* **71**(2), 672–680 (2014).
16. D. G. Nishimura, *Principles of magnetic resonance imaging*, Stanford University (2010).
17. P. Beatty et al., "A method for autocalibrating 2D accelerated volumetric parallel imaging with clinically practical reconstruction times," in *Proc. Int. Soc. for Magn. Reson. Med.*, Vol. **15**, p. 1749 (2007).
18. C. L. Epstein, *Introduction to the Mathematics of Medical Imaging*, SIAM (2007).
19. P. B. Roemer et al., "The NMR phased array," *Magn. Reson. Med.* **16**(2), 192–225 (1990).
20. A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *J. Imaging Sci.* **2**(1), 183–202 (2009).
21. F. Ong et al., "Mridata.org: an open archive for sharing MRI raw data," in *Proc. Int. Soc. Magn. Reson. Med.*, Vol. **26** (2018).
22. N. Dwork, J. W. Gordon, and E. K. Englund, "Accelerated parallel magnetic resonance imaging with compressed sensing using structured sparsity," *J. Med. Imaging* **11**(3), 033504 (2024).
23. G. Giovannetti et al., "A fast and accurate simulator for the design of birdcage coils in MRI," *Magn. Reson. Mater. Phys. Biol. Med.* **15**(1), 36–44 (2002).
24. A. Macovski, "Noise in MRI," *Magn. Reson. Med.* **36**(3), 494–497 (1996).
25. N. Parikh et al., "Proximal algorithms," *Found. Trends Optim.* **1**(3), 127–239 (2014).

Alex McManus is a graduate student in the Department of Applied Mathematics at the University of Colorado Boulder. His research has focused on sampling theory and accurate deep learning.

Stephen Becker is an associate professor of applied mathematics at the University of Colorado Boulder, with courtesy appointments in the electrical, computer and energy engineering, and computer science departments. Previously, he was a Herman Goldstine postdoctoral fellow in mathematical sciences at IBM Research in Yorktown Heights, NY, and a postdoctoral fellow via the Fondation Sciences Mathématiques de Paris at Paris 6. He received his PhD in 2011 from California Institute of Technology under Emmanuel Candès and his bachelor's degrees in math and physics from Wesleyan University.

Nicholas Dwork is an assistant professor in the Department of Biomedical Informatics with a secondary appointment in the Department of Radiology at the University of Colorado School of Medicine. He previously worked as a postdoctoral scholar at the University of California San Francisco in the Surbeck Advanced Imaging Laboratory, advised by Peder Larson. He received a PhD in electrical engineering from Stanford University in 2019; his advisor was John Pauly. He received his BS in electrical engineering from California Polytechnic State University in San Luis Obispo in 2001, and he received his MS in biomedical engineering from the University of California, Los Angeles, in 2004. His current research interests include reducing the scan time of MRI and quantification in optical coherence tomography.