Published in final edited form as:

Magn Reson Med. 2016 April ; 75(4): 1499–1514. doi:10.1002/mrm.25717.

P-LORAKS: Low-Rank Modeling of Local k-Space Neighborhoods with Parallel Imaging Data

Justin P. Haldar¹ and Jingwei Zhuo²

¹Electrical Engineering, University of Southern California, Los Angeles, CA, USA

²Electronic Engineering, Tsinghua University, China

Abstract

Purpose—To propose and evaluate P-LORAKS, a new calibrationless parallel imaging reconstruction framework.

Theory and Methods—LORAKS is a flexible and powerful framework that was recently proposed for constrained MRI reconstruction. LORAKS was based on the observation that certain matrices constructed from fully-sampled k-space data should have low rank whenever the image has limited support or smooth phase, and made it possible to accurately reconstruct images from undersampled or noisy data using low-rank regularization. This paper introduces P-LORAKS, which extends LORAKS to the context of parallel imaging. This is achieved by combining the LORAKS matrices from different channels to yield a larger but more parsimonious low-rank matrix model of parallel imaging data. This new model can be used to regularize the reconstruction of undersampled parallel imaging data, and implicitly imposes phase, support, and parallel imaging constraints without needing to calibrate phase, support, or sensitivity profiles.

Results—The capabilities of P-LORAKS are evaluated with retrospectively undersampled data and compared against existing parallel MRI reconstruction methods. Results show that P-LORAKS can improve parallel imaging reconstruction quality, and can enable the use of new k-space trajectories that are not compatible with existing reconstruction methods.

Conclusion—The P-LORAKS framework provides a new and effective way to regularize parallel imaging reconstruction.

Keywords

Low-Rank Matrix Recovery; Parallel MRI; Support Constraints; Phase Constraints

INTRODUCTION

Relatively slow data acquisition is one of the main shortcomings of modern high-resolution MRI. As a result, techniques that enable high-quality image reconstruction from data sampled below the Nyquist rate have received a substantial amount of attention. Two

Address Correspondence to: Justin P. Haldar, University of Southern California, University Park Campus, 3740 S McClintock Avenue, Hughes Aircraft Electrical Engineering Center (EEB) #442, M/C 2564, Los Angeles, CA 90089, Tel: (213) 740-2358, jhaldar@usc.edu.

popular approaches for achieving this are parallel imaging (1–8) and constrained image reconstruction (9–22), both of which enable the reconstruction of MR images from sub-Nyquist k-space data. Parallel imaging methods sample k-space data simultaneously through a multichannel array of receiver coils. Since each coil has a distinct spatial sensitivity profile, the multichannel data contains more information than single-channel data, which enables reconstruction from sub-Nyquist data (23). In contrast, constrained image reconstruction methods use prior information to reduce the degrees of freedom of the image, which also reduces the required number of measured samples. For example, the use of support and smoothly-varying phase constraints has a long history in MRI image reconstruction from sparsely-sampled data (9–16).

Both parallel imaging and constrained image reconstruction are complementary, and it is no surprise that their combination is more powerful than using either approach on its own (24–48). Some previous work has combined parallel imaging with phase and/or support constraints by invoking “projection-onto-convex sets” methods (24,25), virtual coil concepts that employ conjugate symmetry (49), pseudoinverse/regularized least-squares formulations (26–30), and many other alternative strategies. These previous methods have relied on the use of prior information and/or calibration-based k-space sampling that enable reasonably accurate pre-estimation of an image’s phase, support, coil sensitivity profiles, and/or inter-coil k-space dependencies. For this paper, we say that a sampling trajectory is “calibration-based” if there exists a Nyquist-sampled region of k-space that can be used to calibrate these quantities.

This paper explores the combination of parallel imaging data with low-rank matrix modeling of local k-space neighborhoods (LORAKS) (22,50). LORAKS, a recently described constrained image reconstruction framework designed for single-channel data, makes use of the fact that images with limited spatial support and/or smoothly-varying spatial image phase will have redundancies in k-space that can be exploited to reduce sampling requirements. However, it was shown in Ref. (22) that LORAKS imposes support and phase constraints in a fundamentally different way from previous methods, and can yield substantial improvements in image reconstruction quality. Unlike existing work, the LORAKS framework does not require prior knowledge of the image phase or support, and is flexible enough to be used with both calibrationless and calibration-based k-space sampling trajectories. This is possible because LORAKS uses low-rank matrix modeling to implicitly impose phase and support constraints, rather than traditional approaches that require explicit representations of the image phase and support (9–16).

Low-rank matrix models are powerful for constrained reconstruction because they are flexible enough to represent a wide variety of different datasets, and are also quite parsimonious (i.e., number of degrees of freedom in a low-rank matrix is generally much smaller than the number of matrix entries (51)). Besides LORAKS, various forms of low-rank matrix modeling have also been successfully applied to other settings like dynamic imaging (17–20) and parallel imaging (45–48).

One of the key ingredients for the LORAKS framework was the theoretical observation that, under limited support or smooth phase assumptions, there exist linear “shift-invariant”

interpolation kernels that allow any given k-space value to be linearly predicted from neighboring k-space samples. It is interesting to note that calibration-based k-space based parallel imaging reconstruction methods like SMASH (2), GRAPPA (4), SPIRiT (5), and PRUNO (6) are also based on the use of linear shift-invariant k-space interpolation from a local neighborhood, which suggests that LORAKS might be easily generalized to the parallel imaging context. Different from the LO-RAKS framework (which uses intra-channel relationships between neighboring k-space samples), methods like SMASH, GRAPPA, SPIRiT, and PRUNO made use of linear interpolation kernels that interpolate fully sampled data for each channel based on undersampled measured data samples from multiple channels (i.e., inter-channel relationships). These inter-channel relationships also imply the existence of a nullspace for an appropriately constructed matrix, which has been shown and used in a variety of previous parallel imaging methods (6, 46–48).

This paper introduces the P-LORAKS framework, which is a parallel imaging generalization of LORAKS that combines the intra-channel modeling relationships of LORAKS with the interchannel modeling relationships of parallel imaging methods. As this paper will demonstrate, P-LORAKS is quite flexible: it can be used with a range of different calibration-based and calibrationless sampling trajectories (including specialized trajectories that, to the best of our knowledge, aren't compatible with any existing reconstruction methods), and can also be used in combination with other regularization constraints.

It should be noted that multiple calibrationless parallel imaging reconstruction methods have recently been proposed (43–46), each enabled by different modeling assumptions. Specifically, the CLEAR approach (45) used an image-domain locally low-rank model of multi-channel image patches, the methods in Refs. (43, 44) used a spatial-domain joint sparsity model for the images from different channels, and the SAKÉ approach (46) used a low-rank model of multi-channel k-space neighborhoods. Among these different methods, P-LORAKS has the highest similarity with SAKÉ, since both approaches construct low-rank matrix models based on local neighborhoods of k-space data. However, SAKÉ and P-LORAKS were derived independently based on different theoretical modeling assumptions (the construction of the SAKÉ matrix was primarily motivated using theoretical parallel imaging relationships, rather the theoretical k-space relationships arising from image support or phase assumptions as used by LORAKS),¹ and they can be quite different from each other when P-LORAKS incorporates phase constraints. The results of this paper indicate that the use of phase constraints can give P-LORAKS a substantial performance advantage.

A preliminary account of portions of this work was previously presented in Ref. (53).

¹It should be noted that the SAKÉ journal paper (46) describes an unexplained empirical relationship between the size of the image support and the rank of the SAKÉ matrix. However, we were unaware of this during the development of LORAKS and P-LORAKS, since this relationship was not described in the early SAKÉ abstracts or manuscript preprints. The support-based LORAKS and P-LORAKS formulations were developed independently from SAKÉ, and were primarily motivated by theoretical relationships described in earlier signal processing literature (52). The empirical relationship between rank and support described in the SAKÉ journal paper can be justified using LORAKS theory.

THEORY

LORAKS

Before introducing our proposed P-LORAKS framework, we will first review the high-level details of the original single-channel LORAKS framework. A more complete description can be found in Ref. (22). Without loss of generality, we will only describe LORAKS and P-LORAKS for 2D imaging, noting that higher-dimensional extensions are straightforward. Let $\rho(x, y)$ be a 2D image with Nyquist-sampled Cartesian k-space samples $\rho(n_x, n_y)$, where n_x and n_y are the k-space sampling intervals along the k_x and k_y axes, respectively, and n_x and n_y are integer sample indices. To simplify notation, we will assume for the rest of the paper that the FOV has been normalized such that $k_x = k_y = 1$. We assume that we are interested in reconstructing $\rho(n_x, n_y)$ for all values of n_x between $-N_x$ and $+N_x$ and all values of n_y between $-N_y$ and N_y , where N_x and N_y are positive integers that define the k-space measurement region. We will use the symbol $\mathbf{k}$ to denote the vector of the full set of these noiseless samples. We will also use $\mathbf{d}$ to denote the vector of undersampled and/or noisy measured k-space data, and will use $\mathbf{F}$ to denote the matrix that describes the sampling operation (i.e., $\mathbf{F}\mathbf{k} = \mathbf{d}$ in the absence of noise). For subsampled Cartesian trajectories, $\mathbf{F}$ can be formed by discarding rows from the identity matrix. This is the strategy employed for all examples shown in this paper. For non-Cartesian trajectories, $\mathbf{F}$ will be a matrix that interpolates Cartesian samples onto non-Cartesian locations. Our approach for constructing $\mathbf{F}$ for non-Cartesian data would be identical to the non-Cartesian matrix constructions employed in SAKÉ (46,54). This non-Cartesian $\mathbf{F}$ matrix is practical and easy to construct – please see Refs. (46, 54) for further detail.

The LORAKS framework is based on the observation that it is possible to arrange the values of $\rho(n_x, n_y)$ into low-rank matrices whenever $\rho(x, y)$ has limited spatial support and/or slowly varying phase. Specifically, it was shown in Ref. (22) that if the support of $\rho(x, y)$ does not occupy the entire FOV, then it is possible to construct non-zero functions $f(n_x, n_y)$ such that the convolution of $\rho(n_x, n_y)$ with $f(n_x, n_y)$ is identically equal to zero at all points in k-space. If we further assume that $f(n_x, n_y)$ is bandlimited in k-space (i.e., $f(n_x, n_y) \approx 0$ whenever $n_x^2 + n_y^2 > R^2$, for some appropriate choice of the k-space radius R), this convolution relationship implies that a Hankel-like matrix $\mathbf{C} \in \mathbb{C}^{K \times N_R}$ formed as

$$[\mathbf{C}]_{km} = \tilde{\rho}(n_x^{(k)} - p_m, n_y^{(k)} - q_m), \quad k=1, \dots, K, \quad m=1, \dots, N_R \quad [1]$$

will have approximately low rank. In Eq. [1], we have used $\left\{ \left(n_x^{(k)}, n_y^{(k)} \right) \right\}_{k=1}^K$ to denote K distinct k-space locations, and typically choose these to be the full set of k-space locations (n_x, n_y) from the Cartesian grid that satisfy $-N_x + R \leq n_x \leq N_x - R$ and $-N_y + R \leq n_y \leq N_y - R$. We have also used $\{(p_m, q_m)\}_{m=1}^{N_R}$ to denote an ordered set of distinct elements from the set $\Lambda_R = \{(p, q) \in \mathbb{Z}^2: (p^2 + q^2) \leq R^2\}$, and have used N_R to denote the cardinality of Λ_R . Each row of the $\mathbf{C}$ matrix corresponds to a local neighborhood of k-space, where the neighborhood contains all k-space points within distance R of the neighborhood center $(n_x^{(k)}, n_y^{(k)})$.

Low-rank matrices can similarly be constructed using assumptions about the image phase, based on the fact that real-valued images will have conjugate symmetric Fourier transforms. Specifically, Ref. (22) showed how to use information from opposite sides of k-space to construct two different matrices that would have low-rank for images with slowly-varying phase. In this paper, we focus on the matrix $\mathbf{S} \in \mathbb{R}^{2K \times 2NR}$ from Ref. (22), which was the more powerful of the two different phase-based low-rank matrix constructions. The $\mathbf{S}$ matrix can be defined as

$$\mathbf{S} = \begin{bmatrix} \mathbf{S}_{r+} - \mathbf{S}_{r-} & -\mathbf{S}_{i+} + \mathbf{S}_{i-} \\ \mathbf{S}_{i+} + \mathbf{S}_{i-} & \mathbf{S}_{r+} + \mathbf{S}_{r-} \end{bmatrix}, \quad [2]$$

where the matrices $\mathbf{S}_{r+}$, $\mathbf{S}_{r-}$, $\mathbf{S}_{i+}$, $\mathbf{S}_{i-} \in \mathbb{R}^{K \times NR}$ respectively have elements:

$$[\mathbf{S}_{r+}]_{km} = \tilde{\rho}^r \left(n_x^{(k)} - p_m, n_y^{(k)} - q_m \right), \quad [3]$$

$$[\mathbf{S}_{r-}]_{km} = \tilde{\rho}^r \left(-n_x^{(k)} - p_m, -n_y^{(k)} - q_m \right), \quad [4]$$

$$[\mathbf{S}_{i+}]_{km} = \tilde{\rho}^i \left(n_x^{(k)} - p_m, n_y^{(k)} - q_m \right), \quad [5]$$

$$[\mathbf{S}_{i-}]_{km} = \tilde{\rho}^i \left(-n_x^{(k)} - p_m, -n_y^{(k)} - q_m \right), \quad [6]$$

for $k = 1, \dots, K$ and $m = 1, \dots, NR$. Here, we have used $\tilde{\rho}^r(n_x, n_y)$ and $\tilde{\rho}^i(n_x, n_y)$ to respectively denote the real and imaginary components of $\rho(n_x, n_y)$.

Frequently, the rank of $\mathbf{C}$ will become smaller as the support of the image decreases, while the rank of $\mathbf{S}$ will become smaller as the image phase gets smoother (22). Lower-rank $\mathbf{S}$ matrices can also be associated with images that have smaller spatial supports. As a result, encouraging $\mathbf{C}$ and $\mathbf{S}$ to have low rank during image reconstruction will encourage the reconstructed images to also have these characteristics. The LORAKS reconstruction approach described in Ref. (22) encouraged this low-rank structure using regularization by solving:

$$\hat{\mathbf{k}} = \arg \min_{\mathbf{k}} \|\mathbf{F}\mathbf{k} - \mathbf{d}\|_{\ell_2}^2 + \lambda_C \|\mathbf{C}(\mathbf{k}) - \mathbf{C}_r(\mathbf{k})\|_F^2 + \lambda_S \|\mathbf{S}(\mathbf{k}) - \mathbf{S}_r(\mathbf{k})\|_F^2. \quad [7]$$

In this expression, λ_C and λ_S are regularization parameters that control how strongly the rank constraints are imposed; the matrices $\mathbf{C}(\mathbf{k})$ and $\mathbf{S}(\mathbf{k})$ are formed by arranging the values of $\mathbf{k}$ according to Eq. [1] and Eqs. [2]-[6], respectively; and the matrices $\mathbf{C}_r(\mathbf{k})$ and $\mathbf{S}_r(\mathbf{k})$ are optimal low-rank approximations of $\mathbf{C}(\mathbf{k})$ and $\mathbf{S}(\mathbf{k})$. Specifically, given user-defined rank parameters r_C and r_S , the matrix $\mathbf{C}_r(\mathbf{k})$ is obtained by truncating the singular value decomposition (SVD) of $\mathbf{C}(\mathbf{k})$ at rank r_C , and the matrix $\mathbf{S}_r(\mathbf{k})$ is obtained by truncating the

singular value decomposition (SVD) of $\mathbf{S}(\mathbf{k})$ at rank r_S . Note that, as our notation implies, the matrices $\mathbf{C}_r(\mathbf{k})$ and $\mathbf{S}_r(\mathbf{k})$ are functions of $\mathbf{k}$.

The regularization terms in Eq. [7] measure how well $\mathbf{C}(\mathbf{k})$ and $\mathbf{S}(\mathbf{k})$ are approximated as rank- r_C and rank- r_S matrices, respectively, and are effective at encouraging low-rank structure. Since the LORAKS constraints appear as simple regularizers, it is also straightforward to augment Eq. [7] with other forms of regularization if desired (e.g., sparsity-promoting ℓ_1 -regularization (21)).

P-LORAKS

In parallel imaging, we observe k-space data simultaneously from L different receiver coils, with each channel observing a slightly different image $\rho_\ell(x, y)$. In particular, these images are related according to $\rho_\ell(x, y) = s_\ell(x, y)\rho(x, y)$, where $s_\ell(x, y)$ is the sensitivity profile of the ℓ th coil. Notably, if $\rho(x, y)$ is support limited, then $\rho_\ell(x, y)$ will also be support limited, and could potentially have a much smaller support than $\rho(x, y)$ if the sensitivity profiles are highly localized to specific regions within the FOV (3). This implies that the $\mathbf{C}$ -matrices $\{\mathbf{C}_\ell\}_{\ell=1}^L$ corresponding to each individual coil will each have low rank. Similarly, the $\mathbf{S}$ -matrices $\{\mathbf{S}_\ell\}_{\ell=1}^L$ for each individual coil will also frequently have low rank if $\rho(x, y)$ has slowly varying phase, since sensitivity profiles $s_\ell(x, y)$ usually also have slowly varying phase, and the phase of $\rho_\ell(x, y)$ is simply the sum of the phases of $\rho(x, y)$ and $s_\ell(x, y)$. Representative examples of real multichannel brain images are shown in Fig. 1 to illustrate typical support and phase characteristics.

The most direct approach to applying LORAKS to parallel imaging data would be to use LORAKS to reconstruct each coil image independently. However, in this work, we observe that the bigger P-LORAKS matrices formed as

$$\mathbf{C}^P = [\mathbf{C}_1, \mathbf{C}_2, \dots, \mathbf{C}_L] \quad [8]$$

and

$$\mathbf{S}^P = [\mathbf{S}_1, \mathbf{S}_2, \dots, \mathbf{S}_L] \quad [9]$$

could have even better low-rank characteristics. In particular, it is straightforward to show that $\text{rank}(\mathbf{C}^P) \leq \sum_{\ell=1}^L \text{rank}(\mathbf{C}_\ell)$ and $\text{rank}(\mathbf{S}^P) \leq \sum_{\ell=1}^L \text{rank}(\mathbf{S}_\ell)$, and that the nullspace vectors of $\{\mathbf{C}_\ell\}_{\ell=1}^L$ and $\{\mathbf{S}_\ell\}_{\ell=1}^L$ can be zero-padded to form nullspace vectors of $\mathbf{C}^P$ and $\mathbf{S}^P$. This implies that the conversion of the individual coil matrices into the bigger P-LORAKS matrices does not cause any loss of low-rank matrix structure.

On the other hand, the P-LORAKS matrices potentially have a much larger number of nullspace vectors as the result of inter-channel correlations. For example, existing parallel imaging methods like PRUNO, ESPIRiT, and SAKÉ (6, 46, 48) constructed matrices that are essentially the same as the $\mathbf{C}^P$ matrix, and suggested that these matrices should have

low-rank structure as the result of the special structure of the parallel imaging inverse problem. These arguments were based on the observation (4, 5) that it is frequently possible to find sets of weighting coefficients $w_{\ell mj}$ (i.e., “interpolation kernels”) that satisfy

$$\tilde{\rho}_j(n_x, n_y) \approx \sum_{\ell=1}^L \sum_{m=1}^{N_R} w_{\ell mj} \tilde{\rho}_\ell(n_x - p_m, n_y - q_m) \quad [10]$$

for all possible choices of (n_x, n_y) and for each $j = 1, \dots, L$. In this expression, $\tilde{\rho}_\ell(n_x, n_y)$ is the ideal k-space data from the ℓ th channel, and the (p_m, q_m) values were defined in Eq. [1]. Functionally, Eq. [10] indicates that k-space data in one channel can be “interpolated” by linearly combining information from neighboring k-space data in all the channels. In addition, the interpolation kernels do not depend on the specific k-space location (they are shift invariant).

After simple rearrangement (6, 48), Eq. [10] is equivalent to the existence of non-zero coefficients $\beta_{\ell m}$ such that

$$\sum_{\ell=1}^L \sum_{m=1}^{N_R} \beta_{\ell m} \tilde{\rho}_\ell(n_x - p_m, n_y - q_m) \approx 0 \quad [11]$$

for all possible choices of (n_x, n_y) . It is straightforward to show that any sets of $\beta_{\ell m}$ coefficients satisfying Eq. [11] can be rearranged into approximate nullspace vectors of $\mathbf{C}^P$.

This implies that the rank of $\mathbf{C}^P$ could be much smaller than $\sum_{\ell=1}^L \text{rank}(\mathbf{C}_\ell)$. Similar arguments can be constructed that imply that $\mathbf{S}^P$ could have a much smaller rank than $\sum_{\ell=1}^L \text{rank}(\mathbf{S}_\ell)$.

Similar to the LORAKS-based reconstruction formula given in Eq. [7], the P-LORAKS reconstruction problem can be posed as

$$\hat{\mathbf{k}}^P = \arg \min_{\mathbf{k}^P} \|\mathbf{F}^P \mathbf{k}^P - \mathbf{d}^P\|_{\ell_2}^2 + \lambda_C \|\mathbf{C}^P(\mathbf{k}^P) - \mathbf{C}_r^P(\mathbf{k}^P)\|_F^2 + \lambda_S \|\mathbf{S}^P(\mathbf{k}^P) - \mathbf{S}_r^P(\mathbf{k}^P)\|_F^2. \quad [12]$$

In this expression, $\|\cdot\|_{\ell_2}$ and $\|\cdot\|_F$ are respectively the standard ℓ_2 and Frobenius norms; $\mathbf{k}^P$ and $\mathbf{d}^P$ are the vectors obtained by concatenating the L different fully-sampled k-space ($\mathbf{k}$) and measured k-space ($\mathbf{d}$) vectors for each channel; $\mathbf{F}^P$ is a block-diagonal matrix with L blocks and each block equal to the $\mathbf{F}$ matrix; the matrices $\mathbf{C}^P(\mathbf{k}^P)$ and $\mathbf{S}^P(\mathbf{k}^P)$ are formed by arranging the values of $\mathbf{k}^P$ according to Eqs. [8] and [9], respectively; and $\mathbf{C}_r^P(\mathbf{k}^P)$ and $\mathbf{S}_r^P(\mathbf{k}^P)$ are optimal low-rank approximations of $\mathbf{C}^P(\mathbf{k}^P)$ and $\mathbf{S}^P(\mathbf{k}^P)$, obtained using SVD truncation based on user-defined rank parameters r_C and r_S .

Since the P-LORAKS cost function in Eq. [12] is very similar to the LORAKS cost function, it is possible to use the same algorithm to minimize them both. In this paper, we use the algorithm proposed for LORAKS reconstruction in Ref. (22).

Optimization Algorithm

While both the LORAKS and P-LORAKS cost functions from Eqs. [7] and [12] are non-convex and might seem complicated, local optima can be obtained by a simple majorize-minimize algorithm that is guaranteed to monotonically decrease the cost function value until it converges (22). Majorize-minimize algorithms (55, 56) are simple alternation-based iterative algorithms that are classical and popular for a variety of optimization problems. For example, the well-known expectation-maximization algorithm is a majorize-minimize algorithm, and several algorithms for optimizing ℓ_1 -regularized cost functionals also fall within the majorize-minimize framework. See Refs. (55, 56) for further discussion.

The majorize-minimize algorithm corresponding to P-LORAKS is presented below, and the algorithm for LORAKS from Ref. (22) is recovered when $L = 1$. See Ref. (22) for a more detailed description of how each step was derived, and see Ref. (50) for LORAKS Matlab code.

The algorithm proceeds as follows:

1. Set iteration number $i = 0$, and construct an initial guess for $\mathbf{k}^{\hat{P}}$. For all examples shown in this work, we initialize $\mathbf{k}^{\hat{P}}$ with a zero-filled version of $\mathbf{d}^P$.
2. Based on the estimate of $\mathbf{k}^{\hat{P}}$ from the previous iteration, construct the matrices $\mathbf{C}^P(\mathbf{k}^{\hat{P}})$ and $\mathbf{S}^P(\mathbf{k}^{\hat{P}})$ based on Eqs. [8] and [9].
3. Compute partial SVDs of $\mathbf{C}^P(\mathbf{k}^{\hat{P}})$ and $\mathbf{S}^P(\mathbf{k}^{\hat{P}})$, and construct rank- r_C and r_S approximations $\mathbf{C}_r^P(\hat{\mathbf{k}}^P)$ and $\mathbf{S}_r^P(\hat{\mathbf{k}}^P)$ using SVD truncation.
4. Compute “synthetic” k-space data vectors $\mathbf{k}_C^P$ and $\mathbf{k}_S^P$ based on $\mathbf{C}_r^P(\hat{\mathbf{k}}^P)$ and $\mathbf{S}_r^P(\hat{\mathbf{k}}^P)$, respectively. The entry for $\mathbf{k}_C^P$ corresponding to location (n_x, n_y) and coil ℓ is obtained by summing all of the entries from $\mathbf{C}_r^P(\hat{\mathbf{k}}^P)$ where $\rho_{\ell}(n_x, n_y)$ originally appeared in $\mathbf{C}^P(\mathbf{k}^P)$. The entries of $\mathbf{k}_S^P$ are obtained in a similar, though slightly more complicated manner. In particular, the real part of $\mathbf{k}_S^P$ corresponding to location (n_x, n_y) and coil ℓ is obtained by summing all entries from $\mathbf{S}_r^P(\hat{\mathbf{k}}^P)$ where $\tilde{\rho}_{\ell}^T(n_x, n_y)$ originally appeared in $\mathbf{S}^P(\mathbf{k}^P)$ with positive sign, and then subtracting all entries from $\mathbf{S}_r^P(\hat{\mathbf{k}}^P)$ where $\tilde{\rho}_{\ell}^T(n_x, n_y)$ originally appeared in $\mathbf{S}^P(\mathbf{k}^P)$ with negative sign. The imaginary part of $\mathbf{k}_S^P$ is obtained in an identical manner, except with $\tilde{\rho}_{\ell}^T(n_x, n_y)$ replaced by $\tilde{\rho}_{\ell}^i(n_x, n_y)$.
5. Update the estimate of $\mathbf{k}^P$ according to

$$\hat{\mathbf{k}}^P = \left((\mathbf{F}^P)^H \mathbf{F}^P + \lambda_C \mathbf{P}_C + \lambda_S \mathbf{P}_S \right)^{\dagger} (\mathbf{d}^P + \lambda_C \mathbf{k}_C^P + \lambda_S \mathbf{k}_S^P). \quad [13]$$

In this expression, H is used to denote the matrix conjugate transpose operation, † is used to denote the matrix pseudoinverse, and $\mathbf{P}_C$ and $\mathbf{P}_S$ are diagonal matrices with diagonal entries constructed according to the structure of the $\mathbf{C}^P(\mathbf{k}^P)$ and $\mathbf{S}^P(\mathbf{k}^P)$

matrices. Specifically, the diagonal entry of $\mathbf{P}_C$ corresponding to location (n_x, n_y) and coil ℓ should be set equal to the integer number of times that $\rho_\ell(n_x, n_y)$ appears in $\mathbf{C}^P(\mathbf{k}^P)$, while the diagonal entry of $\mathbf{P}_S$ corresponding to location (n_x, n_y) and coil ℓ should be set equal to the integer number of times that $\tilde{\rho}_\ell^r(n_x, n_y)$ appears (regardless of sign) in $\mathbf{S}^P(\mathbf{k}^P)$.

If k -space is measured on a subsampled Cartesian grid, the matrices appearing in Eq. [13] are all diagonal, which means that the pseudoinverse in Eq. [13] can be easily computed by inspection (the pseudoinverse of a diagonal matrix is obtained by inverting the value of every non-zero entry). If data is acquired with a non-Cartesian trajectory, the matrices appearing in Eq. [13] are all sparse, and the pseudoinverse computation can be easily performed using, e.g., the conjugate gradient algorithm (57).

6. Increment i . Repeat steps 2–5 until convergence is achieved.

Comparisons with the SAKÉ Formulation

Due to their similarities, it is insightful to compare P-LORAKS with the SAKÉ formulation (46). Specifically, while SAKÉ was derived independently based on different assumptions than LORAKS and P-LORAKS, the SAKÉ approach relies on a low-rank matrix that is nearly identical to the $\mathbf{C}^P$ matrix. The main difference between the $\mathbf{C}^P$ matrix and the SAKÉ matrix is that SAKÉ uses a rectangularly-shaped k -space neighborhood, while P-LORAKS uses a circular k -space neighborhood to ensure that the estimated nullspace components have isotropic resolution in the image domain. In practice, this difference does not lead to substantial reconstruction differences, and the SAKÉ matrix can be considered to be essentially the same as the $\mathbf{C}^P$ matrix. It should also be noted that the calibration matrices used in PRUNO (6) and ESPIRiT (48) have the same structure as the SAKÉ matrix and thus also have similarities to the $\mathbf{C}^P$ matrix from P-LORAKS (though PRUNO and ESPIRiT both construct this matrix based on fully-sampled calibration data).

We would also like to note that the SAKÉ cost function is essentially the same as the P-LORAKS cost function as defined in Eq. [12] with $\lambda_C = \infty$ and $\lambda_S = 0$. In particular, the SAKÉ formulation imposes rank constraints quite strictly on the SAKÉ matrix, while P-LORAKS allows the user to modify the regularization parameters in order to adjust the trade-off between data fidelity and the rank constraints. Importantly, SAKÉ does not impose the phase constraints that are provided by the P-LORAKS $\mathbf{S}^P$ matrix. As will be demonstrated, the phase constraints provided by $\mathbf{S}^P$ can be quite powerful.

Algorithmically, the SAKÉ approach uses the Cadzow algorithm (58) for optimization, which is an algorithm similar to projection onto convex sets (except that one of the constraint sets is nonconvex). Unlike the majorize-minimize algorithm employed by P-LORAKS, the Cadzow algorithm will generally not monotonically decrease the cost function and is unlikely to converge to a local minimum. Despite this, the sub-optimal solutions obtained using the Cadzow algorithm can still produce nice reconstruction results.

METHODS

The potential of P-LORAKS was evaluated empirically using two retrospectively-undersampled *in vivo* parallel imaging datasets:

- **T1-Weighted Brain Dataset.** Fully-sampled brain k-space data was collected at our imaging center using a 3D MPRAGE sequence on a Siemens Tim Trio 3T scanner, using a 32 channel headcoil. Data was acquired on a $220 \times 220 \times 152$ Cartesian sampling grid, corresponding to a $210 \text{ mm} \times 210 \text{ mm} \times 154 \text{ mm}$ FOV. A 1D Fourier transform was performed along the frequency-encoding dimension (superior-inferior) to enable independent reconstruction of 2D images. For simplicity, a single 2D slice of 220×152 k-space data was extracted for use in our experiments. This dataset was shown in Fig. 1a.
- **T2-Weighted Brain Dataset.** Fully-sampled brain k-space data was collected at our imaging center using a 2D multislice T2-weighted turbo spin-echo sequence on a Siemens Tim Trio 3T scanner, using a 12 channel headcoil. For each slice, data was acquired on a 256×187 Cartesian sampling grid, corresponding to a $256 \text{ mm} \times 187 \text{ mm}$ FOV, with 1 mm slice thickness. Our experiments use a single slice from this dataset. This dataset was shown in Fig. 1b.

Each dataset was retrospectively undersampled using several different calibration-based and calibrationless sampling trajectories. For the 2D T2-weighted brain dataset, undersampling was performed along the single phase-encoding dimension with full sampling along the readout dimension. For the 3D T1-weighted brain dataset, undersampling was performed simultaneously along both phase encoding dimensions.

All LORAKS and P-LORAKS reconstructions were performed using Matlab. LORAKS reconstruction was performed independently for each channel using the code distributed along with Ref. (50), while P-LORAKS reconstruction was performed using a small modification of that code.

For LORAKS and P-LORAKS, the regularization parameters λ_C and λ_S were set to very small values (i.e., 10^{-6} divided by the number of elements in the matrix) to interpolate/extrapolate un-sampled data without substantially perturbing the measured data, and the neighborhood radius was set to $R = 2$ unless otherwise specified. Note that $R = 2$ corresponds to a circle of diameter 5, which is similar to the neighborhood size we use for SPIRiT reconstruction. The rank parameters r_C and r_S for each reconstruction were optimized to yield minimum normalized root-mean-square reconstruction errors (NRMSEs) with respect to the fully sampled datasets. To reduce the dimension of the reconstruction parameter search space and to illustrate the differences between $\mathbf{C}^P$ -based and $\mathbf{S}^P$ -based reconstruction, LORAKS and P-LORAKS were never implemented with both $\mathbf{C}^P$ -based and $\mathbf{S}^P$ -based constraints simultaneously (i.e., one of the regularization parameters λ_C and λ_S was always set equal to zero in every result that we show), and we show separate results for $\mathbf{C}^P$ -based reconstruction and $\mathbf{S}^P$ -based reconstruction.

SAKÉ reconstructions were also performed for both calibration-based and calibrationless sampling trajectories. As expected, these results were similar to $\mathbf{C}^P$ -based P-LORAKS, and are not shown.

For calibration-based sampling trajectories, we compared LORAKS and P-LORAKS against SPIRiT (5), using code downloaded from <http://www.eecs.berkeley.edu/~mlustig/Software.html>. We used a 5×5 kernel size and the default Tikhonov regularization parameters for calibration and reconstruction.

For calibrationless sampling trajectories, we compared LORAKS and P-LORAKS against the calibrationless joint-sparsity approach described in Ref. (44). The calibrationless joint-sparsity approach optimizes the joint total variation of the images according to

$$\hat{\mathbf{k}}^P = \arg \min_{\mathbf{k}^P} \|\mathbf{F}^P \mathbf{k}^P - \mathbf{d}^P\|_{\ell_2}^2 + \lambda \int_{\text{FOV}} \sqrt{\sum_{\ell=1}^L \left| \frac{\partial}{\partial x} \rho_{\ell}(x, y) \right|^2 + \left| \frac{\partial}{\partial y} \rho_{\ell}(x, y) \right|^2} dx dy. \quad [14]$$

In this expression, $\rho_{\ell}(x, y)$ is the image from the ℓ th coil that is formed by applying the inverse Fourier transform to the fully sampled data $\mathbf{k}^P$, and λ is a user-defined regularization parameter. The data fidelity term in this expression is identical to the data fidelity term used in P-LORAKS. The joint total variation regularization penalty in this expression encourages the set of multichannel images to have sparse edges, with the significant image edges occurring in the same spatial locations for different coils. Due to similarities with our previous work on the joint reconstruction of correlated MRI images (59, 60), minimization of Eq. [14] was achieved using the algorithm described in Ref. (60). The regularization parameter λ was optimized for minimum NRMSE.

For all methods, iterations were halted if the relative change (measured in the ℓ_2 -norm) between consecutive iterates was less than 10^{-4} , or if the total number of iterations exceeded 10^3 . Reconstruction results are visualized after combining the multi-channel images using root sum-of-squares.

RESULTS

Empirical Rank Characteristics of LORAKS and P-LORAKS Matrices

Before evaluating the reconstruction characteristics of P-LORAKS, we first evaluated the modeling capabilities of P-LORAKS in relation to conventional LORAKS. Specifically, we constructed LORAKS and P-LORAKS matrices for several different R values for each fully-sampled dataset. Subsequently, we computed the amount of error we would observe if we made optimal low-rank approximations of the matrices. Optimal low-rank approximations were computed using SVD truncation.

Results of this comparison are shown in Fig. 2. This figure shows that, for fixed total rank and fixed neighborhood radius R , the P-LORAKS matrices are always more accurately approximated by low-rank matrices than the LORAKS matrices (except when the rank is full and the SVDs are not truncated, in which case the error is zero for both P-LORAKS and LORAKS). The gap between P-LORAKS and LORAKS is frequently quite substantial,

which confirms our expectations that P-LORAKS can be used to more accurately and parsimoniously model MRI data than LORAKS.

We also observe that the disparity between P-LORAKS and LORAKS is more pronounced for larger values of R , which suggests that larger values of R might yield better results. However, larger R typically also leads to larger $\mathbf{C}$ and $\mathbf{S}$ matrices, which corresponds to increased computational cost and larger potential for overfitting of the data. The selection of R is discussed theoretically in (22), and will be discussed empirically later in the paper. It should be noted that our choice to use $R = 2$ in most of the following reconstruction results was based on a practical trade-off between reconstruction quality and computation time.

Calibration-Based Reconstruction

Results using calibration-based trajectories are respectively shown for the T2-weighted and T1-weighted brain datasets in Figs. 3 and 4. The calibration region consisted of a 16 fully-sampled phase encoding lines at the center of k-space for the T2-weighted dataset, and a 16×16 fully-sampled region at the center of k-space for the T1-weighted dataset. We show results for three different kinds of sampling trajectories: (a) poisson disk random sampling (61), (b) a conventional 5/8ths partial Fourier sampling pattern (14) that was randomly undersampled according to the poisson disk distribution, and (c) structured sampling obtained by uniform spacing of the phase encoding lines, followed by rounding of each line location to the nearest Cartesian grid point. For the T2-weighted dataset, the number of measured samples is half the number of fully-sampled data points (an acceleration factor of 2). For the T1-weighted dataset, the number of measured samples is a sixth of the number of fully-sampled data points (an acceleration factor of 6).

Poisson disk random undersampling is the trajectory previously used with SAKÉ reconstruction (46), and random undersampling is frequently advocated for the reconstruction of sparse or low-rank signals from undersampled data. As shown in Figs. 3a and 4a, we observe that with this trajectory, the two P-LORAKS methods substantially outperform the reconstructions obtained with single-channel LORAKS. This is consistent with the expected advantages of P-LORAKS over single-channel LORAKS. In addition, $\mathbf{S}$ -based reconstructions were slightly better than $\mathbf{C}$ -based reconstructions for both LORAKS and P-LORAKS, which is consistent with previous observations that the LORAKS-based phase constraints are often more valuable than the LORAKS-based support constraints (22, 50). SPIRiT reconstruction was slightly worse than both P-LORAKS reconstructions for this sampling pattern.

Randomly undersampled partial Fourier acquisition is not common in previous literature, because it generally requires the use of phase constraints in order to account for the unsampled region of k-space. However, a potential advantage of this sampling scheme is that it enables relatively large undersampling factors with more densely-packed samples than would be achieved using a more conventional undersampling pattern. Unsurprisingly, the results shown in Figs. 3b and 4b for this sampling pattern are relatively poor for the methods that do not incorporate phase information: SPIRiT and $\mathbf{C}$ -based LORAKS and P-LORAKS. On the other hand, $\mathbf{S}$ -based P-LORAKS reconstruction yielded consistently strong performance for both datasets. $\mathbf{S}$ -based LORAKS performed well when the

undersampling factor was small (Fig. 3b), since one-half of k-space was almost fully sampled and we would not expect parallel imaging constraints to assist very much in extrapolating the opposite side of k-space. However, **S**-based LORAKS was dominated by **S**-based P-LORAKS when the undersampling factor was more substantial (Fig. 4b).

Structured undersampling is most commonly used for parallel imaging methods like GRAPPA, SENSE, and SPIRiT that do not incorporate more advanced regularization constraints, and is rarely advocated for reconstructions that rely on sparsity or low-rank. However, as shown in Figs. 3c and 4c, both **C**-based and **S**-based P-LORAKS reconstructions performed well with this trajectory, and were substantially better than the single-channel LORAKS reconstructions. Both P-LORAKS results were slightly better than SPIRiT reconstruction for both datasets.

While no single reconstruction approach was uniformly the best across all datasets and all sampling patterns, we observed that **S**-based P-LORAKS had the best performance in 4 out of 6 cases. In the remaining two cases (Figs. 3b and 4c), the performance of **S**-based P-LORAKS was only slightly lower than the best-performing method. Interestingly, structured sampling consistently yielded the best reconstruction performance across all sampling patterns, even though random sampling is usually advocated for nonlinear reconstruction of undersampled data.²

Calibrationless Reconstruction

Results using calibration-based trajectories are respectively shown for the T2-weighted and T1-weighted brain datasets in Figs. 5 and 6. As with the calibration-based reconstructions, we used poisson disk sampling, randomly undersampled partial Fourier sampling, and structured sampling, with acceleration factors of 2 and 6 respectively for the T2-weighted and T1-weighted datasets. Undersampled partial Fourier acquisition is not shown for the T2-weighted dataset because fully-sampled calibration regions formed naturally due to high sampling density, and results were similar to those shown in Fig. 3(b).

Calibrationless data is more challenging to reconstruct than calibration-based reconstruction, and as expected, the results shown in Figs. 5 and 6 are generally worse than the results with calibration-based trajectories. However, consistent with previous results, the P-LORAKS reconstruction methods outperformed the single-channel LORAKS methods. For these cases, **S**-based P-LORAKS uniformly outperformed every other reconstruction method, again suggesting its potential and flexibility. Interestingly, the joint-sparsity reconstruction was largely unsuccessful for all datasets and sampling strategies, confirming that the LORAKS and P-LORAKS constraints are quite different from sparsity-based constraints.

It is worth noting that the best-performing sampling strategies varied substantially for these calibrationless schemes. For the T2-weighted dataset, structured sampling yielded a smaller NRMSE than poisson disk undersampling, though also yielded more visible undersampling artifacts. For the T1-weighted dataset, undersampled partial Fourier acquisition is the only

²An exception to this is in the context of nonlinear sparsely-sampled MR parameter mapping (62, 63), where structured sampling has also been preferred over random sampling.

sampling scheme that yielded reasonably accurate reconstruction results. We hypothesize that the increased sampling density offered by this sampling strategy was instrumental to this enhanced reconstruction quality relative to other sampling schemes.

DISCUSSION

Our results confirmed that the new P-LORAKS methods have consistent advantages over the previous single-channel LORAKS methods, and also highlighted the flexibility and potential usefulness of **S**-based P-LORAKS compared to **C**-based P-LORAKS (which, as described previously, is approximately equivalent to SAKÉ). Our results also confirmed that P-LORAKS can perform well, even when using sampling strategies that are unconventional for this type of regularized reconstruction problem. Specifically, structured uniform sampling is rarely proposed for low-rank matrix reconstruction, though yielded excellent reconstruction quality when used with calibration-based sampling. Similarly, randomly undersampled partial Fourier acquisition is also unconventional, but also yielded small reconstruction NRMSE values when used with **S**-based P-LORAKS.

It's worth mentioning that there are other valid strategies for processing randomly undersampled partial Fourier data. For example, in the first step of a two-step method, a method like SPIRiT, SAKÉ, or **C**-based P-LORAKS might be applied to reconstruct the missing samples within the densely sampled half of *k*-space (excluding the unsampled region of *k*-space). Subsequently, existing approaches for phase-constrained reconstruction (9,13–15,24,25,28–30,49) could be applied in a second step to reconstruct all of *k*-space based on the now “fully-sampled” half of *k*-space obtained from the first step. These strategies might be expected to yield better results than the SPIRiT and **C**-based P-LORAKS reconstructions shown in Figs. 3b, 4b, and 6b.

However, it's also worth pointing out that P-LORAKS potentially enables new forms of partial Fourier sampling that could not easily be handled using this kind of two-step processing. Fig. 7 shows an example of a novel form of partial Fourier sampling pattern. Similar to conventional partial Fourier acquisition, this sampling scheme has denser sampling on one side of *k*-space, and sparser sampling on the opposite side. In contrast to conventional partial Fourier acquisition, the denser and sparser regions are distributed according to an alternating checkerboard pattern. Since all regions of *k*-space are sampled to some degree with this sampling scheme, the use of a two-step partial Fourier reconstruction strategy would not be straightforward. Interestingly, this sampling strategy also can have certain advantages relative to conventional partial Fourier sampling. As seen in Fig. 7, the NRMSE for **S**-based P-LORAKS is smaller with checkerboard sampling than it was in Figs. 4b and 6b for more conventional partial Fourier sampling. In addition, calibrationless checkerboard sampling yielded the smallest NRMSE amongst all of the calibrationless sampling schemes we tried with the T1-weighted brain data.

While this work showed that P-LORAKS can successfully reconstruct calibrationless data, we are not advocating that calibrationless trajectories should be preferred over more conventional calibration-based approaches. For the examples we've shown in this paper, calibration-based sampling strategies still generally led to higher-quality reconstructed

images. However, for fixed undersampling factors, densely sampling one region of k -space must come at the expense of sparser sampling of other k -space regions, and the sampling density trade-offs made by conventional calibration-based approaches are not necessarily optimal within the broader context. Flexible calibrationless methods like LORAKS, P-LORAKS, and other recent approaches (43–46) provide new possibilities in sampling design that could enable better trade-offs between sampling density and the size of the calibration region. We believe that optimal sampling within this framework is a promising direction for future research. In addition, the flexibility of calibrationless methods mean that they are uniquely suited to reconstructing conventional calibration-based data in scenarios where certain samples from the calibration region have been lost (e.g., due to “spike” artifacts from phenomena like arcing) or for reconstructing data in scenarios where calibration data cannot easily be measured (e.g., dynamic imaging with high spatiotemporal resolution and time-varying support, phase, and coil sensitivities).

The results shown in this work were not directly associated with theoretical performance guarantees. While theoretical guarantees may be possible to establish based on existing low-rank matrix reconstruction theory (51), we do not believe that theoretical guarantees are critical for the practical use of P-LORAKS. This is similar to the case for the popular sparsity-based compressed sensing technique (21): while theoretical guarantees exist for certain kinds of compressed sensing problems, these theoretical guarantees are frequently inapplicable to the real reconstruction problems encountered in MRI (64). Despite the lack of guaranteed performance, compressed sensing is still a promising technique for practical applications (65, 66). However, similar to all other nonlinear reconstruction methods that lack theoretical guarantees, P-LORAKS should be evaluated using context-specific task-based validation metrics before it is deployed for routine use in any given application.

As mentioned above, the results shown in this paper were obtained using a neighborhood radius of $R = 2$ as a practical balance between reconstruction quality and computational efficiency. For reference, Fig. 8 shows P-LORAKS reconstruction results for the poisson disk sampling scheme shown in Fig. 3a as a function of both R and the rank constraints r_C and r_S . As can be seen, increasing R from 1 to 4 in this example leads to slight improvements in the best possible reconstruction quality. However, as illustrated in Fig. 9, these improvements are also associated with substantial increases in computation time. In large part, the increase in computation with larger R values is associated with the need to use larger r_C and r_S values to achieve optimal NRMSE. In addition, while we see increasing improvement with increasing R in Fig. 8, this trend will not continue indefinitely. Making R too large will begin to degrade performance as a result of overfitting, similar to the effects of setting the interpolation kernel width too large in SPIRiT or GRAPPA reconstruction (4, 5).

Figure 9 also shows that the current reconstruction times for P-LORAKS are relatively slow in their current form, on the order of minutes for a single 2D slice. However, it should be noted that our code was implemented in Matlab and not optimized for speed. In addition, the computation times shown in Fig. 9 are based on a relatively slow desktop computer (with 2.27GHz dual quadcore processors). Substantial accelerations would be expected from using optimized software running on high-performance hardware. In addition, the algorithm employed in this paper was designed to demonstrate proof-of-principle, and was not

designed specifically for fast convergence. Algorithms for low-rank matrix recovery are continuing to evolve at a rapid pace, and we believe that future algorithmic research will yield better algorithms that will make the P-LORAKS framework even more practical.

The $\mathbf{S}$ -based reconstruction results shown in this work relied heavily on smoothly-varying phase assumptions, and it is reasonable to ask what might happen when image phase changes rapidly. The issue of rapidly-varying image phase was investigated in the original LORAKS paper (22), where it was observed that faster phase variations generally lead to higher-rank $\mathbf{S}$ matrices. However, increasing the value of r_S to account for this would reduce the parsimony of the low-rank matrix model, which can limit potential acceleration factors. On the other hand, setting r_S too low can lead to image artifacts if the rank constraints are enforced too strictly (e.g., when using large values of λ_S). When using small values of λ_S to prioritize data consistency over the rank constraints, both this paper and the previous LORAKS paper (22) have shown examples where $\mathbf{S}$ -based reconstruction was successful despite the existence of abrupt small-scale local phase deviations (e.g., near blood vessels, which can be seen on close inspection of the phase images from Fig. 1).

More generally, it is important to keep in mind that the LORAKS and P-LORAKS frameworks (both $\mathbf{C}$ -based and $\mathbf{S}$ -based) depend on efficient low-rank matrix representations. Acceleration factors may be limited if the images do not have the support, phase, or parallel imaging relationships that lead to parsimonious matrix models. On the other hand, acceleration factors could also be much greater than those shown in this paper if the images of interest have appropriate characteristics.

The results shown in this paper were compared against unregularized methods like SPIRiT, but were not compared against more advanced and better-performing methods like L1-SPIRiT (39). We made this choice because we wanted to highlight the power of the P-LORAKS constraints by themselves. However, it should be noted that just like SPIRiT and LORAKS, the formulation of P-LORAKS makes it easy to include additional regularization terms. For example, Fig. 10 shows example reconstructions obtained using the same calibrationless poisson disk sampling scheme from Fig. 6a, except that reconstruction was performed using P-LORAKS and joint-sparsity regularization simultaneously. While the images shown in this figure are not necessarily of diagnostic quality, there is clearly a dramatic improvement over the reconstruction results shown in Fig. 6a. This suggests that P-LORAKS can be highly complementary to any other form of regularization-based constraints.

Finally, it is worth noting that P-LORAKS is easily extended to dynamic imaging or 3D imaging, and may be especially powerful for dynamic imaging because of the limited x - f support of most dynamic images. See (22) for discussion of such extensions in the context of LORAKS.

CONCLUSION

This paper introduced and evaluated new P-LORAKS methods for reconstructing parallel imaging data while simultaneously leveraging support and phase constraints. P-LORAKS

was demonstrated to have substantial advantages over single-channel LORAKS. In addition, S-based P-LORAKS was demonstrated to have substantial advantages over C-based P-LORAKS (which is almost the same as the SAKÉ formulation). We also showed that P-LORAKS can be used with unconventional k-space sampling schemes, and that P-LORAKS constraints can be synergistically combined with other forms of regularization. We expect this approach and its extensions to prove useful in a range of applications where it would be desirable to obtain high-quality reconstructions from highly-undersampled datasets.

Acknowledgments

This work was supported in part by NSF CAREER award CCF-1350563, NIH grant R01-NS074980, and the USC-Tsinghua Summer Undergraduate Research Program.

References

1. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: Sensitivity encoding for fast MRI. *Magn Reson Med*. 1999; 42:952–962. [PubMed: 10542355]
2. Sodickson DK, Manning WJ. Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays. *Magn Reson Med*. 1997; 38:591–603. [PubMed: 9324327]
3. Griswold MA, Jakob PM, Nittka M, Goldfarb JW, Haase A. Partially parallel imaging with localized sensitivities (PILS). *Magn Reson Med*. 2000; 44:602–609. [PubMed: 11025516]
4. Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, Haase A. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med*. 2002; 47:1202–1210. [PubMed: 12111967]
5. Lustig M, Pauly JM. SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magn Reson Med*. 2010; 65:457–471. [PubMed: 20859996]
6. Zhang J, Liu C, Moseley ME. Parallel reconstruction using null operations. *Magn Reson Med*. 2011; 66:1241–1253. [PubMed: 21604290]
7. Blaimer M, Breuer F, Mueller M, Heidemann RM, Griswold MA, Jakob PM. SMASH, SENSE, PILS, GRAPPA: How to choose the optimal method. *Top Magn Reson Imaging*. 2004; 15:223–236. [PubMed: 15548953]
8. Ying L, Liang ZP. Parallel MRI using phased array coils: Multichannel sampling theory meets spin physics. *IEEE Signal Process Mag*. 2010; 27:90–98.
9. Liang ZP, Boada F, Constable T, Haacke EM, Lauterbur PC, Smith MR. Constrained reconstruction methods in MR imaging. *Rev Magn Reson Med*. 1992; 4:67–185.
10. Plevritis SK, Macovski A. MRS imaging using anatomically based k-space sampling and extrapolation. *Magn Reson Med*. 1995; 34:686–693. [PubMed: 8544688]
11. Madore B, Glover GH, Pelc NJ. Unaliasing by Fourier-encoding the overlaps using the temporal dimension (UNFOLD), applied to cardiac imaging and fMRI. *Magn Reson Med*. 1999; 42:813–828. [PubMed: 10542340]
12. Aggarwal N, Bresler Y. Patient-adapted reconstruction and acquisition dynamic imaging method (PARADIGM) for MRI. *Inverse Probl*. 2008; 24:045015.
13. Margosian P, Schmitt F, Purdy D. Faster MR imaging: imaging with half the data. *Health Care Instrum*. 1986; 1:195–197.
14. Noll DC, Nishimura DG, Macovski A. Homodyne detection in magnetic resonance imaging. *IEEE Trans Med Imag*. 1991; 10:154–163.
15. Huang F, Lin W, Li Y. Partial Fourier reconstruction through data fitting and convolution in k-space. *Magn Reson Med*. 2009; 62:1261–1269. [PubMed: 19780148]
16. Zhao F, Noll DC, Nielsen JF, Fessler JA. Separate magnitude and phase regularization via compressed sensing. *IEEE Trans Med Imag*. 2012; 31:1713–1723.
17. Liang ZP. Spatiotemporal imaging with partially separable functions. *Proc IEEE Int Symp Biomed Imag*. 2007:988–991.

18. Haldar JP, Liang ZP. Spatiotemporal imaging with partially separable functions: A matrix recovery approach. *Proc IEEE Int Symp Biomed Imag.* 2010;716–719.
19. Lingala SG, Hu Y, DiBella E, Jacob M. Accelerated dynamic MRI exploiting sparsity and low-rank structure: k - t SLR. *IEEE Trans Med Imag.* 2011; 30:1042–1054.
20. Zhao B, Haldar JP, Christodoulou AG, Liang ZP. Image reconstruction from highly undersampled $(\mathbf{k}, t)$ -space data with joint partial separability and sparsity constraints. *IEEE Trans Med Imag.* 2012; 31:1809–1820.
21. Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn Reson Med.* 2007; 58:1182–1195. [PubMed: 17969013]
22. Haldar JP. Low-rank modeling of local k -space neighborhoods (LORAKS) for constrained MRI. *IEEE Trans Med Imag.* 2014; 33:668–681.
23. Papoulis A. Generalized sampling expansion. *IEEE Trans Circuits Syst.* 1977; CAS-24:652–654.
24. Samsonov AA, Kholmovski EG, Parker DL, Johnson CR. POCSense: POCS-based reconstruction for sensitivity encoded magnetic resonance imaging. *Magn Reson Med.* 2004; 52:1397–1406. [PubMed: 15562485]
25. Samsonov AA, Velikina J, Jung Y, Kholmovski EG, Johnson CR, Block WF. POCS-enhanced correction of motion artifacts in parallel MRI. *Magn Reson Med.* 2010; 63:1104–1110. [PubMed: 20373413]
26. Tsao J, Boesiger P, Pruessmann KP. k - t BLAST and k - t SENSE: Dynamic MRI with high frame rate exploiting spatiotemporal correlations. *Magn Reson Med.* 2003; 50:1031–1042. [PubMed: 14587014]
27. Sharif B, Derbyshire JA, Faranesh AZ, Bresler Y. Patient-adaptive reconstruction and acquisition in dynamic imaging with sensitivity encoding (PARADISE). *Magn Reson Med.* 2010; 64:501–513. [PubMed: 20665794]
28. Bydder M, Robson MD. Partial Fourier partially parallel imaging. *Magn Reson Med.* 2005; 53:1393–1401. [PubMed: 15906299]
29. Willig-Onwuachi JD, Yeh EN, Grant AK, Ohliger MA, McKenzie CA, Sodickson DK. Phase-constrained parallel MR image reconstruction. *J Magn Reson.* 2005; 176:187–198. [PubMed: 16027017]
30. Lew C, Pineda AR, Clayton D, Spielman D, Chan F, Bammer R. SENSE phase-constrained magnitude reconstruction with iterative phase refinement. *Magn Reson Med.* 2007; 58:910–921. [PubMed: 17969127]
31. Block KT, Uecker M, Frahm J. Undersampled radial MRI with multiple coils. iterative image reconstruction using a total variation constraint. *Magn Reson Med.* 2007; 57:1086–1098. [PubMed: 17534903]
32. Raj A, Singh G, Zabih R, Kressler B, Wang Y, Schuff N, Weiner M. Bayesian parallel imaging with edge-preserving priors. *Magn Reson Med.* 2007; 57:8–21. [PubMed: 17195165]
33. Liang D, Liu B, Wang JJ, Ying L. Accelerating SENSE using compressed sensing. *Magn Reson Med.* 2009; 62:1574–1584. [PubMed: 19785017]
34. Otazo R, Kim D, Axel L, Sodickson DK. Combination of compressed sensing and parallel imaging for highly accelerated first-pass cardiac perfusion MRI. *Magn Reson Med.* 2010; 64:767–776. [PubMed: 20535813]
35. Guerquin-Kern M, Häberlin M, Pruessmann KP, Unser M. A fast wavelet-based reconstruction method for magnetic resonance imaging. *IEEE Trans Med Imag.* 2011; 30:1649–1660.
36. Ramani S, Fessler JA. Parallel MR image reconstruction using augmented Lagrangian methods. *IEEE Trans Med Imag.* 2011; 30:694–706.
37. Ye X, Chen Y, Lin W, Huang F. Fast MR image reconstruction for partially parallel imaging with arbitrary k -space trajectories. *IEEE Trans Med Imag.* 2011; 30:575–585.
38. Chaâri L, Pesquet JC, Benazza-Benyahia A, Ciuciu P. A wavelet-based regularized reconstruction algorithm for SENSE parallel MRI with applications to neuroimaging. *Med Image Anal.* 2011; 15:185–201. [PubMed: 21106436]
39. Murphy M, Alley M, Demmel J, Keutzer K, Vasanawala S, Lustig M. Fast ℓ_1 -SPIRiT compressed sensing parallel imaging MRI: scalable parallel implementation and clinically feasible runtime. *IEEE Trans Med Imag.* 2012; 31:1250–1262.

40. Weller DS, Polimeni JR, Grady L, Wald LL, Adalsteinsson E, Goyal VK. Denoising sparse images from GRAPPA using the nullspace method. *Magn Reson Med*. 2012; 68:1176–1189. [PubMed: 22213069]
41. Knoll F, Clason C, Bredies K, Uecker M, Stollberger R. Parallel imaging with nonlinear reconstruction using variational penalties. *Magn Reson Med*. 2012; 67:34–41. [PubMed: 21710612]
42. She H, Chen RR, Liang D, DiBella EVR, Ying L. Sparse BLIP: Blind iterative parallel imaging reconstruction using compressed sensing. *Magn Reson Med*. 2014; 71:645–660. [PubMed: 23508781]
43. Majumdar A, Ward RK. Calibration-less multi-coil MR image reconstruction. *Magn Reson Imag*. 2012; 30:1032–1045.
44. Chen C, Li Y, Huang J. Calibrationless parallel MRI with joint total variation regularization. *Proc MICCAI*. 2013:106–114.
45. Trzasko JD, Manduca A. CLEAR: Calibration-free parallel imaging using locally low-rank encouraging reconstruction. *Proc Int Soc Magn Reson Med*. 2012:517.
46. Shin PJ, Larson PEZ, Ohliger MA, Elad M, Pauly JM, Vigneron DB, Lustig M. Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion. *Magn Reson Med*. 2014; 72:959–970. [PubMed: 24248734]
47. Gol D, Potter LC. Denoising in parallel imaging via structured low-rank matrix approximation. *Proc Int Soc Magn Reson Med*. 2013:3822.
48. Uecker M, Lai P, Murphy MJ, Virtue P, Elad M, Pauly JM, Vasanawala SS, Lustig M. ESPIRiT – an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA. *Magn Reson Med*. 2014; 71:990–1001. [PubMed: 23649942]
49. Blaimer M, Gutberlet M, Kellman P, Breuer FA, Köstler H, Griswold MA. Virtual coil concept for improved parallel MRI employing conjugate symmetric signals. *Magn Reson Med*. 2009; 61:93–102. [PubMed: 19097211]
50. Halдар, JP. Technical Report USC-SIPI-414. University of Southern California; Los Angeles, CA: 2014. Low-rank modeling of local k -space neighborhoods (LORAKS): Implementation and examples for reproducible research.
51. Recht B, Fazel M, Parrilo PA. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev*. 2010; 52:471–501.
52. Cheung KF, Marks RJ II. Imaging sampling below the Nyquist density without aliasing. *J Opt Soc Am A*. 1990; 7:92–105.
53. Zhuo J, Halдар JP. P-LORAKS: Low-rank modeling of local k -space neighborhoods with parallel imaging data. *Proc Int Soc Magn Reson Med*. 2014:745.
54. Lustig M. Post-cartesian calibrationless parallel imaging reconstruction by structured low-rank matrix completion. *Proc Int Soc Magn Reson Med*. 2011:483.
55. Hunter DR, Lange K. A tutorial on MM algorithms. *Am Stat*. 2004; 58:30–37.
56. Jacobsen MW, Fessler JA. An expanded theoretical treatment of iteration-dependent majorize-minimize algorithms. *IEEE Trans Image Process*. 2007; 16:2411–2422. [PubMed: 17926925]
57. Hestenes MR, Stiefel E. Methods of conjugate gradients for solving linear systems. *J Res Natl Bur Stand*. 1952; 49:409–436.
58. Cadzow JA. Signal enhancement – a composite property mapping algorithm. *IEEE Trans Acoust, Speech, Signal Process*. 1988; 36:49–62.
59. Halдар JP, Liang ZP. Joint reconstruction of noisy high-resolution MR image sequences. *Proc IEEE Int Symp Biomed Imag*. 2008:752–755.
60. Halдар JP, Wedeen VJ, Nezamzadeh M, Dai G, Weiner MW, Schuff N, Liang ZP. Improved diffusion imaging through SNR-enhancing joint reconstruction. *Magn Reson Med*. 2013; 69:277–289. [PubMed: 22392528]
61. Nayak KS, Nishimura DG. Randomized trajectories for reduced aliasing artifact. *Proc Int Soc Magn Reson Med*. 1998:670.
62. Halдар JP, Hernando D, Liang ZP. Super-resolution reconstruction of MR image sequences with contrast modeling. *Proc IEEE Int Symp Biomed Imag*. 2009:266–269.

63. Velikina JV, Alexander AL, Samsonov A. Accelerating MR parameter mapping using sparsity-promoting regularization in parametric dimension. *Magn Reson Med*. 2013; 70:1263–1273. [PubMed: 23213053]
64. Haldar JP, Hernando D, Liang ZP. Compressed-sensing MRI with random encoding. *IEEE Trans Med Imag*. 2011; 30:893–903.
65. Zhang T, Chowdhury S, Lustig M, Barth RA, Alley MT, Grafendorfer T, Calderon PD, Robb FJL, Pauly JM, Vasanawala SS. Clinical performance of contrast enhanced abdominal pediatric MRI with fast combined parallel imaging compressed sensing reconstruction. *J Magn Reson Imag*. 2013; 40:13–25.
66. Sharma SD, Fong CL, Tzung BS, Law M, Nayak KS. Clinical image quality assessment of accelerated magnetic resonance neuroimaging using compressed sensing. *Invest radiol*. 2013; 48:638–645. [PubMed: 23538890]

Figure 1. Magnitude (top rows) and phase (bottom rows) images corresponding to individual channels from (a) a 32-channel T1-weighted brain dataset and (b) a 12-channel T2-weighted brain dataset. The images from each channel each have limited spatial support and slowly varying spatial image phase.

Figure 2.

Normalized root-mean squared errors (NRMSEs) obtained when the LORAKS and P-LORAKS matrices are approximated by matrices having a small total rank. For P-LORAKS (solid lines), total rank refers to the rank at which the SVD was truncated for $\mathbf{C}^P$ or $\mathbf{S}^P$. For LORAKS (dotted lines), total rank is equal to L times the rank at which the SVDs for each $\mathbf{C}_\ell$ and $\mathbf{S}_\ell$ matrix were truncated. Results are shown for (a,b) $\mathbf{C}^P$ and $\mathbf{C}_\ell$ matrices and (c–d) $\mathbf{S}^P$ and $\mathbf{S}_\ell$ matrices, corresponding to (a,c) the 32-channel T1-weighted brain dataset, and (b,d) the 12-channel T2-weighted brain dataset.

Figure 3.

Reconstruction results for calibration-based reconstruction of the T2-weighted brain dataset. (a) Random sampling. (b) Random partial Fourier sampling. (c) Structured sampling. The top rows show reconstructed images using a linear grayscale (normalized so that image intensities are in the range from 0 to 1), while the bottom rows show error images using the indicated colorscale (which ranges from 0 to 0.2 to highlight small errors). NRMSE values are shown underneath each reconstruction, with the best NRMSE values highlighted with bold text.

Figure 4.

Reconstruction results for calibration-based reconstruction of the T1-weighted brain dataset. (a) Random sampling. (b) Random partial Fourier sampling. (c) Structured sampling. The top rows show reconstructed images using a linear grayscale (normalized so that image intensities are in the range from 0 to 1), while the bottom rows show error images using the indicated colorscale (which ranges from 0 to 0.2 to highlight small errors). NRMSE values are shown underneath each reconstruction, with the best NRMSE values highlighted with bold text.

Figure 5.

Reconstruction results for calibrationless reconstruction of the T2-weighted brain dataset. (a) Random sampling. (b) Structured sampling. The top rows show reconstructed images using a linear grayscale (normalized so that image intensities are in the range from 0 to 1), while the bottom rows show error images using the indicated colorscale (which ranges from 0 to 0.2 to highlight small errors). NRMSE values are shown underneath each reconstruction, with the best NRMSE values highlighted with bold text.

Figure 6.

Reconstruction results for calibrationless reconstruction of the T1-weighted brain dataset. (a) Random sampling. (b) Random partial Fourier sampling. (c) Structured sampling. The top rows show reconstructed images using a linear grayscale (normalized so that image intensities are in the range from 0 to 1), while the bottom rows show error images using the indicated colorscale (which ranges from 0 to 0.2 to highlight small errors). NRMSE values are shown underneath each reconstruction, with the best NRMSE values highlighted with bold text.

Figure 7. Reconstruction results for checkerboard sampling. (a) Calibration-based sampling. (b) Calibrationless sampling. The top rows show reconstructed images using a linear grayscale (normalized so that image intensities are in the range from 0 to 1), while the bottom rows show error images using the indicated colorscale (which ranges from 0 to 0.2 to highlight small errors). NRMSE values are shown underneath each reconstruction, with the best NRMSE values highlighted with bold text.

Figure 8.

P-LORAKS reconstruction results as a function of the neighborhood size R and the rank parameters r_C and r_S for (a) C-based P-LORAKS and (b) S-based P-LORAKS. The images on the left show the reconstructions with smallest NRMSE for each R value. The top rows show reconstructions using a linear grayscale (normalized so that image intensities are in the range from 0 to 1), while the bottom rows show error images using the indicated colorscale (which ranges from 0 to 0.2 to highlight small errors). NRMSE values are shown underneath each reconstruction, with the best NRMSE values highlighted with bold text. The plots on the right show how the NRMSE evolves as a function of rank for each different R .

Figure 9.

Computation times as a function of R for the results shown in Fig. 8. For each R , we plot the P-LORAKS reconstruction time corresponding to the value of r_C or r_S that achieved the smallest NRMSE.

Figure 10.

Reconstruction results obtained using P-LORAKS together with joint sparsity regularization. These reconstructions were obtained using the same calibrationless poisson disk sampling pattern shown in Fig. 6a, and have much smaller NRMSE values than any of the reconstructions from Fig. 6a. The top row shows reconstructed images using a linear grayscale (normalized so that image intensities are in the range from 0 to 1), while the bottom row shows error images using the indicated colorscale (which ranges from 0 to 0.2 to highlight small errors). NRMSE values are shown underneath each reconstruction, with the best NRMSE values highlighted with bold text.