Published in final edited form as:

IEEE Trans Med Imaging. 2019 July ; 38(7): 1545–1558. doi:10.1109/TMI.2019.2896180.

OEDIPUS: An Experiment Design Framework for Sparsity-Constrained MRI

Justin P. Haldar [Senior Member, IEEE]

Signal and Image Processing Institute, Ming Hsieh Department of Electrical Engineering,
University of Southern California, Los Angeles, CA, 90089, USA

Daeun Kim [Student Member, IEEE]

Signal and Image Processing Institute, Ming Hsieh Department of Electrical Engineering,
University of Southern California, Los Angeles, CA, 90089, USA

Abstract

This paper introduces a new estimation-theoretic framework for experiment design in the context of MR image reconstruction under sparsity constraints. The new framework is called OEDIPUS (Oracle-based Experiment Design for Imaging Parsimoniously Under Sparsity constraints), and is based on combining the constrained Cramér-Rao bound with classical experiment design techniques. Compared to popular random sampling approaches, OEDIPUS is fully deterministic and automatically tailors the sampling pattern to the specific imaging context of interest (i.e., accounting for coil geometry, anatomy, image contrast, etc.). OEDIPUS-based experiment designs are evaluated using retrospectively subsampled *in vivo* MRI data in several different contexts. Results demonstrate that OEDIPUS-based experiment designs have some desirable characteristics relative to conventional MRI sampling approaches.

Keywords

Experiment Design; Constrained MRI Reconstruction; Sparsity Constraints; Compressed Sensing

I. Introduction

Slow data acquisition speed is a natural consequence of acquiring MRI data with high signal-to-noise ratio (SNR), high spatial resolution, and desirable image contrast characteristics. Slow acquisitions are problematic because they are expensive, limit scanner throughput and temporal resolution, and can be uncomfortable for scan subjects. In practice, these concerns lead to the use of shorter-duration experiments that represent undesirable trade-offs between image resolution and contrast when data is sampled according to the conventional sampling theorem [1]. To mitigate this problem, the research community has been working for decades to develop image reconstruction methods that enable conventional sampling requirements to be relaxed [2]. While many such image reconstruction methods have been proposed over the years, this paper focuses on a popular and widely-used class of methods based on sparsity constraints [3].

While successful applications of sparsity-constrained MRI are widespread throughout the literature, the optimal design of sampling patterns remains a longstanding open problem. Because there is already a large amount of literature on this topic, it would be impractical for us to comprehensively review every contribution. Instead, our literature review focuses only on what we believe is the most relevant previous work.

Some of the earliest sparsity-based methods sample low-resolution data at the conventional Nyquist rate, and use sparsity constraints to recover the unsampled high-frequency data and enable super-resolution reconstruction [4]. More recently, insight provided by compressed sensing theory has inspired the emergence and popularity of pseudo-random Fourier sampling schemes [5].

Random Fourier sampling for MRI has often been motivated by theoretical performance guarantees for random sampling that depend on mathematical concepts of incoherence and/or restricted isometry [6], [7]. Within the MRI community, a fairly common misconception is that random sampling is required for good sparsity-constrained MRI. In reality, random Fourier sampling is neither necessary nor sufficient for good performance [5], [8], [9], particularly in the case of multi-channel imaging for which the reconstruction problem is often not actually underdetermined (and is instead just ill-posed). Most modern approaches have come to rely on heuristic variations of random Fourier sampling that are often tuned for each new application context based on extensive empirical trial-and-error testing [5], [10]–[12]. While these kinds of random sampling variations often work adequately well, their good performance is not guaranteed from a theoretical perspective [8], empirical tuning of random sampling characteristics is often onerous, and it remains unclear whether the results are near-optimal.

This paper describes a principled method called OEDIPUS (oracle-based Experiment Design for Imaging Parsimoniously Under Sparsity constraints) for designing sampling patterns for sparsity-constrained MRI. Compared to randomization-based approaches, OEDIPUS is fully deterministic and automatically tailors the sampling pattern to the specific imaging context (accounting for coil geometry, anatomy, image contrast, etc.) of interest without the need for manual interaction. A preliminary description of this work was previously presented in [13].

The OEDIPUS framework is inspired by classical experiment design approaches. Experiment design is a well-developed subfield of statistics [14]–[16], and experiment design tools have been previously used to improve the efficiency of a wide variety of MRI experiments. These methods are perhaps most visible within the MRI literature in the context of selecting flip angles, repetition times, echo times, and related pulse sequence parameters that influence image contrast. Such methods have already had a major impact in the context of quantitative MRI parameter estimation for applications such as spectroscopy [17], water-fat separation [18], dynamic contrast enhanced imaging [19], diffusion MRI [20], relaxometry [21], and multi-parametric mapping [22], among many others.

However, the use of experiment design techniques to optimize k-space sampling patterns (as done in OEDIPUS) has not been as widely explored, in part due to the computational

complexity of applying experiment design techniques to image reconstruction problems (which are often substantially larger in scale than MRI parameter estimation problems). There are a few notable exceptions. Marseille *et al.* [23] performed optimization of phase encoding positions in 1D under a heuristic surrogate model of an MR image. Reeves *et al.* [24]–[28] optimized phase encoding positions in 2D assuming that the image had limited support in the spatial domain, with perfect *a priori* knowledge of the image support, while later authors explored a similar approach in the context of dynamic imaging [29]. Xu *et al.* [30] performed optimization of phase encoding positions in 1D for the case of parallel imaging using the SENSE model [31]. Samsonov [32] optimized the projection angles for radial imaging under the SENSE model. Haldar *et al.* [33] compared different spatiotemporal (k-t) sampling patterns for parameter mapping in the presence of prior knowledge of the image contrast model. Levine and Hargreaves [34] optimized 2D sampling patterns under support constraints and parallel imaging constraints with the SENSE model. Importantly, none of these approaches directly incorporated transform domain sparsity constraints into the experiment design problem.

The only previous sparsity-based MR experiment design work we are aware of is by Seeger *et al.* [35]. While this work has the strongest similarity to OEDIPUS among existing methods from the literature, it also has substantial differences, including being developed with a completely different context in mind. In particular, the approach described by Seeger *et al.* is intended for real-time adaptive sampling design, in which each subsequent k-space sampling location is chosen greedily on-the-fly using real-time feedback from the data that has already been acquired for that scan. In contrast, OEDIPUS is a non-adaptive and subject-independent approach to experiment design, which enables offline prior computation of the sampling pattern and reuse of the same sampling scheme for different subjects. Another major difference is that Seeger *et al.* assume a Bayesian formulation in which the sparse image coefficients are modeled using the Laplace distribution as a statistical prior, while OEDIPUS is a non-Bayesian approach that does not use any statistical priors. Instead, OEDIPUS takes a deterministic approach based on the constrained Cramér-Rao bound (CRB) [36], [37].

The rest of this paper is organized as follows. Section II introduces notation, reviews the CRB for linear inverse problems with Gaussian noise, and reviews optimal experiment design approaches that aim to minimize the CRB. Section III describes the general OEDIPUS framework as an extension of the methods from Section II. This section also describes the specific OEDIPUS implementation choices we have used in the remainder of the paper. Section IV describes empirical evaluations of OEDIPUS with comparisons against other common k-space sampling methods from the literature. Discussion and conclusions are presented in Section V.

II. Review of “Classical” Experiment Design

A. Linear Inverse Problems with Gaussian Noise

The data acquisition process in MRI is often formulated using a standard finite-dimensional linear model [38]

$$\mathbf{d} = \mathbf{A}\mathbf{f} + \mathbf{n}, \quad (1)$$

where $\mathbf{f} \in \mathbb{C}^N$ are the coefficients of a finite-dimensional representation of the image, $\mathbf{d} \in \mathbb{C}^M$ is the vector of measured data samples, the system matrix $\mathbf{A} \in \mathbb{C}^{M \times N}$ provides a linear model of the MRI data acquisition physics (e.g., potentially including Fourier encoding, sensitivity encoding, and physical effects like field inhomogeneity), and $\mathbf{n} \in \mathbb{C}^M$ is modeled as zero-mean circularly-symmetric complex Gaussian noise. Without loss of generality, we will assume that noise whitening has been applied such that the entries of $\mathbf{n}$ are assumed to be independent and identically distributed with variance $\sigma^2 = 1$.

The objective of MRI reconstruction methods is to estimate the original image $\mathbf{f}$ based only on the noisy measured data $\mathbf{d}$ and prior knowledge of the noise statistics and imaging model $\mathbf{A}$. We will denote such an estimator as $\hat{\mathbf{f}}(\mathbf{d})$.

B. The Cramér-Rao Bound for the Linear Model

The CRB is a theoretical lower bound on the covariance of an unbiased estimator [39]. Assume for now that data has been acquired according to Eq. (1), that $\mathbf{A}$ has full-column rank, and that we have an unbiased estimator $\hat{\mathbf{f}}(\mathbf{d})$, i.e., $E[\hat{\mathbf{f}}(\mathbf{d})] = \mathbf{f}$, where $E[\cdot]$ denotes statistical expectation. The CRB says that

$$\text{cov}(\hat{\mathbf{f}}(\mathbf{d})) \geq (\mathbf{A}^H \mathbf{A})^{-1}, \quad (2)$$

where $\text{cov}(\hat{\mathbf{f}}(\mathbf{d})) \in \mathbb{C}^{N \times N}$ is the covariance matrix of the estimator $\hat{\mathbf{f}}(\mathbf{d})$, and the $\geq$ symbol is used to signify an inequality relationship with respect to the Loewner ordering of positive semidefinite matrices [16]. To be precise, Eq. (2) signifies that $\text{cov}(\hat{\mathbf{f}}(\mathbf{d})) - (\mathbf{A}^H \mathbf{A})^{-1}$ will always be a positive semidefinite matrix. As a result, the CRB matrix $(\mathbf{A}^H \mathbf{A})^{-1}$ represents the smallest possible covariance matrix that can be achieved using an unbiased estimator, since we must always have that $\mathbf{u}^H \text{cov}(\hat{\mathbf{f}}(\mathbf{d})) \mathbf{u} \geq \mathbf{u}^H (\mathbf{A}^H \mathbf{A})^{-1} \mathbf{u}$ for every possible choice of $\mathbf{u} \in \mathbb{C}^N$. To illustrate the importance of this relationship, consider the case where we select $\mathbf{u}$ to be the j th canonical basis vector for $\mathbb{C}^N$, i.e., the j th entry of $\mathbf{u}$ is $[\mathbf{u}]_j = 1$, with $[\mathbf{u}]_n = 0$ for all $n \neq j$. This implies that $\text{var}([\hat{\mathbf{f}}(\mathbf{d})]_j) \geq [(\mathbf{A}^H \mathbf{A})^{-1}]_{jj}$, where $\text{var}([\hat{\mathbf{f}}(\mathbf{d})]_j)$ is the scalar variance of the estimator for the j th entry of $\mathbf{f}$. Note that, because we have chosen j arbitrarily in this illustration, we have that the N diagonal elements of the CRB matrix each provide a fundamental lower bound on the variance for the corresponding N entries of the unbiased estimator $\hat{\mathbf{f}}(\mathbf{d})$.

Importantly, the CRB is a function of the matrix $\mathbf{A}$, which means that it can be possible to use the CRB to compare the SNR-efficiency of different experiments. For readers familiar with the parallel MRI literature, it is worth mentioning that in the case of SENSE reconstruction, the CRB is an unnormalized version of the SENSE g-factor [31], [40], which

is already well known and widely used for its usefulness in comparing different k-space sampling strategies in parallel MRI.

In addition, the CRB is a reconstruction-agnostic lower bound, in the sense that it is completely independent from the actual method used to estimate $\mathbf{f}$. However, it should also be noted that the CRB is only valid for unbiased estimators, while different estimation-theoretic tools are necessary to characterize biased estimators. This fact is particularly relevant to this paper, since most sparsity-constrained reconstruction methods use some form of regularization procedure, and regularized image reconstruction will generally produce biased results [41]. Nevertheless, minimizing the CRB can still be a useful approach to experiment design even if a biased estimator will ultimately be used for reconstruction. In some sense, minimizing the CRB ensures that the data acquisition procedure represented by the matrix $\mathbf{A}$ will encode as much information as possible about $\mathbf{f}$.

C. Optimal Experiment Design using the CRB

Methods for minimizing the CRB have been studied for many years in both statistics [14]–[16] and signal processing [42], [43]. There are generally two classes of popular experiment design approaches: continuous approaches and discrete approaches. Optimal continuous designs are often easier to compute than optimal discrete designs, but are approximate in nature and do not translate easily to the constraints of practical MRI experiments because they do not account for the fact that real experiments acquire a finite, integer number of samples. As a result, we will focus on discrete designs in this paper.

The discrete version of optimal experiment design is often formulated as follows: given a set of P candidate measurements (with $P > M$), choose a subset of M measurements in a way that minimizes the CRB. Formally, let $\mathbf{a}_p \in \mathbb{C}^{1 \times N}$ for $p = 1, \dots, P$ represent the P potential rows of the optimal $\mathbf{A}$ matrix from Eq. (1). Given the corresponding set of potential measurement indices $\Gamma = \{1, \dots, P\}$, the optimal design problem is to select a subset $\Omega \subset \Gamma$ of cardinality M to ensure that the resulting CRB is as small as possible. This leads to the optimal design problem

$$\Omega^* = \arg \min_{\substack{\Omega \subset \Gamma \\ |\Omega| = M}} J\left(\left(\sum_{m \in \Omega} \mathbf{a}_m^H \mathbf{a}_m\right)^{-1}\right), \quad (3)$$

where $J(\cdot): \mathbb{C}^{N \times N} \rightarrow \mathbb{R}$ is a user-selected functional used to measure the size of the CRB.

An ideal choice of $J(\cdot)$ would ensure that for arbitrary positive semidefinite matrices $\mathbf{C} \in \mathbb{C}^{N \times N}$ and $\mathbf{D} \in \mathbb{C}^{N \times N}$, having $J(\mathbf{C}) \leq J(\mathbf{D})$ would imply that $\mathbf{C} \geq \mathbf{D}$, and therefore that $\mathbf{D}$ is uniformly smaller than $\mathbf{C}$. Unfortunately, the Loewner ordering is only a partial ordering, and not all positive semidefinite matrices are comparable. This means that there are generally no universally-optimal experiment designs, and it is necessary to choose $J(\cdot)$ subjectively [16].

Out of many options [16], one of the most popular choices is the *average variance* criterion (also known as *A-optimality*):

$$J\left(\left(\sum_{m \in \Omega} \mathbf{a}_m^H \mathbf{a}_m\right)^{-1}\right) \triangleq \text{Trace}\left[\left(\sum_{m \in \Omega} \mathbf{a}_m^H \mathbf{a}_m\right)^{-1}\right], \quad (4)$$

which takes its name from the fact that the cost function is a lower bound on the sum of elementwise variances, i.e.,

$$\text{Trace}\left[\left(\sum_{m \in \Omega} \mathbf{a}_m^H \mathbf{a}_m\right)^{-1}\right] \leq \sum_{n=1}^N \text{var}([\hat{\mathbf{f}}(\mathbf{d})]_n) \quad (5)$$

for every possible unbiased estimator $\hat{\mathbf{f}}(\mathbf{d})$. Since the mean-squared error (MSE) is proportional to the sum of elementwise variances for any unbiased estimator, Eq. (5) also provides a CRB-based lower bound for the MSE. For the sake of concreteness and without loss of generality, we will focus on the A-optimality criterion in this paper, though note that OEDIPUS can also easily be used with some of the other cost functionals that are popular in the experiment design literature.

Given this choice of $J(\cdot)$, the optimization problem from Eq. (3) is still nontrivial to solve because it is a nonconvex integer programming problem, and global optimality may require exhaustive search through the set of candidate solutions. This is generally infeasible from a practical computational perspective. Specifically, there are P -choose- M possible designs Ω , and this will be a very large number in most scenarios of interest. As a result, Eq. (3) is often minimized using heuristic algorithms that can provide high-quality solutions, but which cannot guarantee global optimality. Popular heuristics include the use of Monte Carlo methods (which randomly construct potential sampling pattern candidates for evaluation) or greedy algorithms (which add, subtract, or exchange candidate measurements to/from the current design in a sequential greedy fashion) [14], [15], [42], [43].

While these kinds of algorithms have been used previously for k-space trajectory design in several image reconstruction contexts [24]–[28], [30], certain computational problems arise when the CRB is applied to conventional image reconstruction approaches (without sparsity constraints). In particular, the CRB matrix is of size $N \times N$, which can be very large if N is chosen equal to the number of voxels in a reasonably-sized image. For example, for a 256×256 image, the CRB matrix will have nearly 4.3×10^9 entries, which will occupy more than 34 GB of memory if each entry is stored in a standard double precision floating point number format. While storing a matrix of this size in memory is feasible on modern high-end computers, other computations (e.g., inverting the matrix) are generally prohibitively expensive. Another problem is that the assumptions underlying the CRB will be violated whenever $N > M$, which means that traditional experiment design approaches are not straightforward to apply for the design of undersampled experiments. These issues are mitigated within the OEDIPUS framework as described in the sequel.

III. Oracle-based Experiment Design for Imaging Parsimoniously Under Sparsity constraints

A. Theoretical components of OEDIPUS

The previous section considered CRB-based design under the standard unconstrained linear model, but modifications are necessary to inform experiment design under sparsity constraints. To describe OEDIPUS, we will assume that it is known in advance that the image vector $\mathbf{f}$ possesses a certain level of sparsity in a known transform domain, i.e., $\|\Psi\mathbf{f}\|_0$

$\leq S$ where $\Psi \in \mathbb{C}^{Q \times N}$ is a known sparsifying transform, S is the known sparsity level, and the ℓ_0 -“norm” $\|\cdot\|_0$ counts the number of nonzero entries of its vector argument. It should be noted that the assumptions that the signal is strictly sparse and that the value of S is known in advance are not conventional in sparsity-constrained MRI reconstruction, though we need them to construct the CRB for this case.

In this work, we will make the additional assumption that Ψ is left-invertible such that $\mathbf{f} = \Psi^\dagger \Psi \mathbf{f}$ for all $\mathbf{f} \in \mathbb{C}^N$, where † is used to denote the matrix pseudoinverse. In the presence of sparsity constraints, this allows us to rewrite Eq. (1) in terms of the vector of transform coefficients $\mathbf{c} = \Psi \mathbf{f} \in \mathbb{C}^Q$ as:

$$\mathbf{d} = \mathbf{A} \Psi^\dagger \mathbf{c} + \mathbf{n}. \quad (6)$$

Recently, Ben-Haim and Eldar [37] derived a version of the constrained CRB that is applicable to Eq. (6) under the constraint that $\|\mathbf{c}\|_0 = S$. In particular, they derived that the CRB in this case is equal to the CRB of the *oracle estimator*, i.e., the estimator that is additionally bestowed with perfect knowledge of which of the S entries of $\mathbf{c}$ are non-zero.

The oracle-based version of the problem can be viewed as a simple unconstrained linear model, and thus fits directly within the framework described in the previous section. In particular, let $\tilde{\mathbf{c}} \in \mathbb{C}^S$ denote the vector of nonzero entries from the $\mathbf{c}$ vector, with the two vectors related through $\mathbf{c} = \mathbf{U} \tilde{\mathbf{c}}$, where $\mathbf{U} \in \mathbb{R}^{Q \times S}$ is a matrix formed by concatenating the S columns of the $Q \times Q$ identity matrix corresponding to the non-zero entries of $\mathbf{c}$. This matrix can be viewed as performing a simple zero-filling operation (i.e., putting the nonzero entries of $\tilde{\mathbf{c}}$ in their correct corresponding positions and leaving the remaining entries equal to zero). Given this notation, the oracle-based version of the problem is then written as

$$\mathbf{d} = \mathbf{A} \Psi^\dagger \mathbf{U} \tilde{\mathbf{c}} + \mathbf{n}. \quad (7)$$

From here, we can use the results of the previous section to easily derive the oracle-based CRB for $\hat{\tilde{\mathbf{c}}}(\mathbf{d})$:

$$\text{cov}(\hat{\tilde{\mathbf{c}}}(\mathbf{d})) \geq \left(\mathbf{U}^H (\Psi^\dagger)^H \mathbf{A}^H \mathbf{A} \Psi^\dagger \mathbf{U} \right)^{-1}, \quad (8)$$

such that the oracle-based CRB for $\hat{\mathbf{c}}(\mathbf{d}) = \mathbf{U}\hat{\mathbf{c}}(\mathbf{d})$ is:

$$\text{cov}(\hat{\mathbf{c}}(\mathbf{d})) \geq \mathbf{U} \left(\mathbf{U}^H (\mathbf{\Psi}^\dagger)^H \mathbf{A}^H \mathbf{A} \mathbf{\Psi}^\dagger \mathbf{U} \right)^{-1} \mathbf{U}^H, \quad (9)$$

and the oracle-based CRB for $\hat{\mathbf{f}}(\mathbf{d}) = \mathbf{\Psi}^\dagger \hat{\mathbf{c}}(\mathbf{d})$ is:

$$\text{cov}(\hat{\mathbf{f}}(\mathbf{d})) \geq \mathbf{\Psi}^\dagger \mathbf{U} \left(\mathbf{U}^H (\mathbf{\Psi}^\dagger)^H \mathbf{A}^H \mathbf{A} \mathbf{\Psi}^\dagger \mathbf{U} \right)^{-1} \mathbf{U}^H (\mathbf{\Psi}^\dagger)^H. \quad (10)$$

Given these sparsity-constrained CRBs, it is easy to formulate and solve an experiment design problem using exactly the same design methods described for the linear model in the previous section. The CRB expressions given above will be valid whenever the $M \times S$ matrix $\mathbf{A}\mathbf{\Psi}^\dagger\mathbf{U}$ has full column rank, which is still feasible for undersampled experiments where $M < N$ (though we still need to satisfy $S \leq M$).

In addition to enabling the use of sparsity constraints and undersampling in experiment design, the expressions in Eqs. (8)-(10) are generally also substantially easier to work with from a computational point of view. In particular, it should be noted that the matrices being inverted in these expressions are only of size $S \times S$, rather than the $N \times N$ matrices that needed to be inverted in the previous case. This can lead to massive computational complexity advantages since we usually have $S \ll N$. For example, for a 256×256 image with $S = 0.15N$, the matrix to be inverted in this case would have roughly 2.25% the memory footprint of the matrix to be inverted in the $N \times N$ case. In this example, this is the difference between a 0.8 GB matrix and a 34 GB matrix.

Even though the final CRB matrices in Eqs. (9) and (10) are generally larger (respectively of size $Q \times Q$ and $N \times N$) than the $S \times S$ matrix being inverted in Eq. (8), it is not hard to show that the matrices in Eqs. (8) and (9) always have the same trace value. In addition, the trace value for the matrix from Eq. (10) will also have this same trace value if $\mathbf{\Psi}$ is a unitary transform (e.g., an orthonormal wavelet transform) and will be proportional to this same trace value if $\mathbf{\Psi}$ corresponds to a tight frame representation [44] (e.g., the undecimated Haar wavelet transform [45] or the discrete curvelet transform [46]). As a result, it often suffices to work directly with the small $S \times S$ matrix from Eq. (8) when performing experiment design using the A-optimality design criterion.

B. Practical Implementation for MRI Applications

1) Constructing the OEDIPUS optimization problem: For the sake of concreteness and without loss of generality, we will describe our implementation of OEDIPUS assuming that we are applying the A-optimality design criterion to the sparsity-constrained CRB from Eq. (8). Given an application context, there are several ingredients we need to select prior to the experiment design process.

The first choice that needs to be made is the sparsifying transform Ψ . While some groups have explored the use of data-dependent sparsifying transforms (e.g., [47], [48]), it is by far more common for the sparsifying transform to be chosen in advance based on past empirical experience with a given MRI application. Our description will focus on the conventional case where Ψ is fixed and provided in advance.

The second choice we need to make relates to the set of candidate measurement vectors $\mathbf{a}_p \in \mathbb{C}^{1 \times N}$ for $p = 1, \dots, P$. This choice is relatively straightforward in the case of ideal single-channel Fourier encoding. Given a field-of-view (FOV) size and the number of voxels within the FOV (which we assume are fixed, and dictated by the specific application context of interest), the candidate measurement vectors will often have entries (assuming an N -dimensional voxel-based finite series expansion of the image [38]) of the form:

$$[\mathbf{a}_p]_{1n} = b_{pn} \exp(-i2\pi \mathbf{k}_p \cdot \mathbf{r}_n), \quad (11)$$

where b_{pn} is a weighting coefficient that depends on the specific choice of voxel basis function (common choices are to use either a Dirac delta voxel function in which $b_{pn} = 1$ or a rectangular voxel basis function in which case b_{pn} takes the form of a sinc function in k -space [38]), $\mathbf{r}_n$ is the spatial coordinate vector specifying the center position of the n th voxel (the common choice is to choose the voxel centers to lie on a rectilinear grid), and $\mathbf{k}_p$ is the p th candidate k -space measurement position. The parameters b_{pn} and $\mathbf{r}_n$ are parameters of the image model that are generally chosen in advance, and which will be invariant to different experiment design choices. As a result, selection of the full set of the candidate measurements reduces to selection of the set of $\mathbf{k}_p$ k -space sampling locations. The set of candidate k -space sampling locations will be sequence-dependent, but typical choices might include the set of all possible k -space sampling locations from a “fully-sampled” Cartesian grid, a “fully-sampled” radial trajectory, a “fully-sampled” spiral trajectory, etc. In each of these cases, we can define “full-sampling” either based on the conventional Nyquist rate, or we can also consider schemes that sample more densely than the Nyquist rate if we want to consider a larger range of possibilities. There is a trade-off here – the more options we consider, the better our optimized schemes can be, though this comes at the expense of increased computational complexity.

Some additional considerations are required if the model is designed to incorporate sensitivity encoding, field inhomogeneity, radiofrequency encoding, or other similar effects. For illustration, in the case of parallel imaging with the SENSE model [31], the measurements take the more general form:

$$[\mathbf{a}_p]_{1n} = c_p(\mathbf{r}_n) b_{pn} \exp(-i2\pi \mathbf{k}_p \cdot \mathbf{r}_n), \quad (12)$$

where $c_p(\mathbf{r}_n)$ is the value of the sensitivity profile at spatial location $\mathbf{r}_n$ for the coil used in the p th measurement.¹ The main complication in applying OEDIPUS to the SENSE model is that, while the Fourier encoding model is independent of the person or object being

imaged, the coil sensitivity profiles often vary slightly from subject to subject, and this means that ideally we should design optimal experiments on a subject-by-subject basis. However, subject-by-subject design is not conducive to the kind of offline subject-independent experiment design that we would like to use OEDIPUS to address. Similar complications exist because of subject-dependent effects for other advanced modeling schemes that account for B_0 inhomogeneity, transmit B_1 inhomogeneity, etc. To mitigate the issue of subject-dependent observations, our formulation for the SENSE case (and similarly for other subject-dependent acquisition models) will assume that we have a set of T different representative sensitivity maps that are appropriate for the given application context, i.e.,

$$[\mathbf{a}_p^t]_{1n} = c_p^t(\mathbf{r}_n) b_{pn} \exp(-i2\pi \mathbf{k}_p \cdot \mathbf{r}_n), \quad (13)$$

for $t = 1, \dots, T$, and will optimize the experiment design with respect to the ensemble of these representative cases. We will formalize this approach mathematically later in this section.

The remaining ingredients that need to be chosen are the sparsity level S and the matrix $\mathbf{U}$ that specifies the transform-domain positions of the non-zero coefficients. However, real MRI images are only approximately sparse rather than exactly sparse, and the locations of the significant transform-domain coefficients will generally vary from subject to subject. To address these issues and enable subject-independent experiment design, we will assume that we are given a set of K exemplar images acquired from a given imaging context, and similar to the previous case, will identify S_k and $\mathbf{U}_k$ values for $k = 1, \dots, K$ that are appropriate for each case. Ultimately, we will attempt to find a design that is optimal with respect to the ensemble of exemplars. In this context, we first apply the sparsifying transform Ψ to each exemplar, and explore sparse approximation quality by hard-thresholding the transform domain coefficients with different choices of S_k . The value of S_k is then chosen as the smallest value that gives adequate sparse approximation quality for that exemplar. Note that “adequate quality” is necessarily subjective and application-dependent. Given the choice of S_k for each exemplar image, we can obtain corresponding $\mathbf{U}_k \in \mathbb{R}^{Q \times S_k}$ matrices by identifying the locations of the S_k largest transform coefficients.

We are now almost ready to perform experiment design, except that we have not specified how to handle optimization with respect to the ensemble of T different measurement models and K different exemplar sparsity constraints. There are a few different options for obtaining a unified treatment. One typical approach would be to minimize the average-case CRB across each of these different conditions, which results in the following optimization problem:

¹In parallel MRI, the same k-space location is measured simultaneously through all of the coils that are present in the receiver array. We have adopted a notation in which each individual measurement is indexed by a separate p value, even if some measurements are required to be sampled simultaneously.

$$\Omega^* = \arg \min_{\substack{\Omega \subset \Gamma \\ |\Omega| = M}} \sum_{k=1}^K \sum_{t=1}^T \text{Trace}[\mathbf{C}_{kt}(\Omega)], \quad (14)$$

where

$$\mathbf{C}_{kt}(\Omega) \triangleq \left(\sum_{m \in \Omega} \mathbf{U}_k^H (\Psi^\dagger)^H (\mathbf{a}_m^t)^H \mathbf{a}_m^t \Psi^\dagger \mathbf{U}_k \right)^{-1}. \quad (15)$$

Another typical choice would be to optimize the worst-case CRB, which leads to the following minimax optimization problem:

$$\Omega^* = \arg \min_{\substack{\Omega \subset \Gamma \\ |\Omega| = M}} \max_{(k,t)} \text{Trace}[\mathbf{C}_{kt}(\Omega)]. \quad (16)$$

Many other choices are also possible and may be preferable in certain contexts, though we leave such an exploration for future work. The choice between minimizing the average-case and the worst-case in experiment design is subjective in nature, and will depend on context-dependent factors. For example, in a large population study with a neuroscientific objective, experimenters may be able to tolerate the situation where a few subjects have low-quality images if the average image quality for the remaining subjects is very high on average. On the other hand, in an emergency medical situation, it may be preferable that every individual image meets certain quality standards, even if that means that the average-case image quality is lower as a result. In the case where $K = T = 1$, both Eq. (14) and (16) lead to identical optimization problems.

2) Solving the OEDIPUS optimization problem: Given a choice of objective function, it remains to find an experiment design that is as close to optimal as possible. While there are many potential ways to minimize these cost functionals, the implementation we have used for this paper is based on a modified sequential backward-selection (SBS) procedure. The standard SBS procedure [42], [43] starts with a hypothetical experiment that collects all P candidate measurements, and then identifies the individual measurement that is the least important to the CRB and deletes it in a greedy fashion. This procedure is then iterated, deleting candidate measurements one-by-one until only M measurement vectors remain. SBS is greedy and will not generally find the globally optimal solution, but it is a simple deterministic algorithm that typically obtains local minima that compare very well in ultimate quality compared to alternative optimization approaches. Additional theoretical comments about the characteristics of SBS can be found in [42], and methods that have the potential to improve on SBS results are discussed in [43].

While the conventional SBS approach may be well-suited to certain applications, it is not immediately applicable to many MRI contexts in which sets of measurements must be acquired simultaneously. For example, in parallel imaging, it does not make any sense to

delete a k-space measurement made with one coil without also deleting the same k-space measurement from the remaining coils. Similarly, it would not make sense to delete one k-space measurement from a given readout without also deleting the other measurements from the same readout. As a result, it is important to be able to adapt SBS to consider sets of candidate measurements rather than individual candidate measurements. Specifically (and similar to [27], [28], [30]), we will assume that our set of P measurements is grouped into L disjoint subsets Ξ_ℓ for $\ell = 1, \dots, L$, and that all measurements within a given subset Ξ_ℓ must either be acquired simultaneously or not acquired at all. We will further assume the practical scenario in which all of these measurement sets have the same cardinality, i.e., $|\Xi_\ell| = C$ for some integer C . This leads to the following modified SBS procedure for OEDIPUS:

1. For $k = 1, \dots, K$ and $t = 1, \dots, T$, construct the sparsity-constrained CRB matrix corresponding to the use of all P candidate measurements:

$$\mathbf{C}_{kt}^{\text{full}} = \left(\sum_{p=1}^P \mathbf{U}_k^H (\Psi^\dagger)^H (\mathbf{a}_p^t)^H \mathbf{a}_p^t \Psi^\dagger \mathbf{U}_k \right)^{-1}. \quad (17)$$

As noted in Sec. III-A, this matrix is often not very large (i.e., of size $S \times S$), and will therefore fit easily within the memory capacities of most modern computers if S is chosen suitably small.

2. For each $\ell = 1, \dots, L$, compute the cost-function value (e.g., the average-case cost from Eq. (14) or the worstcase cost from Eq. (16)) for an experiment that uses all P measurements except the C candidate measurements included in Ξ_ℓ .
3. The cost-function values are compared for each choice of ℓ and the measurement set whose removal causes the smallest degradation in the cost function value is then removed from the set of measurement candidates. Specifically, if i is the index for the removed set of measurement candidates, set $\Gamma = \Gamma \setminus \Xi_i$ and correspondingly set $P = P - C$. Similarly, remove Ξ_i from the set of measurement candidates such that $L = L - 1$.
4. Repeat Steps 1-3 until the desired number of measurements is achieved, i.e., stop when $|\Gamma| = M$.

The main computational complexity associated with this modified SBS approach is that Step 2 requires TLK matrix inversions at each iteration. However, as thoroughly discussed by [42], it should be noted that it is easy and computationally-efficient to use the Sherman-Morrison-Woodbury matrix inversion lemma [49] to use the matrix from Step 1 to quickly compute the matrices needed for Step 2. Specifically, if we use $\mathbf{B}_{\ell t} \in \mathbb{C}^{C \times N}$ to denote the matrix formed by stacking together the potential rows associated with Ξ_ℓ for the t th acquisition model and let $\tilde{\mathbf{B}}_{\ell t}^k = \mathbf{B}_{\ell t} \Psi^\dagger \mathbf{U}_k \in \mathbb{C}^{C \times S}$, then the CRB matrix $\mathbf{C}_{kt}^{\text{full} \setminus \Xi_\ell}$ that would be obtained by keeping all P measurements except those from Ξ_ℓ can be written as

$$\begin{aligned} \mathbf{C}_{kt}^{\text{full} \setminus \Xi_{\ell}} &= \left(\sum_{\substack{p=1 \\ p \notin \Xi_{\ell}}}^P \mathbf{U}_k^H (\boldsymbol{\Psi}^{\dagger})^H (\mathbf{a}_p^t)^H \mathbf{a}_p^t \boldsymbol{\Psi}^{\dagger} \mathbf{U}_k \right)^{-1} \\ &= \mathbf{C}_{kt}^{\text{full}} + \mathbf{G}_{\ell t}^k \left(\mathbf{I}_C - \tilde{\mathbf{B}}_{\ell t}^k \mathbf{G}_{\ell t}^k \right)^{-1} (\mathbf{G}_{\ell t}^k)^H, \end{aligned} \quad (18)$$

where $\mathbf{I}_C$ is the $C \times C$ identity matrix and $\mathbf{G}_{\ell t}^k = \mathbf{C}_{kt}^{\text{full}} (\tilde{\mathbf{B}}_{\ell t}^k)^H$. Note that this final expression only requires the inversion of a $C \times C$ matrix, which is a much faster computation than inverting an $S \times S$ matrix if $C \ll S$ (which it usually will be). In addition, it should be noted that except for the first time that $\mathbf{C}_{kt}^{\text{full}}$ is computed using Eq. (17) in Step 1, all subsequent computations of $\mathbf{C}_{kt}^{\text{full}}$ for Step 1 can be performed recursively from the previous value using Eq. (18).

Further computational simplifications are possible by noting that the trace values required for computing Step 2 are also easily evaluated in a recursive fashion:

$$\begin{aligned} \text{Trace} \left[\mathbf{C}_{kt}^{\text{full} \setminus \Xi_{\ell}} \right] &= \text{Trace} \left[\mathbf{C}_{kt}^{\text{full}} \right] + \\ &\quad \text{Trace} \left[\mathbf{G}_{\ell t}^k \left(\mathbf{I}_C - \tilde{\mathbf{B}}_{\ell t}^k \mathbf{G}_{\ell t}^k \right)^{-1} (\mathbf{G}_{\ell t}^k)^H \right]. \end{aligned} \quad (19)$$

IV. Empirical Evaluations and Comparisons

In the following subsections, we illustrate the performance and flexibility of OEDIPUS relative to traditional sampling design methods in several settings.

A. 2D T2-weighted Brain Imaging

In a first set of evaluations, we designed sampling patterns for the context of 2D T2-weighted brain imaging, based on an acquisition protocol that is routinely used at our neuroscience imaging center. Data was acquired on the same scanner from seven different healthy volunteers on different days.

All acquisition parameters were identical for all subjects. Imaging data was acquired with a TSE pulse sequence on a 3T scanner using a 12-channel receiver coil. We used a fully-sampled 256×160 Nyquist-rate Fourier acquisition matrix, corresponding to a $210\text{mm} \times 131\text{mm}$ FOV, with a 3.5mm slice thickness. Other parameters include $\text{TE}=88\text{ms}$ and $\text{TR}=10\text{sec}$. This data was whitened using standard methods [50], and receiver coil sensitivity maps were obtained using a smoothness-regularized estimation formulation [51]. Gold standard reference images were obtained by applying SENSE reconstruction [31], [50] to the fully sampled data. These gold standard images for two subjects are shown in Fig. 1(a,b).

Sampling patterns were designed using OEDIPUS for two different imaging scenarios: single-channel imaging (a simple Fourier transform with no sensitivity encoding) and 12-channel imaging (including coil sensitivity maps). In both cases, the design was performed with respect to a single slice for simplicity, although it would have been straightforward to use volumetric training data (or, to reduce computational complexity, perhaps a small number of selected slices representing the range of expected variations across the volume of interest) for cases where broader volume coverage may be desired. We denote the OEDIPUS sampling patterns designed for the single-channel and multi-channel cases as SCO and MCO, respectively. Since this is a 2D imaging scenario, OEDIPUS was constrained so that all samples from the same readout line must be sampled simultaneously, meaning that undersampling was only performed in 1D along the phase encoding dimension.

For illustration (and without loss of generality), we show OEDIPUS sampling patterns that were designed for a single exemplar (i.e., $K = 1$), and in the multi-channel case, for a single set of sensitivity maps (i.e., $T = 1$). Specifically, OEDIPUS was performed based on the transform-domain sparsity pattern and sensitivity maps estimated from the first subject. Sparse approximation of this dataset was obtained by keeping the top 15% largest-magnitude transform coefficients (i.e., $S = 0.15N$). The sparsifying transform Ψ was chosen to be the 3-level wavelet decomposition using a Daubechies-4 wavelet filter, whose transform coefficients are shown for the first two subjects in Fig. 1(c). As can be seen, while the images and the exact distributions of wavelet coefficients are distinct for each of the two datasets, the wavelet-domain support patterns for both images still have strong qualitative similarities to one another. This similarity is one of the key assumptions that enables OEDIPUS to design sampling patterns based on exemplar data instead of requiring subject-specific optimization.

MCO sampling patterns were designed for the acceleration factors of $R = 2, 3$, and 4, and an SCO sampling pattern was designed for $R = 2$. We also attempted to design SCO sampling patterns for the acceleration factors of $R = 3$ and 4, but did not obtain meaningful OEDIPUS results because the matrix inversions needed for computing the CRB became numerically unstable. This instability occurred because the matrices became nearly-singular, which is an indication that acceleration factors of 3 or 4 may be too aggressive for unbiased estimation in this particular scenario. Specifically, near-singularity implies that unbiased estimators would have very large variance for certain parameters of interest (i.e., a singular matrix is associated with infinite variance [52]). This is consistent with our empirical results (not shown due to space constraints) for which $R = 3$ or 4 yields unacceptably poor results in this scenario for every case we attempted.

For comparison against SCO and MCO, we also considered two classical undersampling patterns: uniform undersampling and random undersampling. Uniform undersampling is a deterministic sampling pattern that is frequently used in parallel imaging [31], and obtains phase encoding lines that are uniformly spaced on a regular lattice in k-space. Random sampling is often advocated for sparsity-constrained reconstruction [5], [8], [10]–[12]. In this case, we used a common approach in which the central 16 lines of k-space were fully sampled, and the remaining portions of k-space were acquired with uniform Poisson disc random sampling [53]. It is worth mentioning that both the uniform and random sampling

schemes considered in this work are expected to be better-suited to multi-channel scenarios than they are to single-channel cases. In particular, uniform undersampling is not usually advocated in single-channel settings due to the well-known coherent aliasing associated with lattice sampling, and single-channel random undersampling usually employs a more gradual variation in sampling density rather than the abrupt binary change in sampling density that we've implemented. Nevertheless, due to the large number of subjects and scenarios we consider in this paper and the effort associated with manually tuning more nuanced sampling density variations by trial-and-error, we have opted for the sake of simplicity to utilize the same type of simple random sampling design for both single- and multi-channel cases. To avoid the possibility of getting a particularly bad random design, we created ten independently generated realizations of Poisson disc sampling, and report the best quantitative results.

The sampling patterns we obtained are shown in Fig. 2, and demonstrate some interesting features. In particular, we observe that both SCO and MCO generate variable density sampling patterns, which is consistent with the popularity of variable density sampling approaches in the literature [5], [8], [10]–[12]. However, unlike standard randomization-based variable density approaches, we observe that OEDIPUS automatically tailors the sampling pattern to the specific imaging context, without the need for manual interaction. This can be observed from the fact that OEDIPUS automatically produces different sampling density variations for single-channel and multi-channel contexts, with SCO sampling the center of k-space more densely than MCO.

To evaluate these different sampling patterns in a practical reconstruction context, we also performed sparsity-constrained reconstructions of retrospectively-undersampled single-channel and multi-channel datasets derived from the gold standard reference images. We might expect to see especially good empirical reconstruction results for OEDIPUS for the first subject, since the SCO and MCO sampling patterns were designed based on the characteristics of that data. In this case, we are testing on the same data that we trained on, and our results may be overly optimistic due to the common problem of overfitting. On the other hand, the empirical results from the other subjects may be more interesting, since these will demonstrate the capability to design sampling patterns that generalize beyond the training data.

We performed two different types of sparsity-constrained reconstruction. In the first case, we applied standard ℓ_1 -regularized reconstruction [5] according to

$$\hat{\mathbf{f}} = \arg \min_{\mathbf{f} \in \mathbb{C}^N} \|\mathbf{A}\mathbf{f} - \mathbf{d}\|_2^2 + \lambda \|\Psi\mathbf{f}\|_1. \quad (20)$$

This formulation will encourage sparsity in the 3-level Daubechies-4 wavelet decomposition of the image. The regularization parameter λ was set to a small number (i.e., $\lambda = 0.01$) to promote data consistency. We might expect OEDIPUS to perform particularly well in this case, because the sparsifying transform used to design SCO and MCO sampling patterns is matched to the sparsifying transform being used for reconstruction.

While MR images are known to possess wavelet sparsity, it has also been observed that wavelet-based sparse regularization often results in undesirable high-frequency oscillation artifacts. As a result, many authors (e.g., [5], [8]) have advocated the use of total variation (TV) regularization [54] to either augment or replace wavelet regularization. TV regularization is designed to impose sparsity constraints on the finite differences of a reconstructed image. Unfortunately, our OEDIPUS formulation is not easy to apply to TV regularization, because TV is not associated with a left-invertible transform. However, some authors have postulated that CRB-based experiment design may still be beneficial even in cases where the model used to compute the CRB is different from the model used for image reconstruction [23], as long as both models capture the behavior of similar image features of interest. To test this, we also performed TV-regularized reconstruction. Specifically, we solved the same optimization problem as in Eq. (20), but replacing the ℓ_1 -norm with a TV penalty while leaving the regularization parameter λ the same.

In both cases, optimization was performed using a multiplicative half-quadratic algorithm [55], [56] (also known as “Iteratively Reweighted Least Squares” or “Lagged Diffusivity”). Reconstruction quality was measured quantitatively using the normalized root-mean-squared error (NRMSE), defined as $\text{NRMSE} = \|\hat{\mathbf{f}} - \mathbf{f}^*\|_2 / \|\mathbf{f}^*\|_2$, where $\mathbf{f}^*$ is the gold standard.

The empirical NRMSE results for the reconstruction of single-channel data for the first two subjects are shown in Table I and for the remaining subjects in Supplementary Table SI, with illustrative reconstructed images shown in Supplementary Fig. S1,² while NRMSE results for the reconstruction of multi-channel data are shown for the first two subjects in Table II and for the remaining subjects in Supplementary Table SII, with illustrative reconstructed images shown in Supplementary Fig. S2.

For single-channel data, the SCO sampling pattern worked the best for both wavelet- and TV-based reconstruction. The differences in NRMSE between SCO and the other approaches were often substantial, with random sampling being the next best approach. Uniform undersampling had very poor performance in this case, although this is unsurprising because uniform undersampling is known to be a poor choice for sparsity-constrained reconstruction of single-channel data [5].

In the multi-channel case, the MCO sampling worked the best for both wavelet- and TV-based reconstruction at high acceleration factors (i.e., $R = 3$ and 4). We observed that uniform random sampling worked the best with $R = 2$, although MCO follows close behind and all four sampling schemes work similarly well at this low acceleration rate. The good performance of uniform undersampling for multi-channel data at low-acceleration rates occurs despite the fact that it had a slightly worse CRB than MCO, although is consistent with theoretical g-factor arguments from the parallel imaging literature [31], [50].

We also observe that the SCO pattern seems to work better with single-channel data and MCO works better with multi-channel data, as should be expected. In addition, our results

²This paper has supplementary downloadable material provided by the authors, available at <http://ieeexplore.ieee.org> in the supplementary files/multimedia tab.

show that the OEDIPUS sampling patterns designed for one subject generally yielded good performance when applied to similar data from other subjects.

B. 3D T1-weighted Brain Imaging

In a second set of evaluations, we designed sampling patterns for the context of 3D T1-weighted brain imaging, also based on an acquisition protocol that is routinely used at our neuroscience imaging center. Data was acquired on the same scanner from one healthy volunteer and one subject with a chronic stroke lesion. The healthy volunteer was scanned more than 13 months after the stroke subject.

Most acquisition parameters were identical for both subjects. Imaging data was acquired with an MPRAGE pulse sequence on a 3T scanner using a 12-channel receiver coil. Along the two phase-encoding dimensions, we used a fully-sampled 156×192 Fourier acquisition matrix, corresponding to a $208\text{mm} \times 256\text{mm}$ FOV. Because this FOV is much larger than necessary along the left-right dimension, all image reconstructions and experiment designs were performed with respect to a 156×126 voxel grid corresponding to a smaller $208\text{ mm} \times 168\text{ mm}$ FOV. The third (readout) dimension was reconstructed with 1mm resolution, and we extracted one slice along this dimension to be used in our sampling investigation. Other parameters include $\text{TI}=800\text{ms}$, $\text{TE}=3.09\text{ms}$, $\text{TR}=2530\text{ms}$, and flip angle = 10° .

While both datasets were acquired using the same 12-channel array coil, the data for the stroke subject had been hardware-compressed by the scanner down to 4 virtual channels, while we had access to the original 12-channels of data for the healthy subject. To reduce the differences in sensitivity encoding information content between the two datasets, we used standard software coil-compression methods [57], [58] to reduce the data for the healthy subject to 5 virtual channels. From the coil-compressed data, we obtained receiver coil sensitivity maps and gold standard reference images following the same procedures as described in the previous section. These gold standard images are shown in Fig. 3.

Similar to the previous case, we used OEDIPUS to design sampling patterns for both single-channel imaging and multi-channel imaging. Since this is a 3D imaging scenario, we gave OEDIPUS the freedom to undersample in 2D along both phase encoding dimensions. For illustration (and without loss of generality), OEDIPUS sampling patterns were designed using the same approach described in the previous section (i.e., $K = 1$, $T = 1$, $S = 0.15N$, and using the 3-level Daubechies-4 wavelet decomposition) based on the single-slice from the healthy subject. The wavelet coefficients for both reference images are shown in Fig. 3(c). As before, we observe strong qualitative similarities in the wavelet-domain support patterns for both images, despite the presence of a lesion in one image but not in the other.

MCO and SCO sampling patterns were both designed for several acceleration factors between $R = 2$ and 8. Compared to the previous case, the use of 2D undersampling instead of 1D undersampling enables higher acceleration factors for SCO, although high ill-posedness still prevented us from computing meaningful SCO results for $R = 7$ and 8.

Similar to the previous case, we compared SCO and MCO sampling designs against uniform and random sampling. For uniform undersampling, we used CAIPI-type uniform lattice

sampling designs that are known to have good parallel imaging characteristics [59]. For random sampling, we used a standard approach in which the central 16×16 region of k-space was fully sampled, and the remaining portions of k-space were acquired with uniform Poisson disc random sampling. As before, we created ten independently generated realizations of Poisson disc sampling, and report the best results.

The sampling patterns we obtained are shown in Fig. 4. Similar to the previous case, we observe that both SCO and MCO generate variable density sampling patterns, but that SCO and MCO have very different sampling density characteristics from one another. In particular, SCO frequently samples the center of k-space more densely than MCO does, while MCO samples peripheral k-space more densely than SCO. Interestingly, we observe that SCO with $R = 2$ performs full sampling of central k-space. Because the k-space acquisition grid was designed for a larger FOV while reconstruction is performed over a smaller FOV, this corresponds to oversampling the center of k-space (i.e., at a rate higher than the conventional Nyquist rate). On the other hand, at higher acceleration factors, we observe that both the SCO and MCO approaches sample the center of k-space less densely than conventional variable-density random sampling schemes (where the center of k-space is frequently sampled at the Nyquist rate) that are popular for sparsity-constrained MRI. In addition, the SCO pattern visually appears to have more structure than is typical in random designs. This illustrates the fact that not only is OEDIPUS adaptive to the imaging context, but it can also explore a broader range of sampling patterns than traditional sampling design techniques.

Similar to before, single-slice images were reconstructed from retrospectively undersampled data in both single-channel and multi-channel reconstruction contexts using either ℓ_1 -regularization of the wavelet transform coefficients or TV regularization. The empirical NRMSE results for the reconstruction of single-channel data are shown in Table III with illustrative reconstructed images shown in Supplementary Fig. S3, while NRMSE results for the reconstruction of multi-channel data are shown in Table IV with illustrative reconstructed images shown in Supplementary Fig. S4.

Although the use of 2D undersampling admits the use of higher acceleration factors than were achieved in the previous case, our observations are still largely consistent with the observations from Section IV-A. In particular, we observe that SCO generally yields the best performance in the single-channel case for both wavelet and TV-regularization, and for both the healthy subject (which OEDIPUS was trained for) and the stroke subject (which was not used for training). Similarly, MCO generally yields the best performance in the multi-channel case in all of these different scenarios. There are some exceptions to these general rules, e.g., MCO slightly outperforms SCO in many of the cases with $R = 3$ and single-channel data, while MCO is outperformed by uniform (CAIPI) sampling in the multi-channel case with lower acceleration factors (despite having slightly better CRB values). This result is not necessarily surprising, since we have no guarantees that the OEDIPUS sampling design will yield optimal results for any given dataset. However, it is noteworthy that in each of these cases where the most relevant OEDIPUS strategy was outperformed by an alternative approach, the OEDIPUS design is not far behind the leader. In addition, at

higher and higher acceleration factors, we observe that the most relevant OEDIPUS strategy consistently yields the best performance (by wider margins as acceleration increases).

The sampling patterns we've shown were optimized for a single slice near the center of the brain, although it's reasonable to expect that the sampling patterns will still be reasonably good for the neighboring slices from the volume that have similar image features and coil sensitivity characteristics. Reconstruction quality performance across the 3D volume is shown for different sampling patterns in Supplementary Figs. S5 and S6, and match our expectations. In particular, performance is very consistent for the slices that are in a similar region to the slice used for training, although performance sometimes has larger deviations in farther slices that have more distinct characteristics (e.g., image regions that have very different support characteristics, or which have different textural features, etc.). Note that even though our results generalize relatively well across the volume, we still do not recommend using single-slice training data if volumetric reconstruction is ultimately of interest, as this would likely lead to suboptimal results that are biased towards one anatomical region more than others. We believe that it would be much preferred to optimize sampling across exemplar volumes, or at least across a few representative slices reflecting the variations across the volume. The general principal of OEDIPUS is that the training data should be relatively well-matched to the anatomical regions that are of highest importance to the specific application for which the sampling pattern is designed.

V. Discussion and Conclusions

This paper introduced the OEDIPUS framework to designing sampling patterns for sparsity-constrained MRI. Unlike standard randomization based approaches (which have some major theoretical shortcomings [8]), our proposed approach is guided by estimation theory, is deterministic, and automatically tailors itself to the given imaging context at hand.

OEDIPUS was evaluated empirically at many different acceleration factors in several very different imaging contexts. While these empirical evaluations were limited in scope and should not be overinterpreted, our results were very consistent in each case, and suggest several potential conclusions. First, results seem to confirm that the OEDIPUS-derived sampling patterns can be used to design good sampling patterns that outperform some of the classical sampling approaches. While it may be possible to manually tune a variable density random sampling pattern that achieves similar or even better performance to the OEDIPUS patterns, OEDIPUS has the advantage that the sampling density is deterministic and automatically tailored to the specific scenario of interest, without the need for the user to exhaustively explore a range of different random sampling distributions or realizations. Second, sampling patterns designed based on data from one subject appear to generalize well to other subjects. This kind of generalization capability is important, since it would mean that it is not necessary to create unique sampling patterns for each subject, and that offline sampling pattern design may indeed be a viable strategy. And finally, results suggest that OEDIPUS sampling schemes designed based on wavelet sparsity may generalize well to the case of TV reconstruction. Given these results, we believe that OEDIPUS has the potential to be a powerful tool for sampling design across a range of different applications.

The OEDIPUS framework is based on the constrained CRB, which is an estimation theoretic measure for the amount of information provided by a given experiment design. It should be noted that other metrics have also previously been used to measure the expected “goodness” of a sampling design in sparse MRI reconstruction, including mutual incoherence and restricted isometry constants [5], [8]. These other metrics have been previously considered because, when these mathematical constants are good enough, it is possible to guarantee the performance of sparsity-based reconstruction methods [6], [7], and these guarantees will hold for arbitrary sparse images, regardless of their transform domain support characteristics. In this sense, these are measures that can be used to characterize performance for the “worst-case” transform domain support patterns. However, unlike OEDIPUS, these measures are agnostic to “typical” realizations of the inverse problem. Designing for the worst case is likely to sacrifice performance in typical cases. In addition, the conditions under which these kinds of metrics could be used to provide strong performance guarantees are often not met in practical MRI applications [8].

While the implementation described in this paper made use of the greedy SBS method for simplicity, we expect that even better OEDIPUS performance could be further enhanced using appropriate exchange algorithms [43]. With good initializations (e.g., initializing with CAIPI), these kinds of algorithms could also be used to decrease computation time and increase the chances of finding good local optima relative to SBS.

Since OEDIPUS is intended for offline use (i.e., experiment designs only need to be computed once for a given imaging context and can then be reused ad infinitum), its associated computation time is perhaps less important than for other kinds of optimization procedures. Nevertheless, it is worth noting that the computation time can be relatively mild or relatively burdensome depending on the imaging context. At least for the SBS algorithm, the number of candidate sampling subsets L has a very big impact on computation time. For example, in the 2D T2-weighted illustration, we started with a relatively small number of candidates to consider ($L = 160$) and optimization completed in a few hours on a standard desktop computer. On the other hand, we started with a much larger number of candidates in the 3D T1-weighted illustration ($L = 19,656$), and optimization in this case can require many days of computation. The problem of large L is not new in the experiment design literature, and various approaches have been suggested to reduce L . For example, in a different context, Gao and Reeves [27] have suggested the use of simpler classes of sampling patterns (e.g., periodic nonuniform sampling) with fewer candidates to consider. This improves computational complexity though sacrifices optimality.

While our results were shown in the context of static MRI imaging, our approach generalizes in natural ways to other imaging contexts where sparsity constraints may be useful. For example, this approach has the potential to be useful in MRI parameter mapping with sparsity constraints, for which CRBs have already been described [60], and we have performed some initial investigations of OEDIPUS-like sampling [61] in the context of high-dimensional diffusion-relaxation correlation spectroscopic imaging [62], [63].

Although OEDIPUS was formulated to address sparsity constraints, there are natural generalizations to other imaging constraints that are recently popular in the literature. For

example, CRBs have been derived for low-rank matrix recovery [64], structured low-rank matrix recovery [65], and low-rank tensor recovery [66], and it would be straightforward in principle to adapt the OEDIPUS approach to design sampling patterns for popular recent MRI reconstruction methods that leverage such low-rank modeling assumptions [67]–[74]. This promises to be an interesting direction for further research.

The OEDIPUS formulation we used in this work was designed to minimize the average variance, which is a parameter related to NRMSE. Since NRMSE has some well-known limitations [75], the exploration of alternative criteria appears worthwhile.³ We believe that focused application-specific testing is especially important, since the characteristics of nonlinear methods can be hard to generalize from one context to the next [8], and because important context-specific image features may not correlate with NRMSE.

As a final note, we believe it is worth reiterating that, despite common misconceptions, sparsity-constrained MRI generally does not enjoy any meaningful theoretical performance guarantees [8]. OEDIPUS is no exception, and while OEDIPUS was observed to work well in several practical scenarios, we make no claim that OEDIPUS will always be optimal.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported in part by the National Science Foundation (NSF) under CAREER award CCF-1350563 and the National Institutes of Health (NIH) under grants NIH-R21-EB022951, NIH-R01-MH116173, NIH-R01-NS074980, and NIH-R01-NS089212. Computation for some of the work described in this paper was supported by the University of Southern California's Center for High-Performance Computing (<http://hpcc.usc.edu/>).

References

- [1]. Edelstein WA, Mahesh M, and Carrino JA, "MRI: Time is dose—and money and versatility," *J. Am. Coll. Radiol.*, vol. 7, pp. 650–652, 2010. [PubMed: 20678736]
- [2]. Liang Z-P, Boada F, Constable T, Haacke EM, Lauterbur PC, and Smith MR, "Constrained reconstruction methods in MR imaging," *Rev. Magn. Reson. Med.*, vol. 4, pp. 67–185, 1992.
- [3]. Lustig M, Donoho DL, Santos JM, and Pauly JM, "Compressed sensing MRI," *IEEE Signal Process. Mag.*, vol. 25, pp. 72–82, 2008.
- [4]. Liang Z-P, Haacke EM, and Thomas CW, "High-resolution inversion of finite Fourier transform data through a localised polynomial approximation," *Inverse Probl.*, vol. 5, pp. 831–847, 1989.
- [5]. Lustig M, Donoho D, and Pauly JM, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn. Reson. Med.*, vol. 58, pp. 1182–1195, 2007. [PubMed: 17969013]
- [6]. Donoho DL, Elad M, and Temlyakov VN, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Trans. Inf. Theory*, vol. 52, pp. 6–18, 2006.
- [7]. Candès EJ, "The restricted isometry property and its implications for compressed sensing," *C. R. Acad. Sci. Paris, Ser. I*, vol. 346, pp. 589–592, 2008.
- [8]. Haldar JP, Hernando D, and Liang Z-P, "Compressed-sensing MRI with random encoding," *IEEE Trans. Med. Imag.*, vol. 30, pp. 893–903, 2011.

³During the review of this paper and after public dissemination of OEDIPUS [13], [76], a preprint appeared that describes an alternative method for optimizing k-space sampling patterns based on exemplar data [77]. This approach has many similarities to OEDIPUS but is able to address other error metrics besides NRMSE, at the cost of higher computational complexity.

- [9]. Adcock B, Hansen AC, Poon C, and Roman B, "Breaking the coherence barrier: A new theory for compressed sensing," *Forum Math. Sigma*, vol. 5, p. e4, 2017.
- [10]. Wang Z and Arce GR, "Variable density compressed image sampling," *IEEE Trans. Image Process*, vol. 19, pp. 264–270, 2010. [PubMed: 19775971]
- [11]. Knoll F, Clason C, Diwok C, and Stollberger R, "Adapted random sampling patterns for accelerated MRI," *Magn. Reson. Mater. Phy*, vol. 24, pp. 43–50, 2011.
- [12]. Zijlstra F, Viergever MA, and Seevinck PR, "Evaluation of variable density and data-driven k-space undersampling for compressed sensing magnetic resonance imaging," *Invest. Radiol*, vol. 51, pp. 410–419, 2016. [PubMed: 26674209]
- [13]. Haldar JP and Kim D, "OEDIPUS: Towards optimal deterministic k-space sampling for sparsity-constrained MRI," in *Proc. Int. Soc. Magn. Reson. Med*, 2017, p. 3877.
- [14]. Federov VV, *Theory of Optimal Experiments*. New York: Academic Press, 1972.
- [15]. Silvey SD, *Optimal Design: An Introduction to the Theory for Parameter Estimation*. New York: Chapman and Hall, 1980.
- [16]. Pukelsheim F, *Optimal Design of Experiments*. New York: John Wiley & Sons, 1993.
- [17]. Cavassila S, Deval S, Huegen C, van Ormondt D, and Graveron-Demilly D, "Cramér-Rao bounds: an evaluation tool for quantitation," *NMR Biomed*, vol. 14, pp. 278–283, 2001. [PubMed: 11410946]
- [18]. Pineda AR, Reeder SB, Wen Z, and Pelc NJ, "Cramér-Rao bounds for three-point decomposition of water and fat," *Magn. Reson. Med*, vol. 54, pp. 625–635, 2005. [PubMed: 16092102]
- [19]. De Naeyer D, De Deene Y, Ceelen WP, Segers P, and Verdonck P, "Precision analysis of kinetic modelling estimates in dynamic contrast enhanced MRI," *Magn. Reson. Mater. Phy*, vol. 24, pp. 51–66, 2011.
- [20]. Alexander DC, "A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features," *Magn. Reson. Med*, vol. 60, pp. 439–448, 2008. [PubMed: 18666109]
- [21]. Jones JA, Hodgkinson P, Barker AL, and Hore PJ, "Optimal sampling strategies for the measurement of spin-spin relaxation times," *J. Magn. Reson. B*, vol. 113, pp. 25–34, 1996.
- [22]. Zhao B, Haldar JP, Liao C, Ma D, Griswold MA, Setsompop K, and Wald LL, "Optimal experiment design for magnetic resonance fingerprinting: Cramer-Rao bound meets spin dynamics," Preprint, 2017, arXiv:1710.08062.
- [23]. Marseille GJ, Fuderer M, de Beer R, Mehlkopf AF, and van Ormondt D, "Reduction of MRI scan time through nonuniform sampling and edge-distribution modeling," *J. Magn. Reson. B*, vol. 103, pp. 292–295, 1994. [PubMed: 8019779]
- [24]. Reeves SJ, "Selection of k-space samples in localized MR spectroscopy of arbitrary volumes of interest," *J. Magn. Reson. Imag*, vol. 5, pp. 245–247, 1995.
- [25]. Gao Y and Reeves SJ, "Optimal k-space sampling in MRSI for images with a limited region of support," *IEEE Trans. Med. Imag*, vol. 19, pp. 1168–1178, 2000.
- [26]. Gao Y and Reeves SJ, "Efficient backward selection of k-space samples in MRI on a hexagonal grid," *Circuits Systems Signal Process*, vol. 19, pp. 267–278, 2000.
- [27]. Gao Y and Reeves SJ, "Fast k-space sample selection in MRSI with a limited region of support," *IEEE Trans. Med. Imag*, vol. 20, pp. 868–876, 2001.
- [28]. Gao Y, Stakowski SM, Reeves SJ, Hetherington HP, Chu W-J, and Lee J-H, "Fast spectroscopic imaging using online optimal sparse k-space acquisition and projections onto convex sets reconstruction," *Magn. Reson. Med*, vol. 55, pp. 1265–1271, 2006. [PubMed: 16680731]
- [29]. Sharif B, Derbyshire JA, Faranesh AZ, and Bresler Y, "Patient-adaptive reconstruction and acquisition in dynamic imaging with sensitivity encoding (PARADISE)," *Magn. Reson. Med*, vol. 64, pp. 501–513, 2010. [PubMed: 20665794]
- [30]. Xu D, Jacob M, and Liang Z-P, "Optimal sampling of k-space with Cartesian grids for parallel MR imaging," in *Proc. Int. Soc. Magn. Reson. Med*, 2005, p. 2450.
- [31]. Pruessmann KP, Weiger M, Scheidegger MB, and Boesiger P, "SENSE: Sensitivity encoding for fast MRI," *Magn. Reson. Med*, vol. 42, pp. 952–962, 1999. [PubMed: 10542355]

- [32]. Samsonov AA, "Automatic design of radial trajectories for parallel MRI and anisotropic fields-of-view," in Proc. Int. Soc. Magn. Reson. Med, 2009, p. 765.
- [33]. Haldar JP, Hernando D, and Liang Z-P, "Super-resolution reconstruction of MR image sequences with contrast modeling," in Proc. IEEE Int. Symp. Biomed. Imag., 2009, pp. 266–269.
- [34]. Levine E and Hargreaves B, "On-the-fly adaptive k -space sampling for linear MRI reconstruction using moment-based spectral analysis," IEEE Trans. Med. Imag., vol. 37, pp. 557–567, 2018.
- [35]. Seeger M, Nickisch H, Pohmann R, and Schölkopf B, "Optimization of k -space trajectories for compressed sensing by Bayesian experimental design," Magn. Reson. Med, vol. 63, pp. 116–126, 2010. [PubMed: 19859957]
- [36]. Gorman JD and Hero AO, "Lower bounds for parametric estimation with constraints," IEEE Trans. Inf. Theory, vol. 26, pp. 1285–1301, 1990.
- [37]. Ben-Haim Z and Eldar YC, "The Cramér-Rao bound for estimating a sparse parameter vector," IEEE Trans. Signal Process, vol. 58, pp. 3384–3389, 2010.
- [38]. Fessler JA, "Model-based image reconstruction for MRI," IEEE Signal Process. Mag, vol. 27, pp. 81–89, 2010. [PubMed: 21135916]
- [39]. Kay SM, Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory. Upper Saddle River: Prentice Hall, 1993.
- [40]. Robson PM, Grant AK, Madhuranthakam AJ, Lattanzi R, Sodickson DK, and McKenzie CA, "Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k -space-based parallel imaging reconstructions," Magn. Reson. Med, vol. 60, pp. 895–907, 2008. [PubMed: 18816810]
- [41]. Fessler JA, "Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): Applications to tomography," IEEE Trans. Image Process, vol. 5, pp. 493–506, 1996. [PubMed: 18285134]
- [42]. Reeves SJ and Zhe Z, "Sequential algorithms for observation selection," IEEE Trans. Signal Process, vol. 47, pp. 123–132, 1999.
- [43]. Broughton R, Coope ID, Renaud P, and Tappenden R, "Determinant and exchange algorithms for observation subset selection," IEEE Trans. Image Process, vol. 19, pp. 2437–2443, 2010. [PubMed: 20388597]
- [44]. Kovacevic J and Chebira A, "Life beyond bases: The advent of frames (part I)," IEEE Signal Process. Mag, vol. 24, pp. 86–104, 2007.
- [45]. Starck J-L, Fadili J, and Murtagh F, "The undecimated wavelet decomposition and its reconstruction," IEEE Trans. Image Process, vol. 16, pp. 297–309, 2007. [PubMed: 17269625]
- [46]. Candès E, Demanet L, Donoho D, and Ying L, "Fast discrete curvelet transforms," Multiscale Model. Simul, vol. 5, pp. 861–899, 2006.
- [47]. Ravishanker S and Bresler Y, "MR image reconstruction from highly undersampled k -space data by dictionary learning," IEEE Trans. Med. Imag, vol. 30, pp. 1028–1041, 2011.
- [48]. Lingala SG and Jacob M, "Blind compressive sensing dynamic MRI," IEEE Trans. Med. Imag, vol. 32, pp. 1132–1145, 2013.
- [49]. Golub G and van Loan C, Matrix Computations, 3rd ed. London: The Johns Hopkins University Press, 1996.
- [50]. Pruessmann KP, Weiger M, Börnert P, and Boesiger P, "Advances in sensitivity encoding with arbitrary k -space trajectories," Magn. Reson. Med, vol. 46, pp. 638–651, 2001. [PubMed: 11590639]
- [51]. Allison MJ, Ramani S, and Fessler JA, "Accelerated regularized estimation of MR coil sensitivities using augmented Lagrangian methods," IEEE Trans. Med. Imag, vol. 32, pp. 556–564, 2013.
- [52]. Stoica P and Marzetta TL, "Parameter estimation problems with singular information matrices," IEEE Trans. Signal Process, vol. 49, pp. 87–90, 2001.
- [53]. Nayak KS and Nishimura DG, "Randomized trajectories for reduced aliasing artifact," in Proc. Int. Soc. Magn. Reson. Med, 1998, p. 670.
- [54]. Rudin LI, Osher S, and Fatemi E, "Nonlinear total variation based noise removal algorithms," Physica D, vol. 60, pp. 259–268, 1992.

- [55]. Nikolova M and Ng MK, "Analysis of half-quadratic minimization methods for signal and image recovery," SIAM J. Sci. Comput, vol. 27, pp. 937–966, 2005.
- [56]. Halдар JP, "Constrained imaging: Denoising and sparse sampling," Ph.D. dissertation, University of Illinois at Urbana-Champaign, Urbana, IL, USA, 2011 [Online]. Available: <http://hdl.handle.net/2142/24286>
- [57]. Buehrer M, Pruessmann KP, Boesiger P, and Kozerke S, "Array compression for MRI with large coil arrays," Magn. Reson. Med, vol. 57, pp. 1131–1139, 2007. [PubMed: 17534913]
- [58]. Huang F, Vijayakumar S, Li Y, Hertel S, and Duensing GR, "A software channel compression technique for faster reconstruction with many channels," Magn. Reson. Imag, vol. 26, pp. 133–141, 2008.
- [59]. Breuer FA, Blaimer M, Mueller MF, Seiberlich N, Heidemann RM, Griswold MA, and Jakob PM, "Controlled aliasing in volumetric parallel imaging (2D CAIPIRINHA)," Magn. Reson. Med, vol. 55, pp. 549–556, 2006. [PubMed: 16408271]
- [60]. Zhao B, Lam F, and Liang Z-P, "Model-based MR parameter mapping with sparsity constraints: Parameter estimation and performance bounds," IEEE Trans. Med. Imag, vol. 33, pp. 1832–1844, 2014.
- [61]. Kim D, Wisnowski JL, and Halдар JP, "Improved efficiency for microstructure imaging using high-dimensional MR correlation spectroscopic imaging," in Proc. Asilomar, 2017, pp. 1264–1268.
- [62]. Kim D, Doyle EK, Wisnowski JL, Kim JH, and Halдар JP, "Diffusion-relaxation correlation spectroscopic imaging: A multidimensional approach for probing microstructure," Magn. Reson. Med, vol. 78, pp. 2236–2249, 2017. [PubMed: 28317261]
- [63]. Kim D, Wisnowski J, Nguyen C, and Halдар JP, "Probing in vivo microstructure with T1-T2 relaxation correlation spectroscopic imaging," in Proc. IEEE Int. Symp. Biomed. Imag, 2018, pp. 675–678.
- [64]. Tang G and Nehorai A, "Lower bounds on the mean-squared error of low-rank matrix reconstruction," IEEE Trans. Signal Process, vol. 59, pp. 4559–4571, 2011.
- [65]. Zachariah D, Sundin M, Jansson M, and Chatterjee S, "Alternating least-squares for low-rank matrix reconstruction," IEEE Signal Process. Lett, vol. 19, pp. 231–234, 2012.
- [66]. Liu X and Sidiropoulos ND, "Cramér-Rao lower bounds for low-rank decomposition of multidimensional arrays," IEEE Trans. Signal Process, vol. 49, pp. 2074–2086, 2001.
- [67]. Halдар JP and Liang Z-P, "Spatiotemporal imaging with partially separable functions: A matrix recovery approach," in Proc. IEEE Int. Symp. Biomed. Imag, 2010, pp. 716–719.
- [68]. Lingala SG, Hu Y, DiBella E, and Jacob M, "Accelerated dynamic MRI exploiting sparsity and low-rank structure: k-t SLR," IEEE Trans. Med. Imag, vol. 30, pp. 1042–1054, 2011.
- [69]. Halдар JP and Liang Z-P, "Low-rank approximations for dynamic imaging," in Proc. IEEE Int. Symp. Biomed. Imag, 2011, pp. 1052–1055.
- [70]. Shin PJ, Larson PEZ, Ohliger MA, Elad M, Pauly JM, Vigneron DB, and Lustig M, "Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion," Magn. Reson. Med, vol. 72, pp. 959–970, 2014. [PubMed: 24248734]
- [71]. Halдар JP, "Low-rank modeling of local fc-space neighborhoods (LORAKS) for constrained MRI," IEEE Trans. Med. Imag, vol. 33, pp. 668–681, 2014.
- [72]. Halдар JP and Kim TH, "Computational imaging with LORAKS: Reconstructing linearly predictable signals using low-rank matrix regularization," in Proc. Asilomar, 2017, pp. 1870–1874.
- [73]. Trzasko JD and Manduca A, "A unified tensor regression framework for calibrationless dynamic, multi-channel MRI reconstruction," in Proc. Int. Soc. Magn. Reson. Med, 2013, p. 603.
- [74]. He J, Liu Q, Christodoulou AG, Ma C, Lam F, and Liang Z-P, "Accelerated high-dimensional MR imaging with sparse sampling using low-rank tensors," IEEE Trans. Med. Imag, 2016, In Press.
- [75]. Kim TH and Halдар JP, "The Fourier radial error spectrum plot: A more nuanced quantitative evaluation of image reconstruction quality," in Proc. IEEE Int. Symp. Biomed. Imag, 2018, pp. 61–64.

- [76]. Haldar JP and Kim D, "OEDIPUS: An experiment design framework for sparsity-constrained MRI," Preprint, 2018, arXiv:1805.00524.
- [77]. Gozcu B, Mahabadi RK, Li Y-H, Ilıcak E, Cukur T, Scarlett J, and Cevher V, "Learning-based compressive MRI," Preprint, 2018, arXiv:1805.01266.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Fig. 1.

Gold standard reference images for the 2D T2-weighted brain imaging experiment. The top row shows data from the first subject, while the bottom row shows data from the second subject.

Fig. 2.
Sampling patterns for the 2D T2-weighted brain imaging experiment.

Fig. 3.

Gold standard reference images for the 3D T1-weighted brain imaging experiment. The top row shows data from the healthy subject, while the bottom row shows data from the stroke subject. The slice we've chosen for the stroke subject includes both a stroke lesion and a fiducial marker, both of which are absent in the image for the healthy subject. The fiducial marker appears as a white circle on the top left side of the image, while the lesion appears as a dark spot in the white matter near the middle of the brain on the left side of the image.

Fig. 4.
Sampling patterns for the 3D T1-weighted brain imaging experiment.

TABLE I

TABLE OF NRMSE VALUES FOR RECONSTRUCTED SINGLE-CHANNEL 2D T2-WEIGHTED BRAIN DATA. RESULTS ARE SHOWN FOR WAVELET (WAV) AND TOTAL VARIATION (TV) RECONSTRUCTION APPROACHES, AND FOR UNIFORM (UNI), RANDOM (RAND), AND OEDIPUS SAMPLING PATTERNS.

		First Subject (Training)				Second Subject			
		Uni	Rand	SCO	MCO	Uni	Rand	SCO	MCO
$R = 2$	Wav	0.792	0.178	0.126	0.310	0.855	0.204	0.143	0.339
	TV	0.801	0.114	0.095	0.185	0.801	0.155	0.114	0.305

TABLE II

TABLE OF NRMSE VALUES FOR RECONSTRUCTED MULTI-CHANNEL 2D T2-WEIGHTED BRAIN DATA. RESULTS ARE SHOWN FOR WAVELET (WAV) AND TOTAL VARIATION (TV) RECONSTRUCTION APPROACHES, AND FOR UNIFORM (UNI), RANDOM (RAND), AND OEDIPUS SAMPLING PATTERNS.

		First Subject (Training)				Second Subject			
		Uni	Rand	SCO	MCO	Uni	Rand	SCO	MCO
$R = 2$	Wav	0.046	0.051	0.056	0.049	0.053	0.061	0.065	0.056
	TV	0.043	0.046	0.052	0.046	0.050	0.055	0.060	0.053
$R = 3$	Wav	0.087	0.090		0.072	0.105	0.110		0.084
	TV	0.071	0.076		0.064	0.085	0.092		0.075
$R = 4$	Wav	0.160	0.149		0.096	0.194	0.176		0.113
	TV	0.126	0.117		0.082	0.145	0.144		0.097

TABLE III

TABLE OF NRMSE VALUES FOR RECONSTRUCTED SINGLE-CHANNEL 3D T1-WEIGHTED BRAIN DATA. RESULTS ARE SHOWN FOR WAVELET (WAV) AND TOTAL VARIATION (TV) RECONSTRUCTION APPROACHES, AND FOR UNIFORM (UNI), RANDOM (RAND), AND OEDIPUS SAMPLING PATTERNS.

		Healthy Subject (Training)				Stroke Subject			
		Uni	Rand	SCO	MCO	Uni	Rand	SCO	MCO
$R = 2$	Wav	0.470	0.153	0.117	0.118	0.413	0.122	0.085	0.086
	TV	0.325	0.141	0.104	0.105	0.238	0.114	0.075	0.076
$R = 3$	Wav	0.601	0.249	0.199	0.166	0.616	0.217	0.136	0.135
	TV	0.512	0.203	0.145	0.141	0.427	0.178	0.110	0.112
$R = 4$	Wav	0.831	0.324	0.270	0.650	0.856	0.295	0.210	0.613
	TV	0.748	0.246	0.182	0.289	0.867	0.233	0.140	0.253
$R = 5$	Wav	0.815	0.389	0.324	0.735	0.913	0.354	0.273	0.699
	TV	0.732	0.288	0.213	0.376	0.869	0.272	0.177	0.371
$R = 6$	Wav	0.820	0.439	0.402	0.767	0.896	0.383	0.330	0.763
	TV	0.787	0.317	0.264	0.535	0.889	0.297	0.248	0.600

TABLE IV

TABLE OF NRMSE VALUES FOR RECONSTRUCTED MULTI-CHANNEL 3D T1-WEIGHTED BRAIN DATA. RESULTS ARE SHOWN FOR WAVELET (Wav) AND TOTAL VARIATION (TV) RECONSTRUCTION APPROACHES, AND FOR UNIFORM (Uni), RANDOM (Rand), AND OEDIPUS SAMPLING PATTERNS.

		Healthy Subject (Training)				Stroke Subject			
		Uni	Rand	SCO	MCO	Uni	Rand	SCO	MCO
$R = 2$	Wav	0.070	0.074	0.074	0.073	0.053	0.056	0.055	0.054
	TV	0.069	0.071	0.071	0.069	0.052	0.054	0.052	0.052
$R = 3$	Wav	0.092	0.095	0.103	0.096	0.070	0.072	0.074	0.071
	TV	0.088	0.089	0.092	0.089	0.067	0.068	0.069	0.067
$R = 4$	Wav	0.112	0.114	0.126	0.114	0.087	0.089	0.090	0.086
	TV	0.103	0.105	0.109	0.103	0.080	0.082	0.082	0.079
$R = 5$	Wav	0.122	0.129	0.142	0.129	0.104	0.107	0.105	0.098
	TV	0.112	0.117	0.121	0.113	0.089	0.095	0.093	0.089
$R = 6$	Wav	0.144	0.145	0.154	0.141	0.161	0.126	0.117	0.111
	TV	0.128	0.129	0.130	0.122	0.105	0.109	0.102	0.098
$R = 7$	Wav	0.181	0.161		0.153	0.316	0.151		0.122
	TV	0.143	0.140		0.130	0.146	0.128		0.106
$R = 8$	Wav	0.194	0.178		0.162	0.336	0.177		0.135
	TV	0.156	0.152		0.138	0.164	0.145		0.114