Available online at www.sciencedirect.com
SCIENCE DIRECT
Magnetic Resonance Imaging xx (2004) xxx–xxx
MAGNETIC RESONANCE IMAGING
ELSEVIER

Detection of pseudoperiodic patterns using partial acquisition of magnetic resonance images

Oren Boiman[a], Sharon Peled[b,1], Yehezkel Yeshurun[a,*]

[a]Sackler Faculty of Exact Sciences, School of Computer Science, Tel Aviv University, Ramat Aviv 69978, Israel
[b]Department of Radiology, Tel Aviv Sourasky Medical Center, Ramat Aviv 69978, Israel

Received 2 February 2004; accepted 11 August 2004

Abstract

Improving the resolution of magnetic resonance imaging (MRI), or, alternatively, reducing the acquisition time, can be quite beneficial for many applications. The main motivation of this work is the assumption that any information that is a priori available on the target image could be used to achieve this goal. In order to demonstrate this approach, we present a novel partial acquisition strategy and reconstruction algorithm, suitable for the special case of detection of pseudoperiodic patterns. Pseudoperiodic patterns are frequently encountered in the cerebral cortex due to its columnar functional organization (best exemplified by orientation columns and ocular dominance columns of the visual cortex). We present a new MRI research methodology, in which we seek an activity pattern, and a pattern-specific experiment is devised to detect it. Such specialized experiments extend the limits of conventional MRI experiments by substantially reducing the scan time. Using the fact that pseudoperiodic patterns are localized in the Fourier domain, we present an optimality criterion for partial acquisition of the MR signal and a strategy for obtaining the optimal discrete Fourier transform (DFT) coefficients. A by-product of this strategy is an optimal linear extrapolation estimate. We also present a nonlinear spectral extrapolation algorithm, based on projections onto convex sets (POCSs), used to perform the actual reconstruction. The proposed strategy was tested and analyzed on simulated signals and in MRI phantom experiments.

Keywords: Constrained reconstruction; Pattern detection; MRI; Spectral extrapolation

1. Introduction

Incorporating all prior knowledge in the process of estimation always leads to better results. This trivial truth could be used efficiently to improve the resolution (or, alternatively, the acquisition time) of magnetic resonance imaging (MRI) processes, when there exists an a priori model of the target. A prominent case where such a model does exist is manifested by the pseudoperiodic patterns that are typical of columnar functional organization in the cerebral cortex [1–4]. An example is the ocular dominance columnar structure, found in layer 4 of the striate cortex. Ocular dominance columns are attributed to be stereoscopic processing units, but there is no single computational model for their action [5]. These columns exhibit a pseudoperiodic pattern of alternating stripes (500–1000 μm wide) of left-eye dominated cells and right-eye dominated cells.

Our work presents a framework that enables inclusion of an underlying model in order to improve upon imaging results. Using the columnar model that results in a pseudoperiodic pattern, we have designed a specific model-dependent MRI experiment, which may be able to detect such patterns in a fraction of the time required by straightforward imaging. While the presented technique is general enough to fit to other kinds of patterns, it may provide exceptionally high scan-time reduction for pseudoperiodic patterns. Some MRI methods are inherently limited in spatial resolution due to timing constraints. A prominent example is the single-shot echo-planar imaging (EPI) technique, widely used for functional MRI (fMRI) studies of task-related brain activity. For such methods, the option of tailoring the sensitivity of the data acquisition sequence to specific patterns may be particularly useful.

* Corresponding author. Tel.: +972 3 640 9367; fax: +972 3 640 9358.
E-mail address: hezy@post.tau.ac.il (Y. Yeshurun).

[1] Current address: Harvard Center for Neurodegeneration and Repair, 220 Longwood Avenue, Boston, MA 02115, USA.

0730-725X/$ – see front matter © 2004 Published by Elsevier Inc.
doi:10.1016/j.mri.2004.08.016

In the Theory Section, we discuss pseudoperiodicity and describe an optimality criterion for partial acquisition of MR signals. This criterion yields an optimal linear extrapolation estimate of the data. We then introduce an algorithm based on projections onto convex sets (POCSs) used to reconstruct the data. In the Computer Simulations and MRI Phantom Experiments Sections, we present an implementation of the suggested method using both computer-generated images and MRI scans. We conclude by discussing the implications of our method.

2. Theory

Our goal is to perform high-resolution imaging of pseudoperiodic patterns. We assume that, as in single-shot EPI, the appropriate high-resolution data cannot be fully acquired due to signal decay over time. We approach the problem by acquiring only a part of the data, that is, an optimally sampled fraction of the discrete Fourier transform (DFT) coefficients. The percent of acquisition is determined by the scanner capabilities: sample as many DFT coefficients as possible before signal loss is too great. In order to choose the optimal DFT coefficients to sample, a model-based preprocessing is performed and the results are fed to the scanner. After partial acquisition, the raw data are transmitted to postprocessing used to reconstruct the full image from the partial data.

The entire acquisition and reconstruction strategy used in this work is depicted in Fig. 1. The preprocessing stage computes an acquisition strategy whose initial estimation is based on prior knowledge (sample model images). The initial estimation feeds into an iterative improvement procedure, using simulated annealing. A product of this process is an approximately optimal acquisition strategy for linear extrapolation and an optimal linear extrapolator. Using this acquisition strategy, an MRI partial acquisition experiment can be conducted. After partial acquisition, a postprocessing reconstruction takes place. The raw k-space data feeds a nonlinear iterative extrapolation algorithm, based on POCS (discussed in this section) and constrained by prior knowledge. The linear extrapolator ('by-product' of the preprocessing) can also be used to enhance the nonlinear extrapolation. The restored image is the output of the algorithm. In this section, we outline the qualities of pseudoperiodic signals and describe our new strategy for partial acquisition of such signals. Next, we describe a new POCS-based algorithm for spectrum extrapolation, suitable for partially acquired signals. We end the section by comparing the POCS extrapolation to Bayesian extrapolation techniques and describe a generalization of POCS, which enables the use of nonconvex projections.

2.1. Pseudoperiodic signals

In this work, we deal with a broad class of random signals, denoted pseudoperiodic signals. This class includes the set of almost-periodic random signals [6]. In practice, the only requirement of a random signal to be regarded as pseudoperiodic is that its covariance function oscillates, with a 'pseudoperiod', though not necessarily in the form of a sinusoid. Consequently, we expect to find several peaks in the spectrum of such processes. In other words, most of the energy of pseudoperiodic processes is concentrated in several peaks in the frequency domain. We can use this characteristic to approximate a signal by sampling a few Fourier coefficients (partial acquisition) and restore the other coefficients (spectrum extrapolation) using prior information about the signal. Pseudoperiodic patterns include harmonic processes and many more 'real-life' examples, which makes them a practical model.

2.2. Optimal acquisition of pseudoperiodic processes

The acquisition model we use is the linear model $\mathbf{y}=\mathbf{P}\mathbf{x}+\mathbf{n}$, where $\mathbf{x}$ is the original signal vector, $\mathbf{P}$ is the acquisition system function, $\mathbf{n}$ is an additive noise vector and $\mathbf{y}$ is the acquired image. $\mathbf{x}_e$ denotes the estimate of $\mathbf{x}$.$\mathbf{n}$ is assumed to be zero-mean white Gaussian noise with

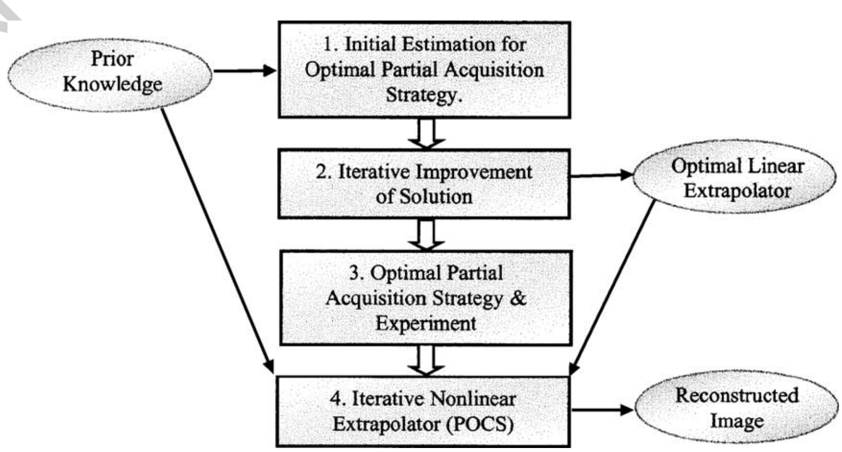

Fig. 1. The overall scheme: boxes represent algorithms and ellipses represent data. Stages 1 and 2 are the preprocessing steps. In Stage 3, we perform the actual acquisition. Stage 4 is a postprocessing stage used to reconstruct the image.

variance σ^2 so that $\mathbf{R}_{nn}=\sigma^2\mathbf{I}$, where $\mathbf{R}_{nn}$ is the autocorrelation matrix of $\mathbf{n}$. We will denote by $\mathbf{R}_{xx}$ the autocorrelation matrix of $\mathbf{x}$. We assume that $\mathbf{x}$ has been normalized to be a zero-mean vector, so that $\mathbf{R}_{xx}$ is also the covariance matrix of $\mathbf{x}$. $\mathbf{P}$ depends on the used acquisition strategy. In case of a full acquisition, $\mathbf{P}$ denotes the matrix appropriate for the exponential signal decay. In general, the error measure we use is the mean square error (MSE).

In our context of high-resolution restoration of pseudo-periodic signals, we would like to assure that low-error estimations correspond to high-resolution estimates. Practically, this means we might give up the optimality in the MSE sense for an increased resolution of the estimate. Suppose we are interested in identifying small, pseudoperiodic details of size σ_R. Let $\mathbf{G}$ be a real, linear, FIR high-pass filter with an appropriate cut-off frequency. The convolution of $(\mathbf{x}-\mathbf{x}_e)$ with $\mathbf{G}$ corresponds to the exclusion of coarse details in the difference signal. Let $\mathbf{W}$ denote the operator of convolution with $\mathbf{G}$. The error measure

$$E\left[\left(x-x_e\right)^{T*}\mathbf{W}^{T*}\mathbf{W}\left(x-x_e\right)\right] \tag{1}$$

is a weighted mean-square error that discounts the error of coarse details. A better approach is to balance between coarse and fine details. This trade-off is expressed in the filter vector:

$$G_\lambda=\delta_K+\lambda\mathbf{G},\qquad \lambda\geq 0 \tag{2}$$

$\mathbf{G}_\lambda$ is a high-enhance filter (e.g., unsharp filter). Let $\mathbf{W}_\lambda$ denote the operator of convolution with $\mathbf{G}_\lambda$. Thus, the weighted mean-square error used is

$$E\left[\left(x-x_e\right)^{T*}\mathbf{W}_\lambda^{T*}\mathbf{W}_\lambda\left(x-x_e\right)\right] \tag{3}$$

We will later denote $\mathbf{W}_\lambda$ simply by $\mathbf{W}$. We next describe the utilization of principal component analysis for the problem of determining the partial acquisition DFT coefficients. This technique is suboptimal and leads to several implementation problems. Afterwards, we describe a new technique for finding the approximately optimal DFT coefficients along with an optimal linear extrapolation filter.

2.2.1. Principal component analysis

Principal component analysis using the Karhunen–Loeve (KL) transform can be used to determine the DFT coefficients in a partial acquisition problem. It is well known that DFT is the KL transform for periodic random processes [6]. For circular wide-sense stationary or almost-periodic processes, the DFT and the KL transforms are very similar. This implies that it is possible to approximately express the KL basis vectors using a few DFT coefficients. However, finding the optimal DFT coefficients given the KL basis is not trivial. In Ref. [7], it is suggested to approximate the KL basis using a partial acquisition of the DFT where the optimal DFT coefficients are found by an exhaustive search. This method is not practical even for short signals (e.g., signals of length 64), and the authors recommend the use of simulated annealing for solving this problem. In Ref. [8], a branch and bound algorithm is used to improve the efficiency of the technique. However, it is not clear that the KL basis is optimal for the case of partial acquisition of the DFT. In Ref. [9], it is suggested to directly acquire the KL expansion instead of the DFT. This method should be optimal in theory, but it has some practical problems, and it requires the use of nonstandard MRI acquisition techniques. In the next subsection, we describe an efficient algorithm to find an approximation of the optimal subset of DFT coefficients to acquire.

2.2.2. An optimal solution

As before, let us denote by $\mathbf{x}$ and $\mathbf{y}$ the spectrum of the original signal and the acquired signal, respectively. Let b denote the number of acquired DFT coefficients. Let us describe the partial acquisition model by the expression

$$\mathbf{y}_{b\times 1}=\mathbf{PEx}_{N\times 1}+\mathbf{En}_{N\times 1} \tag{6}$$

where $\mathbf{E}$ is a binary ($b\times$N) acquisition matrix. That is, the vector $\mathbf{Ex}$ contains only b nonzero DFT coefficients. $\mathbf{P}$ is a constant ($b\times b$) diagonal degradation matrix modeling the decay of the signal during acquisition time. Given $\mathbf{P}$ and $\mathbf{E}$, we seek an optimal filter $\mathbf{G}$, which can be viewed as a linear extrapolator, such that the estimate $\mathbf{x}_e=\mathbf{Gy}$ would bring the weighted MSE to a minimum. The optimal operator $\mathbf{G}$ is the Wiener filter:

$$\mathbf{G}_{opt}=\mathbf{R}_{xx}\mathbf{E}^{T*}\mathbf{P}^{T*}\left(\mathbf{PER}_{xx}\mathbf{E}^{T*}\mathbf{P}^{T*}+\mathbf{ER}_{nn}\mathbf{E}^{T*}\right)^{-1} \tag{7}$$

and the resulting error is

$$\begin{aligned} err_{opt} &= Tr\left\{\mathbf{W}\left(\mathbf{R}_{xx}-\mathbf{GPER}_{xx}\right)\mathbf{W}^{T*}\right\} \\ &= Tr\left\{\mathbf{W}\left(\mathbf{R}_{xx}-\mathbf{R}_{xx}\mathbf{E}^{T*}\mathbf{P}^{T*}\left(\mathbf{PER}_{xx}\mathbf{E}^{T*}\mathbf{P}^{T*}\right.\right.\right. \\ &\qquad \left.\left.\left.+\mathbf{ER}_{nn}\mathbf{E}^{T*}\right)^{-1}\mathbf{PER}_{xx}\right)\mathbf{W}^{T*}\right\} \end{aligned} \tag{8}$$

The MSE is minimal for every given $\mathbf{E}$. Consequently, the optimal acquisition matrix $\mathbf{E}$ is given by minimizing the MSE with respect to $\mathbf{E}$. A similar optimality criterion is used in Ref. [10]. Minimizing the expression cannot be achieved by derivation because the space of matrices $\mathbf{E}$ is discrete.

Exhaustive search is not practical either, from the same reasons discussed previously. Minimizing the MSE is equivalent to the maximization of f where

$$f = Tr\left\{\mathbf{W}\left(\mathbf{R}_{xx}\mathbf{E}^{T*}\mathbf{P}^{T*}\left(\mathbf{PER}_{xx}\mathbf{E}^{T*}\mathbf{P}^{T*} + \mathbf{ER}_{nn}\mathbf{E}^{T*}\right)^{-1}\mathbf{PER}_{xx}\right)\mathbf{W}^{T*}\right\} = Tr\left\{\left[\mathbf{E}^{T*}\mathbf{P}^{T*}\left(\mathbf{PER}_{xx}\mathbf{E}^{T*}\mathbf{P}^{T*} + \mathbf{ER}_{nn}\mathbf{E}^{T*}\right)^{-1}\mathbf{PE}\right]\mathbf{R}_{xx}\mathbf{W}^{T*}\mathbf{WR}_{xx}\right\} \quad (9)$$

For many common processes (almost stationary processes in particular), most of the energy of $\mathbf{R}_{xx}$ is concentrated on the main diagonal. In this case, the matrix $\mathbf{PER}_{xx}\mathbf{E}^{T*}\mathbf{P}^{T*}$ is a matrix with a dominant main diagonal. If the additive noise is white, $\mathbf{R}_{nn}=\sigma^2\mathbf{I}$ and $\mathbf{ER}_{nn}\mathbf{E}^{T*}=\sigma^2\mathbf{I}_b$, where $\mathbf{I}_b$ is the identity matrix of size $(b\times b)$. We can make the approximation

$$\mathbf{E}^{T*}\mathbf{P}^{T*}\left(\mathbf{PER}_{xx}\mathbf{E}^{T*}\mathbf{P}^{T*} + \mathbf{ER}_{nn}\mathbf{E}^{T*}\right)^{-1}\mathbf{PE} \approx \mathbf{E}^{T*}\mathbf{P}^{T*}Diag \times \left(\mathbf{PER}_{xx}\mathbf{E}^{T*}\mathbf{P}^{T*} + \mathbf{ER}_{nn}\mathbf{E}^{T*}\right)^{-1}\mathbf{PE} \quad (10)$$

The resulting approximation is a diagonal matrix. This allows us to write

$$f \approx Tr\left\{\left[\mathbf{E}^{T*}\mathbf{P}^{T*}Diag\left(\mathbf{PER}_{xx}\mathbf{E}^{T*}\mathbf{P}^{T*} + \mathbf{ER}_{nn}\mathbf{E}^{T*}\right)^{-1}\mathbf{PE}\right]\mathbf{R}_{xx}\mathbf{W}^{T*}\mathbf{WR}_{xx}\right\} = Tr\left\{\left[\mathbf{E}^{T*}\mathbf{P}^{T*}Diag\left(\mathbf{PER}_{xx}\mathbf{E}^{T*}\mathbf{P}^{T*} + \sigma_n^2\mathbf{I}_b\right)^{-1}\mathbf{PE}\right]Diag\left(\mathbf{R}_{xx}\mathbf{W}^{T*}\mathbf{WR}_{xx}\right)\right\} = Tr\left\{\left[Diag\left(\mathbf{PER}_{xx}\mathbf{E}^{T*}\mathbf{P}^{T*} + \sigma_n^2\mathbf{I}_b\right)^{-1}\right] \times \left[\mathbf{PE}Diag\left(\mathbf{R}_{xx}\mathbf{W}^{T*}\mathbf{WR}_{xx}\right)\mathbf{E}^{T*}\mathbf{P}^{T*}\right]\right\} \quad (11)$$

Assume for the moment that $\mathbf{P}=\mathbf{I}_b$. In this case, the expression we seek to maximize can be thought of as a partial sum of diagonal elements. $\mathbf{E}$ determines which elements are accumulated in the partial sum. Denote

$$a_i = \left\{Diag\left(\mathbf{R}_{xx} + \sigma_n^2\mathbf{I}\right)\right\}_{i,i}$$

$$b_i = \left\{Diag\left(\mathbf{R}_{xx}\hat{\mathbf{W}}^{T*}\hat{\mathbf{W}}\mathbf{R}_{xx}\right)\right\}_{i,i} \quad (12)$$

Then, Eq. (11) is maximized by choosing $\mathbf{E}$ so that the b maximal elements of the ratio b_i/a_i will be acquired. When $\mathbf{P}=\mathbf{I}_b$, there is no significance to the row order of $\mathbf{E}$. Otherwise, $\mathbf{E}$ should be found by other methods described in Ref. [11]. By choosing $\mathbf{E}$, this method also provides an approximation to the optimal linear extrapolation filter $\mathbf{G}_{opt}$. If the assumption of the approximation in Eq. (10) is not valid, other methods should be used in order to maximize the expression in Eq. (9). In this work, the maximization was implemented using a simulated annealing process, whose iterations were initialized by the described approximation. The optimal extrapolation filter is linear and it relies heavily on the accuracy of $\mathbf{R}_{xx}$. Note that in practice, we cannot guarantee the precision of $\mathbf{R}_{xx}$ because it is generated from a model. Moreover, our simulations show that even with an exact covariance matrix, linear extrapolation offers a low-quality restoration in our setting, which makes it unsuitable for our strict requirements. Nevertheless, it is an optimal linear acquisition and extrapolation and it outperforms any other *linear* extrapolation technique. Next, we describe a nonlinear spectral extrapolation algorithm, which does not make direct use of $\mathbf{R}_{xx}$. Thus, our scheme provides a clear separation between acquisition and extrapolation, which is useful where no high-resolution reference can be acquired (as assumed in Ref. [12]) and only an approximate model is supplied.

2.3. Spectrum extrapolation

In this subsection, we give an overview of the POCS algorithm and its use in restoration problems. Afterwards, we describe a new POCS-based algorithm for spectrum extrapolation of partially acquired signals. This algorithm utilizes prior knowledge, specific for our application.

2.3.1. Projection onto convex sets

Projections onto convex set are an iterative algorithm for finding elements that lie at the intersection of closed convex sets. That is, if $C_1,\ldots, C_n$ are closed convex sets and we wish to find *any* x such that:

$$x\in C_1 \cap \ldots \cap C_n \quad (13)$$

Projections onto convex set provide an iterative method for finding such an element x, assuming that the intersection is not empty. The algorithm relies on the knowledge of the projections P_i onto the convex sets C_i. It is shown [13] that the cyclic control sequence

$$f_{k+1} = P_n P_{n-1}\cdots P_1 f_k, \qquad k = 0,1,\ldots \quad (14)$$

converges to an element in the intersection of $C_1,\ldots, C_n$. In our context, y is a distorted signal and we seek the original signal x. The prior knowledge consists of n properties, each restrict x to lie in a convex set:

$$\begin{aligned} &x\in C_1 \\ &\vdots \\ &x\in C_n \end{aligned} \quad (15)$$

Projections onto convex set are used to seek a *feasible* solution, that is, a solution that is consistent with all the prior knowledge on x. Note that the projection operators are generally nonlinear. Thus, POCS provides a straightforward method to incorporate nonlinear prior knowledge in the restoration process. Ignoring issues of run-time, adding constraints that decrease the size of the intersection set improves the quality of the solution. In the next subsection, we present the prior knowledge we use in this work and the appropriate projections.

2.3.2. Applied projections

We now present the actual projections that were used in this work. We denote by $\mathbf{x}$ the original signal and by $\mathbf{y}$ the acquired signal. We denote by $\mathbf{x}_e$ the 'current' (in terms of iterative algorithm) estimation of $\mathbf{x}$. In order to simplify the presentation, we treat $\mathbf{x}$ as a 1D vector, though an extension to 2D is straightforward. We use the Fourier domain and the image domain interchangeably, where the domain is understood from the context. Proofs of correctness of the projections and the convexity of sets, not detailed here, can be found in Refs. [13–16].

2.3.2.1. Data constraint. The data constraint restricts $\mathbf{x}_e$ to agree with the acquired k-space data. Let $\mathbf{d}$ be a binary vector such that $d[i]=1$ if the ith element in k-space was acquired. Thus, a projection P_{data} is defined by

$$P_{data}(\mathbf{x}_e)_i = \mathbf{x}_e[i](1-d[i]) + \mathbf{y}[i]d[i] \tag{16}$$

2.3.2.2. Bounded support constraint. The bounded support constraint restricts $\mathbf{x}_e$ to a bounded support, dependent on the signal. For instance, the size and the location of the scanned object in an MRI experiment are usually known from preliminary scans. Let $\mathbf{b}$ be a binary vector such that $b[i]=1$ if the ith element in $\mathbf{x}_e$ is a part of the signal's support. A projection $P_{support}$ is defined by

$$P_{support}(\mathbf{x}_e)_i = \mathbf{x}_e[i]\mathbf{b}[i] \tag{17}$$

2.3.2.3. Real signal constraint. The real signal constraint restricts $\mathbf{x}_e$ to be a real vector. This constraint is equivalent to the k-space's conjugate symmetry constraint. A projection P_{real} is given by

$$P_{real}(\mathbf{x}_e)_i = real(\mathbf{x}_e[i]) \tag{18}$$

2.3.2.4. Nonnegative signal constraint. The nonnegative signal constraint restricts $\mathbf{x}_e$ to a vector with nonnegative elements. It is usually the case in image restoration that pixel values cannot be negative, particularly in MRI. A projection $P_{nonnegative}$ is given by

$$P_{nonnegative}(\mathbf{x}_e)_i = \begin{cases} \mathbf{x}_e[i] & \mathbf{x}_e[i] \geq 0 \\ 0 & otherwise \end{cases} \tag{19}$$

2.3.2.5. Reference signal constraint. The reference signal constraint restricts $\mathbf{x}_e$ to lie in a sphere centered at a vector $\mathbf{r}$, known to be 'close' to $\mathbf{x}$. It is well known in MRI that a reference scan can be used to improve the extrapolation of a partially acquired image (e.g., keyhole imaging). The set

$$C_{reference} = \{\mathbf{x}_e \mid \|\mathbf{x}_e - \mathbf{r}\| \leq \varepsilon\} \tag{20}$$

specifies a sphere and is therefore convex. A projection $P_{reference}$ on the sphere $C_{reference}$ is given by

$$P_{reference}(\mathbf{x}_e) = \begin{cases} \mathbf{x}_e & \|\mathbf{x}_e - \mathbf{r}\| \leq \varepsilon \\ \mathbf{r} + \frac{\mathbf{x}_e - \mathbf{r}}{\|\mathbf{x}_e - \mathbf{r}\|}\varepsilon & otherwise \end{cases} \tag{21}$$

2.3.2.6. Bounded energy constraint. The bounded energy constraint restricts the energy of $\mathbf{x}_e$. The bound B can be estimated using preliminary scans. The set

$$C_{energy} = \{\mathbf{x}_e \mid \|\mathbf{x}_e\| \leq B\} \tag{22}$$

defines the constraint. Thus, the bounded energy constraint is a specific case of the reference signal constraint. A projection P_{energy} is given by

$$P_{energy}(\mathbf{x}_e) = \begin{cases} \mathbf{x}_e & \|\mathbf{x}_e\| \leq B \\ \frac{\mathbf{x}_e}{\|\mathbf{x}_e\|}B & otherwise \end{cases} \tag{23}$$

2.3.2.7. Bounded MSE constraint. The bounded MSE constraint restricts $\mathbf{x}_e$ to the set C_{MSE}

$$C_{MSE} = \{\mathbf{x}_e \mid \|\mathbf{P}\mathbf{x}_e - \mathbf{y}\| \leq E\} \tag{24}$$

where $\mathbf{P}$ is the degradation matrix, and E is an MSE error bound.

The projection $\mathbf{w}=P_{MSE}(\mathbf{x}_e)$ can be found by minimizing $\|\mathbf{w}-\mathbf{x}_e\|$ given $\|\mathbf{P}\mathbf{w}-\mathbf{y}\| \leq E$.

The result of the minimization is:

$$\mathbf{w} = \left(\mathbf{I} + \lambda \mathbf{P}^{T*}\mathbf{P}\right)^{-1}\left(\mathbf{x}_e + \lambda \mathbf{P}^{T*}\mathbf{y}\right) \tag{25}$$

where λ is a nonnegative number chosen to satisfy $\|\mathbf{P}\mathbf{w}-\mathbf{y}\|=E$. If $\mathbf{P}$ is shift invariant, the projection $\mathbf{w}$ is easier to obtain in the Fourier domain:

$$\mathbf{w}[i] = \frac{\mathbf{x}_e[i] + \lambda \overline{\mathbf{P}}[i,i]\mathbf{y}[i]}{1 + \lambda|\mathbf{P}[i,i]|^2} \tag{26}$$

and we set λ such that

$$\begin{aligned} \|\mathbf{P}\mathbf{w}-\mathbf{y}\|^2 &= \sum_i \left(\frac{\mathbf{P}[i,i]\mathbf{x}_e[i] + \lambda|\mathbf{P}[i,i]|^2\mathbf{y}[i]}{1+\lambda|\mathbf{P}[i,i]|^2} - \mathbf{y}[i]\right)^2 \\ &= \sum_i \frac{\left(\mathbf{P}[i,i]\mathbf{x}_e[i] - \mathbf{y}[i]\right)^2}{\left(1+\lambda|\mathbf{P}[i,i]|^2\right)^2} \\ &= \sum_i \frac{\mathbf{d}[i]}{\left(1+\lambda|\mathbf{P}[i,i]|^2\right)^2} \end{aligned} \tag{27}$$

Recalling that λ is nonnegative and noting that the expression in Eq. (27) is a continuous and monotone-decreasing function of λ, we can find a unique solution to Eq. (27) using the Newton method. The bounded MSE constraint is a generalization of the data constraint. It is a more accurate constraint in case the noise is not negligible or the degradation operator cannot be ignored.

2.3.2.8. Smoothness constraint. The smoothness constraint restricts $\mathbf{x}_e$ to the set C_{smooth}

$$C_{smooth} = \{\mathbf{x}_e | \, \| \mathbf{S}\mathbf{x}_e \| \leq E\} \tag{28}$$

where $\mathbf{S}$ is taken as a spatially invariant finite difference operator. This is a specific case of the bounded MSE constraint ($\mathbf{y}=\mathbf{0}$), and the projection is obtained in the same manner.

2.4. Bayesian techniques and generalized projections

Bayesian techniques are widely used for solving restoration problems. In this subsection, we compare the POCS algorithm to an application of Bayesian techniques for the problem of spectral extrapolation used in Ref. [17]. Maximizing the a posteriori probability (MAP) of the estimate $\mathbf{x}_e$ given the acquired signal $\mathbf{y}$ is obtained by maximizing the expression $P(\mathbf{y}|\mathbf{x}_e)P(\mathbf{x}_e)$.

In the linear degradation model $\mathbf{y}=\mathbf{P}\mathbf{x}+\mathbf{n}$, the likelihood is given by the probability density of the noise

$$P(\mathbf{y}|\mathbf{x}_e) = P(\mathbf{n}) = P(\mathbf{y}-\mathbf{P}\mathbf{x}_e) \tag{29}$$

The prior probability density $P(\mathbf{x}_e)$ expresses the statistical model of the signal. In Ref. [17], a MAP algorithm called BAISE is used for MR image reconstruction from a partially acquired k-space. The BAISE algorithm relies on the following prior knowledge:

1. $\mathbf{n}$ is an additive noise, whose real and imaginary parts are white Gaussian processes.
2. The scanned object O has a known bounded support. Any signal outside the support of the object is due to noise.
3. The object is real. Any imaginary component in the image is due to noise.
4. The partial derivatives of the object (approximated by a finite difference) have a Lorentzian distribution.

The Lorentzian distribution characterizes partial derivatives (edges) of general real-world images. It captures general image characteristics like flat regions and sharp edges. Let $\mathbf{x}_R$ and $\mathbf{x}_I$ be the real and imaginary parts of $\mathbf{x}_e$, respectively. Thus, the resulting expression for $P(\mathbf{x}_e|\mathbf{y})$ [ignoring the denominator $P(\mathbf{y})$] is

$$\begin{aligned} P(\mathbf{x}_e|\mathbf{y}) = {} & c_1 \exp\left[-\frac{1}{2\sigma^2}\left(\mathbf{y}-\mathbf{P}\mathbf{x}_e\right)^{T*}\left(\mathbf{y}-\mathbf{P}\mathbf{x}_e\right)\right] \\ & \times c_2 \exp\left[-\frac{1}{2\sigma^2}\sum_{i\notin 0}\mathbf{x}_R[i]^2\right] \times \prod_{i\in 0}\frac{1}{\pi\left(a+\delta\mathbf{x}_R[i]^2/a\right)} \\ & \times c_3 \exp\left[-\frac{1}{2\sigma^2}\sum_{\forall i}\mathbf{x}_I[i]^2\right] \end{aligned} \tag{30}$$

where $c1$, $c2$ and $c3$ are normalization factors and a is the Lorentzian distribution parameter. By maximizing $\log(P(\mathbf{x}_e|\mathbf{y}))$ and discarding additive constants we get

$$\begin{aligned} & \left[-\frac{1}{2\sigma^2}\left(\mathbf{y}-\mathbf{P}\mathbf{x}_e\right)^{T*}\left(\mathbf{y}-\mathbf{P}\mathbf{x}_e\right)\right] + \left[-\frac{1}{2\sigma^2}\sum_{i\notin 0}\mathbf{x}_R[i]^2\right] \\ & -\sum_{i\in O}\log\left(a+\delta\mathbf{x}_R[i]^2/a\right) + \left[-\frac{1}{2\sigma^2}\sum_{\forall i}\mathbf{x}_I[i]^2\right] \end{aligned} \tag{31}$$

Maximizing Eq. (31) is equivalent to the minimization of

$$\begin{aligned} & \left(\mathbf{y}-\mathbf{P}\mathbf{x}_e\right)^{T*}\left(\mathbf{y}-\mathbf{P}\mathbf{x}_e\right) + \sum_{i\notin 0}\mathbf{x}_R[i]^2 + \sum_{\forall i}\mathbf{x}_I[i]^2 \\ & + 2\sigma^2\sum_{i\in O}\log\left(a+\delta\mathbf{x}_R[i]^2/a\right) \end{aligned} \tag{32}$$

The first three terms are the following norms (squared):

1. MSE.
2. Energy of the noise outside the object's support.
3. Energy of the noise in the imaginary components.

Ignoring the last term, this MAP algorithm is equivalent to the minimization of the sum of these squared norms. Minimizing a function $f(\mathbf{x}_e)$ can be performed by POCS if $\boldsymbol{f}$ is a convex function of $\mathbf{x}_e$. If we can estimate that the optimal solution $\mathbf{x}_{opt}$ satisfies $f(\mathbf{x}_{opt})\leq l$ (and $\boldsymbol{l}$ is 'tight'), then the set $C=\{\mathbf{x}_e | \, \|f(\mathbf{x}_e)\| \leq l\}$ is convex and we can project on it. This is indeed the case for the first three terms in Eq. (32). Unfortunately, the last term in Eq. (32), representing the Lorentzian distribution, is not a convex function, so it cannot incorporate directly to the POCS approach. Utilizing the Lorentzian constraint can be done using either an optimization approach or a generalized projection approach.

2.4.1. Optimization approach

In this approach, a gradient search is used such that the search direction is projected on the convex constraint. That is

$$\mathbf{x}_e^{(n+1)} = P_1 \cdot \cdot P_k\left(\mathbf{x}_e^{(n)} - \lambda\nabla F\left(\mathbf{x}_e^{(n)}\right)\right) \tag{33}$$

where λ is the gradient step size. This form of projected gradient can be very slow to converge [18].

2.4.2. Generalized projections approach

In this approach, the POCS method is generalized to incorporate a projection onto nonconvex sets — in this case, the Lorentzian distribution of partial derivatives. Using the cyclic control, it is shown in Ref. [19] that if up to two nonconvex projections are used, the sum of distances of the estimate from the constraint sets is known to converge. Projecting $\mathbf{x}_e$ on the set

$$C_L = \{\mathbf{x} | \sum_{i \in O} \log\left(a + \delta \mathbf{x}_R[i]^2 / a\right) \leq \varepsilon_L\} \tag{34}$$

is performed by minimizing the Lagrangian

$$L(\mathbf{w},\lambda) = -\| \mathbf{w} - \mathbf{x}_e \|^2 + \lambda\left(\sum_{i \in O} \log\left(a + \delta \mathbf{w}[i]^2 / a\right) - \varepsilon_L\right) \tag{35}$$

The resulting expression for $\mathbf{w}$ is

$$\mathbf{w}[i] = \begin{cases} \mathbf{x}_R[i] + \lambda\left(\frac{(\mathbf{w}[i]-\mathbf{w}[i-1])}{a^2+(\mathbf{w}[i]-\mathbf{w}[i-1])^2} + \frac{(\mathbf{w}[i]-\mathbf{w}[i+1])}{a^2+(\mathbf{w}[i]-\mathbf{w}[i+1])^2}\right) & i \in O \\ \mathbf{x}_R[i] & \text{otherwise} \end{cases} \tag{36}$$

λ should be chosen to satisfy the equality in Eq. (34). The generalized POCS approach is the method used in the simulations and phantom experiments described in this article.

3. Computer simulations

3.1. Overview

In this section, we discuss the methods and results of computer simulations used to test the theory presented in the Theory Section. Two-dimensional data were generated and MR partial acquisition was simulated on the generated data. The generalized POCS algorithm presented in the previous section was applied to restore the data where the selection of the DFT coefficients to acquire was done according to the optimality criterion in Eq. (8). The optimal DFT coefficients were first chosen according to the approximation in Eq. (12), and then an optional simulated annealing process was used to improve the approximation. The chosen DFT coefficients determined the linear extrapolation filter in Eq. (7). The POCS iterations were initialized either by the partially acquired data or by the linear extrapolation estimate. We end this section by discussing the results of a control simulation used for sensitivity analysis.

3.2. Simulation data

The 2D data used for the simulation were the single image depicted in Fig. 2. The image is a sum of a base image of a human brain and a template containing several pseudoperiodic patterns. The pseudoperiodic patterns were generated using a multinormal stationary distribution with a decaying sinusoidal covariance (i.e., the first row of the Toeplitz covariance matrix is defined by a sinusoidal multiplied by a Hanning window). The brain image and the template were normalized to a zero mean and unit standard deviation and then they were added. Overall, the 2D data contain coarse details and very fine details with a spatial frequency close to the Nyquist frequency.

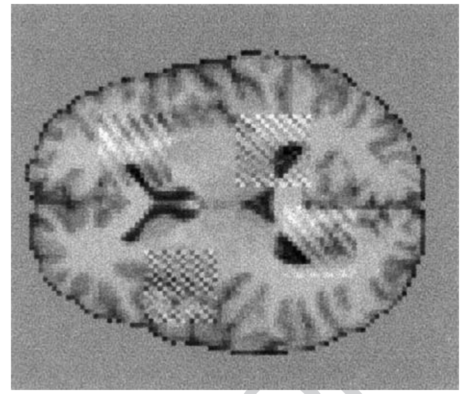

Fig. 2. Two-dimensional data used in the computer simulations.

3.3. Partial acquisition simulation

White Gaussian noise was added to the real and imaginary parts of the 2D DFT of the image so that the total SNR was 5. The SNR was computed relative to the standard deviation of the human brain image. Partial acquisition was simulated by using only a subset of the columns of the 2D DFT. This is not the optimal partial acquisition strategy, but it is a practical method of partial acquisition that suits the data acquisition strategy of EPI, the prevalent rapid acquisition technique used in fMRI. The subsets used comprised, respectively, 10%, 20%, 30% and 40% of the 2D DFT columns. The first and second statistical moments of the data, needed for the linear extrapolation, were computed using the rows of the data image as a sample set. The linear extrapolation was used to determine which columns to acquire in the 2D-DFT of the image. The error measure used for the linear extrapolation was a weighted MSE of the form in Eq. (2).

3.4. Signal restoration

The image was restored using the generalized iterative POCS algorithm, described previously. The iterative algorithm was initialized by the zero-filled DFT of the partially acquired image or the linear extrapolation estimate. The following constraints were used in the algorithm:

1. Reference image constraint: The reference image was set as the subsampled (ratio 1:2) original image.
2. Bounded MSE constraint: The image was restricted to agree with the acquired data up to a norm determined by the level of noise.
3. Smoothness constraint: The energy of the image's partial derivatives (approximated by a finite difference) was restricted to lie below a threshold set by corresponding energy of the original image.

4. Bounded energy constraint: The image's energy was restricted not to be higher than the energy of the original image.
5. Bounded support constraint: The support of the image was restricted to the support of the human brain depicted in Fig. 2.
6. Real image constraint.
7. Lorentzian distribution of the partial derivatives: The nonconvex constraint was set by approximating the Lorentzian distribution of the partial derivatives of the original image. The ε parameter was set according to the Lorentzian likelihood of the original image.
8. Reference image constraint: The second reference image used was the linear extrapolation estimate. Constraint no. 8 was applied only when the linear extrapolation estimate was used to initialize the POCS iterations. The ε parameter was chosen according to the true norm-distance between the original image to the linear extrapolation estimate. Herein, it is assumed that these bounds are found using a thorough parameter tuning. Some useful heuristics for tuning these parameters are described in Discussion. Moreover, a control experiment (presented in Control Simulation Section) was conducted in order to test the sensitivity of these parameters.

3.5. Results

Fig. 3 displays several samples of images reconstructed using the POCS algorithm (initialized by the acquired data), given different acquisition percents. Even using as little as 10% of the acquired image, the reconstruction enables the detection of the pseudoperiodic pattern. Fig. 4a shows the relative reconstruction error as a function of the acquisition percent. The relative MSE is the ratio of the MSE to the norm of the signal:

$$\text{Rel MSE} = \frac{\| I_r - I \|}{\| I \|},$$

where I denotes the original image, and I_r denotes the restored image.

Another error measure is used to quantify the quality of fine-detail restoration as follows

$$\text{MSE ratio} = \frac{\| I_r - I \|}{\| I_r - I_b \|}, \tag{37}$$

I_b denotes the brain image without the added template. The MSE ratio is zero for an exact restoration and might get to infinity if the restored image is identical to the brain image without the fine details. The MSE ratio as a function of the acquisition percent is shown in Fig. 4b. According to these results, an increase in the acquisition percent has a steady effect (almost linear) on the relative MSE and the MSE ratio. Projections onto convex set initialized by the acquired data perform better than the linear extrapolator and worse than POCS initialized by the linear extrapolation results.

3.6. Control simulation

A control experiment was conducted in order to test the sensitivity of the quality of reconstruction as a function of

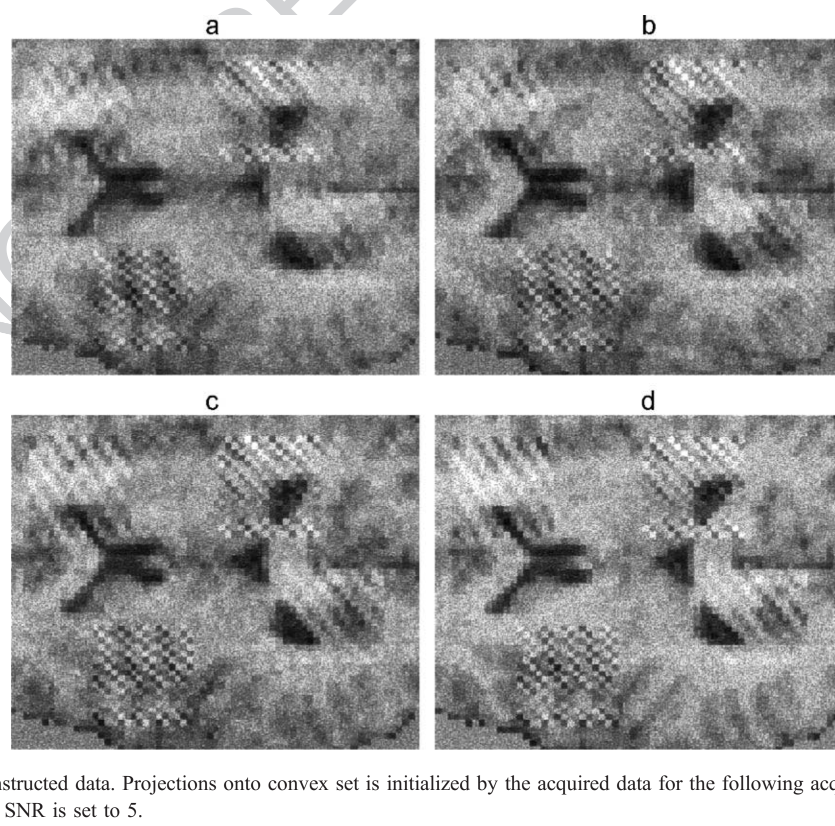

Fig. 3. Samples of 2D reconstructed data. Projections onto convex set is initialized by the acquired data for the following acquisition percents: (a) 10%, (b) 20%, (c) 30% and (d) 40%. SNR is set to 5.

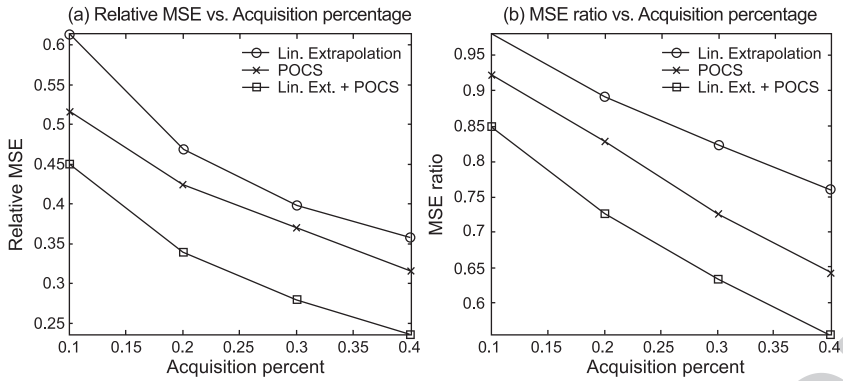

Fig. 4. Results of the 2D computer simulations. The figures show the reconstruction error as a function of the acquisition percent. The compared algorithms are the linear extrapolator, the POCS algorithm initialized by the acquired data (POCS) and the POCS algorithm initialized by the linear extrapolation estimate (Lin. Ext.+POCS). SNR is set to 5. (a) The relative MSE of the reconstruction. (b) The MSE ratio of the reconstruction.

the quality of the prior knowledge. The POCS algorithm is dependent on several parameters that define the constraints. The tightness of each constraint is represented by a threshold parameter (epsilon). We should choose this threshold according to the certainty of our estimations. The control simulation applied a single quality factor to all the thresholds. A quality factor of one does not change the thresholds. A quality factor >1 relaxes the constraints and a quality factor <1 tightens them. The other settings for the control simulation are as described above, where the SNR was set to 5% and the acquisition percent was set to 20%. The results of the simulation are presented in Fig. 5. The results show that setting the thresholds with an error of ~25% has little effect on the quality of reconstruction. Moreover, it is clear that whenever there is uncertainty in setting the thresholds, the upper bound should be selected. The reason is that reducing the thresholds might create a situation where the intersection of the convex sets is empty, or that it does not contain the original image. Increasing the thresholds increases the size of the intersection set, but it has a moderate effect on the reconstruction quality.

4. MRI phantom experiments

4.1. Overview

In this section, we present the results of applying the restoration algorithm presented in previous sections to raw data of MRI scans. We also discuss the general methods used in the acquisition and in the preprocessing of the raw data.

4.2. General methods

All MRI scans were conducted on a 1.5-T GE Signa Horizon LX MM scanner using a standard head coil. Two plastic gratings were constructed: each grating contained 10 compartments separated by thin plastic separators. In the first grating, each compartment was 2 mm wide. In the second grating, each compartment was 1 mm wide. Both gratings were placed in a water-filled fish bowl with a volume of 3000 ml. In order to lower the relaxation times of the water, 10 ml of Gd-DPTA was added to the water (concentration of 0.3%). The resulting T_2* was measured to be ~320 ms. In order to generate k-space data from which to choose lines, two spin-echo EPI scans were conducted with

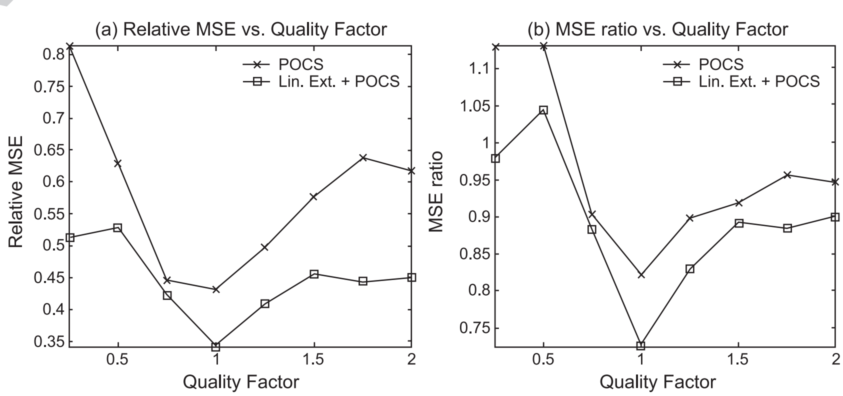

Fig. 5. Results of the control simulation. The compared algorithms are the POCS algorithm initialized by the acquired data (POCS) and the POCS algorithm initialized by the linear extrapolation estimate (Lin. Ext.+POCS). SNR is 5. Acquisition percent is 20%. (a) Relative MSE as a function of the quality factor. (b) MSE ratio as a function of the quality factor.

parameters: 128×512 (readout×phase); FOV=240 mm; slice thickness, 3 mm; bandwidth, 62.5 kHz; TR, 2 s. The gratings were set perpendicular to the phase-encoding direction. Two scans were conducted: a full scan and a half scan. One scan covered the entire extent of k-space, starting in the higher frequencies. The other scan covered half of k-space (plus several more phase-encode lines near the middle of k-space), starting in the lower frequencies. Effective TE for half k-space scans was 27.5 ms and for full scans was 645.3 ms. In addition, lower resolution EPI scans (128×256) were also acquired and were used as a reference in the restoration algorithm. Data from both scans were combined to create a single k-space as described in the next section.

4.3. Preprocessing scanned data

Raw k-space data were acquired and preprocessed before applying the restoration algorithm. For each type of scan (full scan or half scan), 50 repetitions of the scans were acquired and recorded. These were used for supplying images of increased SNR (by averaging). The whole restoration algorithm was applied four times, each time using a different number of repetitions (n=1,9,25,49), and thus with a different SNR. Given an averaged k-space in the two scanning modes, the two images were registered to the reference (T_1 weighted) image by a simple correlation. Afterwards, the two images were normalized in order to fit the norm of the reference image. Then, the two images were merged into a single k-space by taking from each image the part of k-space scanned first (i.e., without significant signal loss due to T_2 and T_2* signal decay). Finally, a phase correction algorithm using phase estimation based on the middle k_y lines in k-space [20] was applied to result in a single image. The same steps were also taken for the low-resolution images. SNR was computed as the ratio of the standard deviation of the noise (as estimated in the merged image) to the phantom standard deviation (i.e., the standard deviation of the phantom in the merged image). Fig. 6 displays the merged image, generated from all the k_y lines in two spin-echo EPI scans (Fig. 6a,b), and the reference image acquired using a T_1 scan (Fig. 6c,d). After averaging 49 images, resulting in an SNR of 9.4, the EPI image exhibits the periodic line pattern only on the right part of the 1-mm phantom. The line pattern is more consistently apparent in the 2-mm phantom, although both are far from the clarity of the pattern in the T_1 reference image. Note that due to the fact that the partial acquisition simulation is done on this EPI data set, we cannot expect to reconstruct better results.

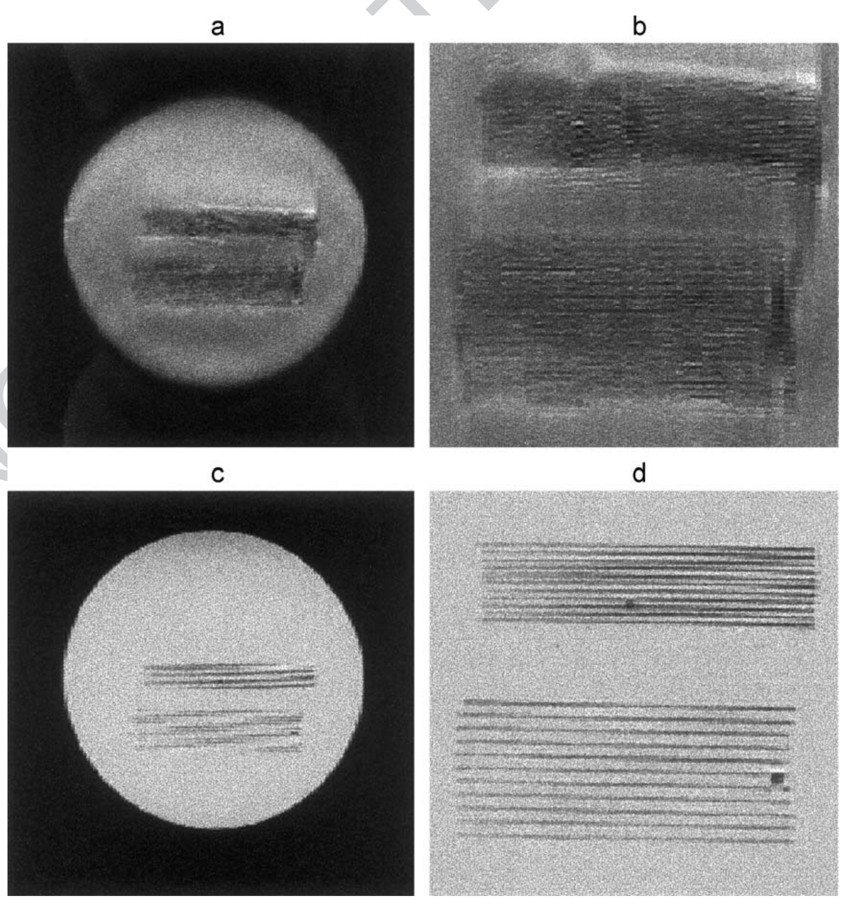

Fig. 6. (a) Echo-planar imaging acquisition results (SNR is 9.4): merged image (128×512 pixels); (b) zoom on the structure in the merged image (65×150 pixels); (c) T_1 acquisition results: entire image (512×512 pixels); (d) zoom on structure (260×150 pixels).

4.4. Simulation of nonuniform sampling

Deciding which k_y lines to acquire was accomplished as discussed in the Theory Section. Analysis was performed relative to a reference (T_1-weighted) image (i.e., model image). Linear approximation was used to initialize a simulated annealing optimization. The resulting k_y lines (partial acquisition strategy) were 'sampled' from the merged image. These data were fed to the postprocessing algorithm.

4.5. Projections used with POCS

The projections were designed carefully so as not to bias the results toward the reference image in a trivial manner (e.g., bound the distance of the restored image from the reference image). Each projection contributes genuine (though approximate) information on the image to restore. The following projections were used with the POCS scheme:

1. Bounded support constraint: The support of the image was estimated using the reference image.
2. Real signal and nonnegative signal constraints: Note that in general, EPI-acquired images are not real and positive, but this projection assumes that a phase correction algorithm was previously performed. Alternately, it is possible to incorporate into POCS the phase correction as a projection [21], prior to the real signal and nonnegative signal projections.
3. Bounded energy constraint: The image's energy was restricted to be not higher than the energy of the reference image.
4. Smoothness constraint: The energy of the image's partial derivatives (approximated by a finite difference) was restricted to lie below a threshold estimated using the reference image.
5. Lorentzian distribution of the partial derivatives: The nonconvex constraint was set by estimating the Lorentzian distribution parameters of the partial derivatives of the reference image.
6. Bounded MSE constraint: The image was restricted to agree with the acquired data up to a norm determined by the 'distance' of the reference image to the acquired data. Three projections were designed. The first constrained the acquired data in the low frequencies of k-space. The second constrained the acquired data in the high frequencies of k-space. The third constraint was applied on the low-resolution reference scan. The main

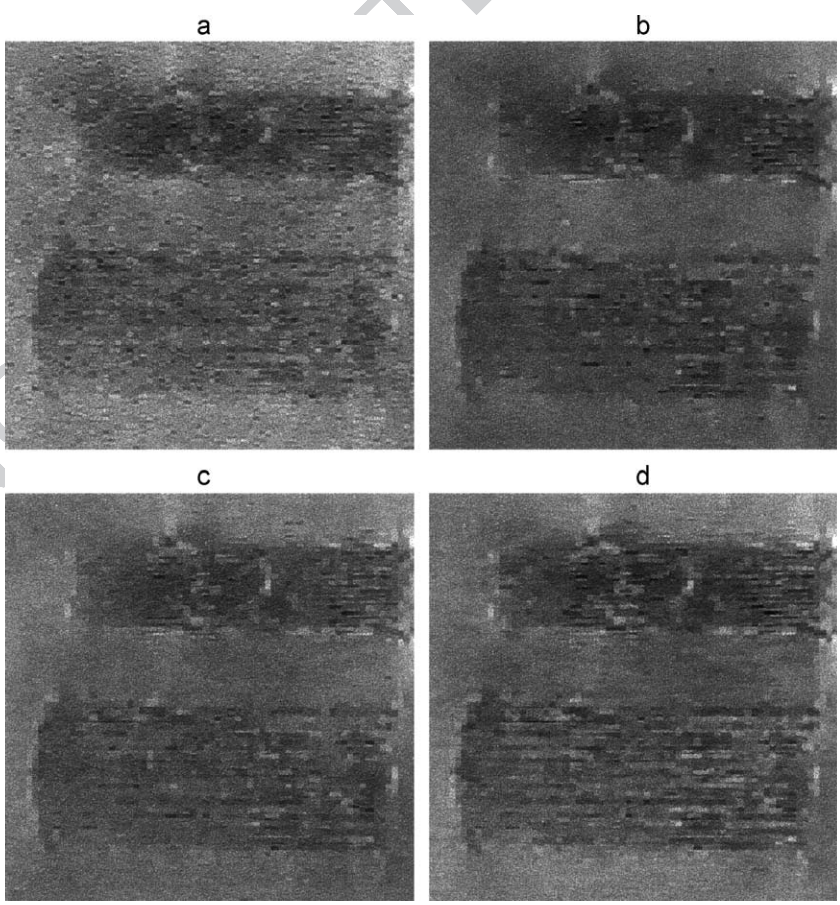

Fig. 7. Restoration using 15% partial acquisition, zoom on restored structure (65×150 pixels). SNR (of merged image) in the phantom region is 2 (a), 4.3 (b), 7 (c) and 9.43 (d).

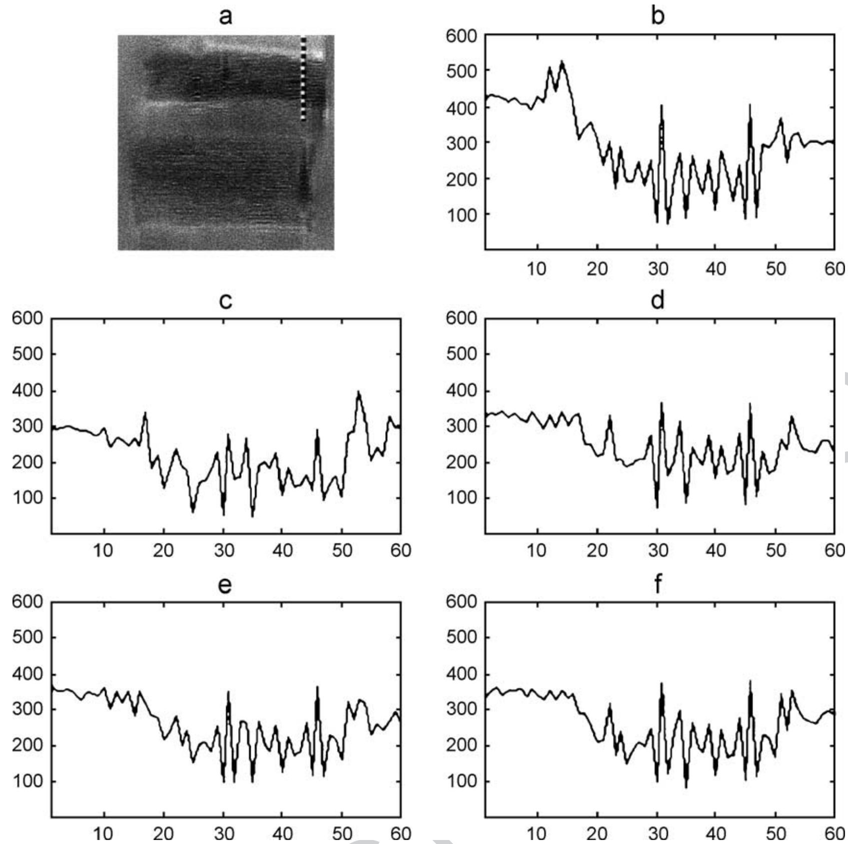

Fig. 8. Cross-section intensity vs. pixel number in acquired image (SNR is 9.43) and restored images. The cross-sectional profile is perpendicular to the small grating. The profile's location is marked in a. Profile of fully acquired grating is shown in b; profiles of restored images is shown in graphs c–f using the following acquisition percents: 10% (c), 15% (d), 20% (e) and 30% (f).

intention of this low-resolution constraint is to provide a suitable 'background' for the details.

4.6. Results

Fig. 7 displays the restoration results using 15% partial acquisition at different levels of SNR. After a certain point, SNR does not appear to affect the quality of reconstruction. Fig. 8 displays cross-sectional profiles through the small grating for different acquisition percentages. Gradual accuracy is achieved by increasing the acquisition percent. However, the peak locations and the prime features of the pattern can be processed from the 15% partially acquired and restored image. Even by using a 10% partial acquisition, one can compute approximations to the peak locations, the number of peaks and the periodicity of the pattern. In other words, for pseudoperiodic models, it is possible to find a small fraction of the DFT coefficients that captures most of the energy.

5. Discussion

In this section, we discuss the qualities of the proposed technique. First, we compare its scan-time reduction capabilities, comparing to other widely used techniques. Then, we elaborate on some issues regarding the practical use of this method and its application for fMRI uses.

5.1. Scan-time reduction

The acquisition and restoration technique we used in this work allowed a scan-time reduction of up to 90% with an acceptable level of restoration, depending on the SNR. This significant reduction is possible due to the assumption of pseudoperiodic patterns in the MR image that can be well localized in the Fourier domain. A high-resolution restoration was achieved without a high-resolution reference (as assumed in Ref. [12]) by relying on an approximate model. In Ref. [17], nonuniform acquisition of the k-space is used to reduce the scan time by up to ~52%, even then, at the expense of fine details. In Ref. [22], 75% scan-time reduction is achieved in MR spectroscopic imaging (MRSI), with a significant loss of detail. This apparent high scan-time reduction is achieved due to the flexibility in the choice of k-space coefficients to acquire, which is inherent in the acquisition of MRSI data. Fast imaging methods usually limit the choice of coefficients that can be efficiently acquired; EPI, for example, is limited to rows in k-space. In Ref. [7], up to 87.5% of scan-time reduction is used for a small (16×16 pixel) phantom. The SNR of the used images was relatively high (SNR=26). This technique is not practical for larger images because the acquired phase-encode lines are chosen by brute force search. In Ref. [8], a branch and bound optimization is used on larger images. In that work, lower scan-time reductions are reported. In Ref. [9], better performance of this method is reported when the KL basis vectors are acquired directly at the expense of SNR. Obviously, given a small number of phase-encode lines, MSE has to be sacrificed for resolution. Our method directs the search of optimal phase-encode lines suitable for the acquisition of images containing fine, pseudoperiodic details. Although our approach supplies an optimal linear extrapolator, we do not use it for extrapolation because of its low-quality restoration and its dependence on the precision of the autocorrelation matrix of the data. Nevertheless, it is optimal and it outperforms any other linear extrapolator with the same acquisition mode. Instead of the linear filter, a POCS algorithm is used for k-space extrapolation. This allows the use of many useful prior knowledge constraints, improving the quality of extrapolation.

5.2. Issues regarding the practical implementation of the method

Three issues may seem to limit the utility of the technique presented in this work:

1. Calculation of the covariance matrix of the data.
2. Setting the correct bounds and thresholds in the POCS constraints.
3. Implementing the technique for nondirectional (isotropic) pseudoperiodic patterns.

It was shown above that the covariance matrix can be computed from an image of a weighting different from the acquired image. This is probably also true for images of different modalities, as long as image-processing techniques are used to register between the modalities. In the extreme case, where no reference image can be achieved, the covariance matrix can be approximated. For instance, a 'typical' reference image can be constructed by using all the prior knowledge of the image, for example, energy of the signal, finite support of the signal, mean value of pixels, etc. This reference image can be used to compute the covariance matrix. The covariance matrix is needed for determining the acquisition strategy and for the linear extrapolation. If the covariance matrix used is a rough approximation of the correct covariance matrix, the linear extrapolation estimate will be highly inaccurate. In this case, it is better not to use the linear extrapolation estimate as a reference image constraint in the POCS algorithm. We have demonstrated that our POCS algorithm is rather independent of its initialization and the use of the linear extrapolation estimate as a reference image. Consequently, the covariance matrix should be used only for determining the partial acquisition strategy (in this case, the phase-encode lines to acquire) and might therefore be approximated using the prior knowledge on the image.

The second issue involves the POCS algorithm. Using incorrect prior knowledge can cause the POCS algorithm to diverge or to converge to a nonacceptable solution. However, we showed that a limited relaxation of the POCS constraints does not greatly affect the quality of the restored image. Therefore, the thresholds used in constraints should be relaxed if their value is not certain. The POCS iterations themselves offer information regarding the correct value of the thresholds. Changing a threshold, which causes POCS to diverge, implies too strict a constraint. In contrast, changing of a threshold, causing the constraint to be never active (size of projection is zero), implies too relaxed a constraint. The convergence properties of POCS can be inferred from the norm of the difference between consecutive iterations and by the sum of projection sizes. The process of tuning the parameters can be automated, using the heuristics described above.

The third issue is dependent on the acquisition method and on the geometric shape of the pseudoperiodic pattern. In all the simulations, we assumed a rectangular acquisition of the 2D DFT and that the pattern had a clear orientation. In the case of an isotropic pattern, it is preferable to use a spiral acquisition methodology. The only part of the proposed technique, which is not trivial to adapt to spiral imaging, is the selection of the DFT coefficients to acquire. Finding an optimal method for selecting the DFT coefficients to acquire with spiral imaging is an issue for further research.

5.3. Using the method for fMRI

The initial motivation for this work was to enable high-resolution (submillimeter) fMRI. This is not possible in general, but this work suggests that pseudoperiodic patterns can be acquired at a high resolution using a small number of phase-encode lines. Pseudoperiodic patterns are abundant in parts of the mammalian cortex that are organized in functional columns. Using a low number of phase-encoding lines (~10% of the k-space), a large ensemble can be acquired and used for averaging (increasing the SNR). In the Computer Simulations Section, we used an image in which the pseudoperiod was apparent. Obviously, in fMRI experiments, the pseudoperiodic pattern appears only in averaged difference images. Because of the linearity of the Fourier transform, the application of our technique is straightforward: partially acquire the DFT coefficients of each 'state', form the state's difference and average the difference images. Note that the calculation is done upon the partially acquired DFT coefficients. The resulting image is the input to the POCS algorithm. Clearly, the statistical model of the signal and the prior knowledge, which is used in POCS, should be based on the averaged difference image and not on the original signal. Other statistical analysis techniques (e.g., correlation with a time sequence) can be implemented as well.

6. Summary

In this paper, a new technique for partial acquisition and reconstruction of MR images was suggested and demonstrated. While being general, the technique was specifically designed to detect a spatial pattern by utilizing the attributes of pseudoperiodic patterns in order to allow high reduction of the scan-time. This approach represents a new MRI research methodology in which an experiment is designed to allow detection of pseudoperiodic spatiotemporal patterns that characterize, for example, the activity patterns of cortical columns. The detection ignores all the details not related to the pattern—this contrasts the typical methodology according to which analysis is made after a complete acquisition. Utilizing the proposed technique, it has been shown that the quality of the restored images is acceptable even at a very low SNR. The algorithm relies on prior knowledge in the form of constraints that can be relaxed without a significant degradation of the restored image. The algorithm presented is practical and efficient enough to be implemented in real time, and thus offers a new option for visualizing pseudoperiodic patterns in general.

Acknowledgment

The authors thank Dr. T. Hendler and Dr. M. Graif for their support.

References

[1] Purves D, Riddle DR, LaMantia AS. Iterated patterns of brain circuitry (or how the cortex gets its spots). Trends Neurosci 1992;15(10):362–8.
[2] Horton JC, Dagi LR, McRane EP. Arrangement of ocular-dominance columns in human visual cortex. Arch Ophtalmol 1990;108:1025–31.
[3] Menon RS, Bradley GG. Submillimeter functional localization in human striate cortex using BOLD contrast at 4 Tesla: implications for the vascular point-spread function. Magn Reson Med 1999;41:230–5.
[4] Albright TD, Robert D, Gross CG. Columnar organization of directionally selective cells in visual area MT of the macaque. J Neurosci 1984;51:16–31.
[5] Yeshurun Y, Schwartz EL. Cepstral filtering on a columnar image architecture: a fast algorithm for binocular stereo segmentation. IEEE Trans Pattern Anal Mach Intell 1989;11:759–67.
[6] Therrien CW. Discrete random signals and statistical signal processing. Prentice-Hall.
[7] Cao Y, Levin DN. Feature-recognizing MRI. Magn Reson Med 1993;30:305–17.
[8] Cao Y, Levin DN. Locally focused magnetic resonance imaging. Proc-Int Conf Image Proc 1994.
[9] Weaver JB, Healy Jr DM. Acquisition of the Karhunen–Loeve expansion to reduce MR imaging times. Proc-Int Conf Image Proc 1994.
[10] Reeves SJ. Selection of observations in magnetic resonance spectroscopic imaging. Proc-Int Conf Image Proc 1995:641–4.
[11] Boiman, O. Partial Acquisition of MR Images with Pseudo Periodic Patterns. M.Sc. thesis, Tel-Aviv University, Israel; 2001.
[12] Liang Z-P, Lauterbur PC. A generalized series approach to MR spectroscopic imaging. IEEE Trans Med Imag 1991;10(2):132–7.
[13] Youla DC, Webb H. Image restoration by the method of convex projections: Part 1 — Theory. IEEE Trans Med Imag 1982;MI-1(2):81–94.
[14] Sezan MI, Stark H. Image restoration by the method of convex projections: Part 2 — Applications and numerical results. IEEE Trans Med Imag 1982;MI-1(2):95–101.
[15] Stark H, editor. Image recovery: theory and applications. Academic Press; 1987.
[16] Mammone RJ. Computational methods of signal recovery and recognition. Wiley.
[17] Marseille GJ. MRI scan time reduction through non-uniform sampling. Doctoral thesis, Technical University of Delft, The Netherlands; 1997.
[18] Gill PE, Murray W, Wright MH. Practical optimization. Academic Press.
[19] Levi A. Image restoration by the method of projections with applications to the phase and magnitude retrieval problems. Ph.D. dissertation, Department of Electrical, Computer and System engineering. Troy (NY):Rensselaer Polytechnic Institute; 1983.
[20] Margosian P. Proceedings of 4th SMRM Conference. 1985. p. 1024.
[21] Haacke EM, Lindskog ED, Lin W. A fast, iterative, partial-Fourier technique capable of local phase recovery. J Magn Reson 1991;92:126–45.
[22] Plevritis SK, Macovski A. MRS imaging using anatomically based k-space sampling and extrapolation. Magn Reson Med 1995;34:686–93.