k-t FASTER: Acceleration of Functional MRI Data Acquisition Using Low Rank Constraints

Mark Chiew,[1]* Stephen M. Smith,[1] Peter J. Koopmans,[1] Nadine N. Graedel,[1] Thomas Blumensath,[2] and Karla L. Miller[1]

Purpose: In functional MRI (fMRI), faster sampling of data can provide richer temporal information and increase temporal degrees of freedom. However, acceleration is generally performed on a volume-by-volume basis, without consideration of the intrinsic spatio-temporal data structure. We present a novel method for accelerating fMRI data acquisition, k-t FASTER (FMRI Accelerated in Space-time via Truncation of Effective Rank), which exploits the low-rank structure of fMRI data.
Theory and Methods: Using matrix completion, $4.27\times$ retrospectively and prospectively under-sampled data were reconstructed (coil-independently) using an iterative nonlinear algorithm, and compared with several different reconstruction strategies. Matrix reconstruction error was evaluated; a dual regression analysis was performed to determine fidelity of recovered fMRI resting state networks (RSNs).
Results: The retrospective sampling data showed that k-t FASTER produced the lowest error, approximately 3–4%, and the highest quality RSNs. These results were validated in prospectively under-sampled experiments, with k-t FASTER producing better identification of RSNs than fully sampled acquisitions of the same duration.
Conclusion: With k-t FASTER, incoherently under-sampled fMRI data can be robustly recovered using only rank constraints. This technique can be used to improve the speed of fMRI sampling, particularly for multivariate analyses such as temporal independent component analysis. **Magn Reson Med 74:353–364, 2015.** © **2014 Wiley Periodicals, Inc.**

Key words: fMRI; k-t acceleration; compressed sensing; low-rank acceleration; matrix completion; resting state networks

INTRODUCTION

Strategies for accelerating functional MRI (fMRI) data acquisition have received strong interest since the technique was introduced for neuroimaging. Faster sampling of blood oxygenation level dependent (BOLD) signals provides improved statistical power and richness of the temporal modeling in brain dynamics measured with fMRI (1). Furthermore, higher sampling bandwidths can

[1]FMRIB Centre, University of Oxford, Oxford, United Kingdom.
[2]IVSR, University of Southampton, Southampton, United Kingdom.
*Correspondence to: Mark Chiew, Ph.D., FMRIB Centre, John Radcliffe Hospital, Oxford UK OX3 9DU. E-mail: mchiew@fmrib.ox.ac.uk

reduce aliasing of physiological noise (2) and provide finer characterizations of hemodynamic responses.

One application that gains significantly from accelerated whole brain fMRI is the estimation of brain connectivity, for example, in resting state networks (RSNs) (3,4). Higher sampling rates can reduce the imaging durations required to achieve RSN estimation (to a limit, dictated by the low-frequency nature of the fluctuations), or more importantly, improve estimation of RSNs with greater temporal dimensionality over the same total scan time. Often, fMRI data are modeled as a linear mixture of components (commonly using statistically independent component models), and robust unmixing of these RSN components can require many time points (5). Additionally, recent work has suggested that higher sampling rates can reveal frequency-specific RSN characteristics unavailable to methods with lower sampling bandwidths (6).

Acceleration in fMRI has until recently been predominantly parallel-imaging based, working on an image-by-image basis without exploiting any structure in the temporal domain. The classic approaches of SENSE (7) and GRAPPA (8) can shorten echo train lengths in two-dimensional (2D) echo-planar imaging (EPI) acquisitions, although these methods only enable small reductions in volume measurement times. Faster sampling with 3D EPI can be achieved using parallel imaging in both phase-encoded directions (9). Methods that use coil-sensitivity information in conjunction with regularized inverse problem solvers (10,11) have demonstrated extremely fast sampling rates of whole brain volumes at 100 ms resolution. These methods, however, can suffer from low (and spatially variable) spatial resolution. Simultaneous multislice imaging (1,12,13) has emerged as a successful acceleration strategy that exploits coil sensitivity information without major signal-to-noise ratio (SNR) penalties or loss of spatial resolution. Inevitably, however, all of these methods have limits on acceleration, dependent on coil numbers and geometry.

In contrast to time-independent methods, k-t acceleration uses information across both space and time to accelerate imaging. In other words, k-t methods exploit redundancy and/or structure in k-t space (or equivalently in x-f space) to reduce sampling requirements. One of the first such methods demonstrated in fMRI, UNFOLD (14), shares signal bandwidth across multiple aliased spatial locations. The k-t BLAST method (15) further exploits spatio-temporal structure and prior information to reduce sampling requirements in dynamic imaging. While these methods have been successful in applications such as cardiac imaging, where the signals of interest have strong

periodicity, they have not been widely applied in fMRI due to its broad spectrum temporal characteristics (16).

Compressed sensing methods (17) can also exploit properties of k-t space to facilitate under-sampling, but rather than requiring a specific distribution of the x-f signals in relation to the signal aliasing patterns, acceleration is facilitated by a more general consideration of data sparsity under a known transform. Although sparse k-t imaging methods date from the very onset of compressed sensing in MRI (18), adoption in fMRI has been limited due to considerable low-frequency signal content which limits sparsity in the temporal frequency domain. The k-t FOCUSS method (19) combines the features of prior information (k-t BLAST) with sparsity-promoting reconstruction (compressed sensing), and has been applied to fMRI using a principal component analysis (PCA)-based temporal sparsifying transform (20,21). Some groups have also used compressed sensing in a time-independent manner for accelerating fMRI on an image-by-image basis, focusing only on its limited spatial sparsity characteristics (22,23).

Another class of k-t methods fit low-rank models to the data, assuming that the measured data can be represented as a summation of a small number of spatial maps modulated by associated time-series (24,25). These methods do not require any specific temporal structure, and use training data to estimate temporal components (similar to k-t FOCUSS) followed by estimation of the associated spatial components. Under-sampling is possible because a low-rank model has fewer unknowns than a full-rank model, and hence requires fewer data points to recover. These methods, however, require robust estimation of a temporal basis before reconstruction, which can be difficult using a separate acquisition or a subset of the data.

Recently, a new mathematical framework for reconstructing randomly under-sampled low-rank matrices has been developed, called matrix completion (26). Analogous to compressed sensing, matrix completion allows recovery of a randomly under-sampled low-rank matrix. In contrast to the low-rank model-fitting methods described above, matrix completion estimates the temporal and spatial subspaces simultaneously, and no training data are required (27). Furthermore, the theoretical matrix completion problem has been studied with noisy sampling (28), with evidence suggesting that it can be more robust than other methods with data that are only approximately low-rank. A general framework using this approach has been developed for accelerated MR imaging (29,30), and has recently also been explored in task-based fMRI (31).

In this study, we introduce the novel application of matrix completion to accelerate the acquisition of resting state fMRI data, using a 3D sampling strategy and a new reconstruction approach (32). This is motivated by the observation that resting fMRI data are nearly always modeled as a mixture of distinct temporal (functional) processes that are distributed across space in functional networks, which is precisely the definition of a low-rank space-time (k-t) matrix. In fact, independent component analysis (ICA)-based analysis of resting state fMRI data often involves PCA dimensionality reduction (33), which

explicitly enforces the low-rank representation on the fMRI data matrix by discarding low variance principal components, assumed to be primarily unstructured (e.g., thermal) noise. In effect, matrix completion is used to directly estimate the principal components (spatial maps and time-series) of the fMRI data.

We call this approach k-t FASTER: FMRI Accelerated in Space-time via Truncation of Effective Rank. We explore its efficacy in recovering k-t matrices and estimating RSNs without coil sensitivity information in both retrospective and prospective under-sampling experiments with a 4.27× acceleration factor.

THEORY

The proposed k-t FASTER acceleration method relies on low-rank matrix completion strategies (28) to reconstruct under-sampled image time-series data. What makes fMRI so well suited to this type of acceleration is that its data are inherently low rank, composed of a relatively small number of spatially coherent temporal processes (i.e., spatial maps and associated time-series), so that each voxel carries information about many other voxels. In reality, random (e.g., thermal) and structured (e.g., physiological, motion) noise sources are also present, causing the k-t matrix to be only approximately low rank (although many of these can also be well represented as low rank) (34).

Figure 1 illustrates the difference between standard acquisition strategies and the proposed k-t FASTER strategy. In one common approach to resting-state fMRI, the rank r k-t matrix is fully (and inefficiently) sampled and then PCA dimensionality reduction is used to "compress" the data before applying ICA (35). In this context, compression refers to the fact that the resulting (denoised) matrix can be represented with fewer coefficients. Recognizing the compatibility of PCA compression with the final fMRI analysis, k-t FASTER efficiently under-samples the acquired k-t matrix and uses matrix completion methods to directly approximate the principal components. Alternatively, k-t FASTER can be viewed as a technique that uses nonlinear algorithms to "fill-in" the nonsampled k-t matrix locations using a low-rank constraint. This matrix can then be used (for example) to estimate RSNs by searching the dimensionality-reduced space for independent components, which have been shown to correspond to functionally meaningful networks of brain activity (36).

Matrix Completion

The matrix completion (MC) problem can be intuitively understood by comparison to compressed sensing (CS). In CS, under-sampled images are recovered by assuming the data can be represented in a known basis using a limited number of non-zero coefficients, i.e., using "sparsity" to constrain reconstruction. In MC, the property of matrix rank replaces sparsity, where a matrix of low rank has few non-zero singular values. Thus the concept of compressibility in CS for sparse signals applies similarly to low-rank matrices: they contain less information, or fewer degrees of freedom (DOF), than their full-rank counterparts and can therefore be under-sampled.
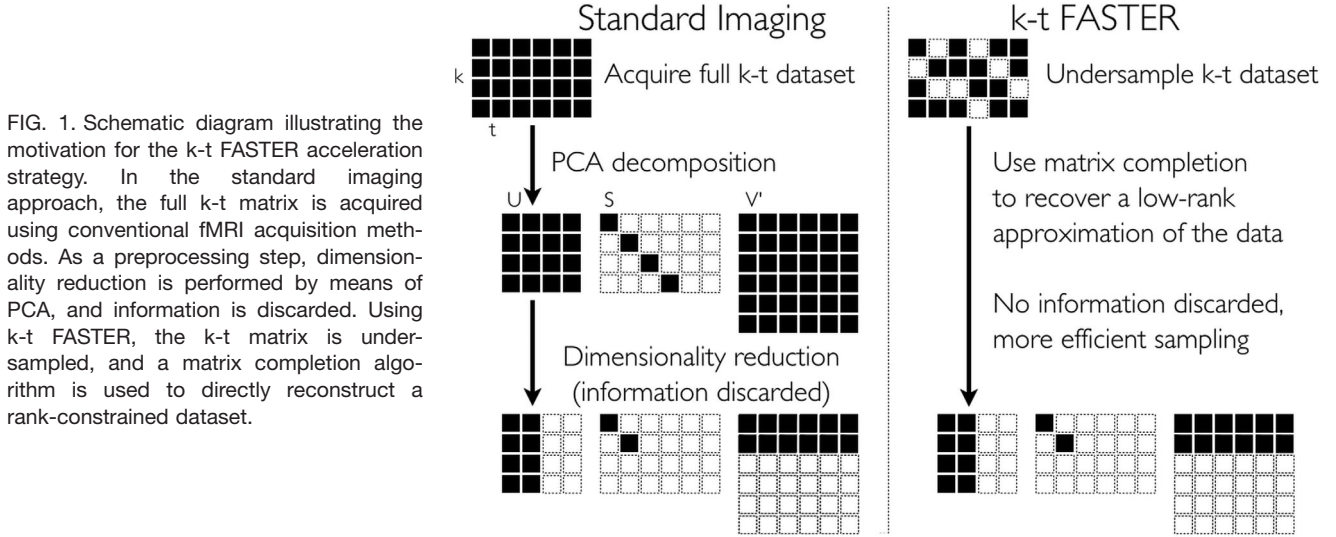
FIG. 1. Schematic diagram illustrating the motivation for the k-t FASTER acceleration strategy. In the standard imaging approach, the full k-t matrix is acquired using conventional fMRI acquisition methods. As a preprocessing step, dimensionality reduction is performed by means of PCA, and information is discarded. Using k-t FASTER, the k-t matrix is undersampled, and a matrix completion algorithm is used to directly reconstruct a rank-constrained dataset.

However, where CS typically requires a prespecified basis and estimates sparse coefficients, MC has no such requirement.

Following the framework laid out in previous low-rank imaging work (24,29,30), consider the single-channel low-rank signal model:

$$X = U\Sigma V^* + N \qquad [1]$$

where U is an $m \times r$ matrix of normalized spatial components (maps), V is an $n \times r$ matrix normalized temporal components (time-courses), Σ is an $r \times r$ diagonal matrix of weights, and N is an $m \times n$ matrix of additive noise. The first term in Eq. [1] is an r-component outer product model, which is written as a rank r singular value decomposition (SVD) of a matrix with m voxels and n time points. Because the SVD can be used to compute the PCA decomposition of a matrix, Eq. [1] shows the link between low-rank signal models and PCA dimensionality reduction.

The total DOF in the signal are the number of independent coefficients in the SVD: $r(m + n - r)$. Notably, with an undersampling factor R, each additional measured time point adds m/R additional samples, but only r additional DOF, meaning that the ratio of samples to unknown matrix DOF decreases with increasing n (assuming $m/R > r$, which is typical of imaging datasets). Although theoretical sampling requirements have been characterized (37), practical evaluations of matrix recovery algorithms have investigated performance in regimes where the number of samples is nearly the same as the total DOF (38).

In this single channel model, acceleration is limited by the intrinsic structure of the data, not by external quantities such as the number or geometry of receiver coils. A plot of the singular values, ordered by decreasing variance, for a set of representative fMRI data are shown in Figure 2. The moderately low-rank structure is apparent, with the non-zero asymptote reflecting the presence of noise.

Iterative Hard Thresholding with Matrix Shrinkage

Several strategies have been developed to solve the MC problem. Early work in low rank dynamic MRI used a matrix factorization method (29,30) that constructed low rank estimates by iteratively incrementing the rank of estimates until the target rank is reached. This algorithm however, was designed for substantially lower target ranks (e.g., $r = 5$ for dynamic contrast enhanced breast imaging) than appropriate for fMRI, where analyses typically aim to estimate tens to hundreds of functional components, with higher dimensionalities reflecting greater richness of information about brain dynamics (39).

In this study, we introduce the use of reconstruction approach based on the generalized Iterative Hard Thresholding algorithm (40), which solves the constrained optimization problem:
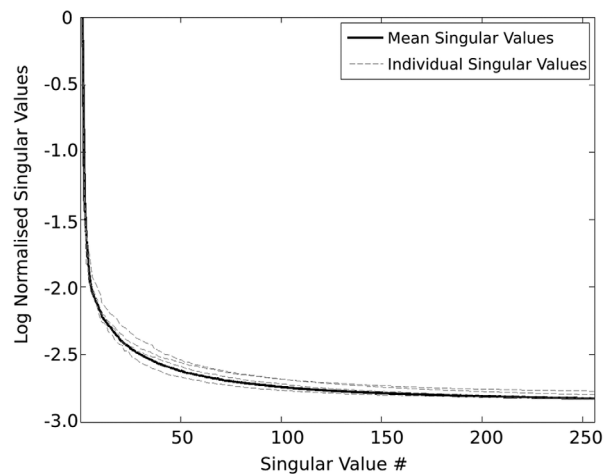

FIG. 2. Normalized singular values from fMRI data with 512 time points. These singular value decays reflect approximately low rank nature of fMRI data, where after mean centering, $>80\%$ of the signal variance is captured in the top 25% components. Dashed lines represent fMRI data from six different subjects, with mean values indicated with the solid line.

$$\min \| Y - \Phi\left(\hat{X}\right) \|_2$$
$$s.t. \ rank\left(\hat{X}\right) = r \qquad [2]$$

where $\hat{X}$ is the matrix estimate, Y is a vector of measurements, and Φ is a measurement operator that provides vectorized matrix samples. The solution to this problem comes from drawing the optimal estimate from the set of all rank r matrices (a union of linear subspaces) (41) by means of orthogonal projections, and is calculated by the following procedure:

$$X^{n+1} = S_r(X^n + \mu(\Omega(X) - \Omega(X^n))) \qquad [3]$$

where X^n is the n^{th} matrix estimate, μ is a step size parameter, Ω is the matrix sampling operator, and S_r is the hard thresholding (orthogonal projection) operator:

$$\sigma_i \rightarrow \begin{cases} \sigma_i & if \ i \leq r \\ 0 & otherwise \end{cases} \qquad [4]$$

where σ_i is the i^{th} singular value of the matrix, and r is the fixed rank cutoff.

An alternative approach to the direct optimisation of the nonconvex rank constraint in Eq. [2] is a relaxed approach that replaces the nonconvex rank constraint with a convex nuclear norm constraint (42), which has several advantages. A simple approach to the optimization of this relaxed cost function is very similar to Eq. [3], but uses the alternative shrinkage operator:

$$\sigma_i \rightarrow \begin{cases} \sigma_i - \tau & if \ \sigma_i \geq \tau \\ 0 & otherwise, \end{cases} \qquad [5]$$

where τ is some predefined parameter.

To combine advantages of both of these approaches, the thresholding procedure was modified to incorporate a shrinkage step which changes the thresholding operator (Eq. [4]) to:

$$\sigma_i \rightarrow \begin{cases} \sigma_i - \tau & if \ i \leq r \\ 0 & otherwise \end{cases} \qquad [6]$$

where τ is an additional parameter which shrinks the remaining singular values by a fixed amount. Here, we use a floating threshold, so that $\tau = c\sigma_{r+1}$ with a proportionality constant $0 < c \leq 1$ (see Appendix A1 for algorithm details). Note that when $c = 0$, this reduces to the standard hard thresholding operator. This modification makes the algorithm similar to the Fixed Point Continuation approach developed by Goldfarb and Ma (43), which also makes use of both hard thresholding and soft shrinkage operations, and has demonstrably better performance in recovery of matrices that have moderate (instead of very low) rank. The floating shrinkage parameter is also similar to the nonconvex spectral penalty introduced in Lingala et al (44). The final step of the algorithm is replacement of the originally sampled data back into the matrix estimate, which was found to improve reconstruction fidelity. However, this modifies the distribution of singular values, and the final matrix estimate is no longer strictly rank r (see Supporting Fig. S1, which is available online).

One advantage of the iterative hard threshold and matrix shrinkage (IHT+MS) algorithm is its simplicity, as well as good empirical performance in fMRI data, which has (a) approximately low rank structure, (b) very large matrix sizes, and (c) high target ranks. With target ranks on the order of 10^2, incremented rank methods can be inefficient, whereas IHT+MS can use SVD approximation methods to speed up computation and reduce memory requirements. As fMRI k-t matrices tend to have a very large spatial dimension and smaller temporal dimension, the singular values and right singular vectors of X can be computed from the eigenvalues and eigenvectors of $X^H X = V\Sigma^2 V^*$, which is computationally inexpensive due to the small temporal dimension. The remaining left singular vectors can be estimated by computing $U = XV\Sigma^{-1}$.

METHODS

The k-t FASTER method was evaluated in two ways: using retrospectively under-sampled data, and in experiments using a prospectively under-sampled 3D EPI acquisition.

Imaging Parameters

Two datasets of six subjects each were collected on healthy subjects with informed consent in accordance with local ethics. The first dataset used a simultaneous multislice EPI acquisition (1,12) to enable retrospective sampling of data with a known true k-t matrix at high temporal resolution. These data were acquired on a 3 Tesla (T) (Siemens Verio, 32-channel head coil) with an acceleration factor of 8 to achieve whole-brain coverage with volume TR = 836 ms, spatial matrix size of $106 \times 106 \times 64$ at 2 mm isotropic resolution, and 1075 time points. These data were used to simulate under-sampled 3D acquisitions similar to our second datasets.

The second dataset used prospective under-sampling to provide a realistic comparison of k-t FASTER with a duration-matched fully sampled acquisition. These data were collected on a Siemens 7T using a 32-channel head coil (Nova Medical) with a 3D segmented EPI acquisition modified to allow arbitrary under-sampling of the k_z dimension. The 3D EPI was acquired using an in-plane acceleration factor of 2 to achieve an appropriate TE for fMRI at 7T, and was reconstructed in-plane using GRAPPA (8), independently of and before MC. Potentially, g-factor noise amplification could produce sub-optimal reconstructions by increasing error bounds that depend on noise characteristics. However, in these data collected with a low acceleration factor and 32-channel coil geometry, noise amplification is not expected to have a significant impact. The acquisition used a spatial matrix size of $100 \times 98 \times 64$ at 2mm isotropic resolution, with 72 time points in the fully sampled case and 304 time points for the accelerated data. Each subject was scanned twice with matched 5-min scan durations in the fully sampled and accelerated cases respectively.

Retrospective Sampling and Reconstruction

We explored retrospective sampling and reconstruction of k-t matrices to evaluate the k-t FASTER approach in

FIG. 3. Representative diagrams of the different k_z-t under-sampling patterns studied. Pseudorandom under-sampling (**a**), sheared grid under-sampling (**b**), and temporal under-sampling (i.e., standard low-temporal resolution) (**c**). In all cases, in-plane (k_x-k_y) data are fully sampled.
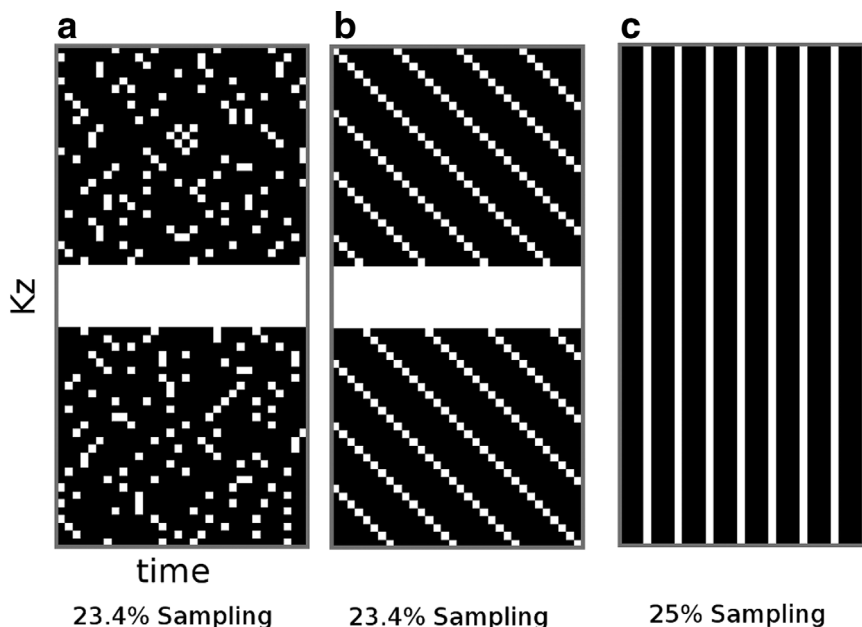
comparison with three additional methods under conditions where the full k-t matrix is known. We identify this as "ground truth" data, although because it is real measured data it still contains noise and artifacts. For all methods, these data were Fourier transformed along the spatial dimensions to produce a simulated 3D k-t space for retrospective under-sampling and subsequent reconstruction. These simulations mimicked our prospectively under-sampled data by under-sampling in the k_z dimension (i.e., measuring planes in k_x, k_y).

For k-t FASTER reconstruction, sampling was performed by selecting a pseudo-random subset of k_z planes. Sampling of 23.4% was achieved by keeping 8 central k_z planes and 7 pseudo-randomly selected outer k_z planes (in total 15 of 64 planes) from each time point (Fig. 3a). Reconstruction of the data was performed using the IHT+MS algorithm with target rank of 128, chosen a priori. The shrinkage thresholding parameter c was set to 0.5, which was tuned to produce the best average results in the retrospectively sampled data, with a step size μ of 0.8.

Acceleration by means of the "keyhole" method uses data sharing or interpolation to recover missing entries (45) and continues to receive attention as an acceleration technique for fMRI (46). We performed a keyhole-type reconstruction, denoted here as "k-t INTERP," of our pseudo-randomly under-sampled data (fully sampled central k-space, under-sampled outer k-space) by linearly interpolating across time for every under-sampled k-space location to estimate the missing data.

We also compared k-t FASTER with the linearly reconstructed low-rank partially separable functions (PSF) model (24), which we call "k-t PSF." In the k-t PSF approach, a central portion of fully sampled k-space is used to estimate a temporal basis, which is then used to least squares fit the grid sampled points. The central 8 k_z planes were used as training data, combined with sheared grid sampling with an 8× under-sampling factor (Figure 3b, overall sampling factor of 23.4%). A rank

constraint of 26 was used for the matrix reconstruction error and RSN data analysis, which produced the lowest Frobenius norm errors after comparison with a range of candidate target ranks (see Supporting Fig. S2 for a comparison of k-t PSF reconstruction error versus rank).

A final comparison with a simulated long-TR fully (spatially) sampled acquisition of the same total imaging duration, which we label "k-t SLOW." To generate this, the source data were temporally decimated by selecting every 4th time point (Fig. 3c), resulting in 25% sampling of the full k-t matrix.

Prospective Sampling and Reconstruction

In the prospectively under-sampled data, the k-t FASTER method using pseudo-random under-sampling was compared with the slow, fully sampled 3D EPI acquisition (k-t SLOW), in separate acquisitions from the same subjects. The k-t FASTER reconstruction here differs slightly from the retrospectively sampled case, using a lower target rank of 72 due to the reduced number of time points in the acquisition. Furthermore, reconstruction was performed on a coil-by-coil basis, and then sum-of-squares combined. In all other respects, except where explicitly noted, sampling and reconstruction parameters were identical to the retrospectively sampled cases.

RSN Data Analysis

Following k-t FASTER image reconstruction, the inverse Fourier transform and magnitude operation was applied to produce a magnitude image time-series, which underwent motion correction and linear registration (FSL FLIRT) (47) to a standard space for comparison. Dual regression (48,49), a two-stage multiple linear regression framework, was used to generate spatial z-statistic maps. It first performs a spatial regression of each dataset against the canonical maps to extract a time-series of regression coefficients corresponding to each map, then performs a second temporal regression of each dataset
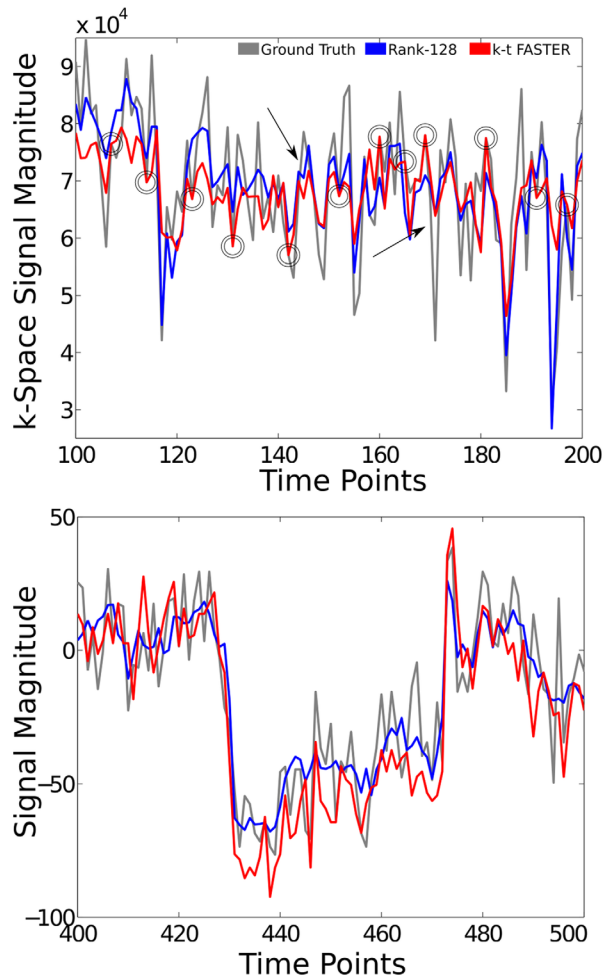
FIG. 4. **a**: Example portion of a magnitude k-space time-series from a k-t FASTER reconstruction (red), compared with the ground truth (grey) and a rank-128 ground truth time-series (blue). The k-t FASTER data reproduce many dynamic features of the ground truth in between sampled points. Arrows point to features that are common to the k-t FASTER and rank-128 time-series, suggesting that they reflect the rank constraints rather than errors intrinsic to the k-t FASTER method. **b:** Example voxel time-series segment similarly showing good recovery of dynamic features in the k-t FASTER reconstruction.

against this set of time-series. The output is a set of z-statistic maps that reflect the degree to which each spatial regressor, corresponding to a canonical RSN, is expressed with a unique time-course in the data. Data from the retrospective sampling comparisons were evaluated using dual regression against a set of 82 resting state spatial regressors (representing canonical network spatial maps) derived from high-dimensional group-level ICA of a large, high quality resting fMRI dataset from the Human Connectome Project (50,51). Output maps were null-corrected using a Gaussian and Gamma mixture model (4).

Voxel-by-voxel correlation coefficients were computed for the output z-statistic maps against input regressor maps and the ground truth output maps (retrospective sampling only) to quantify the reconstruction fidelity of each map in a single summary statistic. This metric

captures the requirement for correct attribution of unique time-courses to corresponding spatial networks during matrix recovery to correctly produce a subject-specific version of a given RSN.

RESULTS

Retrospective Sampling Results

To illustrate the low-rank matrix reconstruction fidelity, Figure 4a shows a portion of a k-space magnitude time-series from a representative dataset. The ground truth data are shown alongside a rank-128 dimensionality reduced ground truth (i.e., result of a PCA dimensionality reduction with rank 128) and the k-t FASTER time-series, with overall correlations of 0.81 and 0.72, respectively. Many transient features of the reconstructed k-space time-series can be seen between the sampled points, which would be lost by slower traditional sampling or methods that simply interpolate between the sampled points. Furthermore, although the transient features are recovered imperfectly, many of the reconstruction artefacts are shared between the k-t FASTER and rank-128 data (arrows), which emphasizes the fact the k-t FASTER output reproduces (and rejects) many of the same variance components as a rank-128 decomposition of the ground truth that might be used in RSN analysis. Similarly, Figure 4b shows good recovery of dynamic features in a real-space voxel time-series segment.

Overall k-t matrix reconstruction fidelity for the various methods was evaluated using the fractional Frobenius norm error, defined as:

$$\mathrm{err}_{\mathrm{F}}(\hat{X}) = 100 \cdot \frac{\|\hat{X} - X\|_{\mathrm{F}}}{\|X\|_{\mathrm{F}}} \% \qquad [7]$$

where $\hat{X}$ is the estimated matrix, and X is the ground truth. Figure 5 shows the $\mathrm{err}_{\mathrm{F}}$ values for all six retrospectively sampled datasets respectively, including the $\mathrm{err}_{\mathrm{F}}$ obtained by performing a rank-128 dimensionality reduction of the ground truth. With this metric, the k-t FASTER method produces the lowest $\mathrm{err}_{\mathrm{F}}$ for all datasets, comparable to the amount of error present in a rank 128 truncated ground truth. The k-t PSF reconstruction produced the next lowest errors, followed by k-t INTERP. A paired t-test between the $\mathrm{err}_{\mathrm{F}}$ values for k-t FASTER versus the k-t PSF method shows a significant difference ($p < 0.05$). Further analysis of matrix reconstruction fidelity can be found in Supporting Figs. S3 and S4.

Although the Frobenius norm results are very promising, the error metric is effectively scaled by the variance of the matrix components, and preferentially weights the first few major components and may not capture the fidelity of lower variance components that are important in defining RSNs. Ultimately, the reconstruction methods need to be evaluated in the context of RSN recovery using dual regression.

The dual regression results show better representation of RSN spatial maps for k-t FASTER compared with k-t PSF, k-t INTERP, and k-t SLOW. Figure 6a shows the correlations coefficients for each reconstruction method, across all six datasets and 82 dual regression maps. The grand means (± standard deviation) for each method
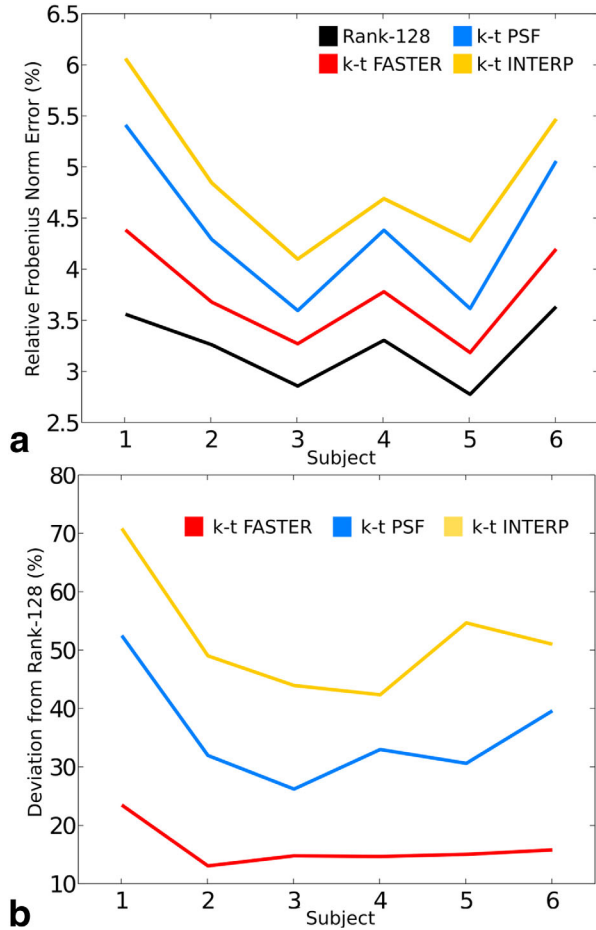
a

b

FIG. 5. **a**: Relative Frobenius norm errors for all six subjects, across the three different reconstruction methods and the rank-128 dimensionality reduction of the ground truth data. **b:** Frobenius norm errors normalized to the error in the rank-128 truth data, highlighting the excellent performance of the k-t FASTER method compared with the k-t PSF and k-t INTERP methods.

were: k-t SLOW 0.13 ± 0.05, k-t PSF 0.10 ± 0.04, k-t INTERP 0.18 ± 0.06, and k-t FASTER 0.22 ± 0.08. The mean difference between k-t FASTER and the other methods is significant (paired t-tests, $P < 0.05$ corrected). The relatively low correlations values reflect the fact that these correlations were performed against canonical maps derived from a large set of group-averaged data, which are expected to show substantial mismatch with individual subject maps.

The results in Figure 6a could be driven by smoothing or de-noising in k-t FASTER, which would inflate correlations with the canonical RSNs (which tend to be smooth as a result of averaging across a large cohort). It is thus important to ensure that dataset-specific features are being retained in the reconstruction. Whereas Figure 6a shows results from regression with the canonical RSN maps used in the dual regression, Figure 6b shows correlations against the z-statistic maps output from a dual regression analysis of the ground truth (subject-specific) maps. Under this metric, means ($\pm$ standard deviation) of the spatial correlations were: k-t SLOW 0.64 ± 0.06, k-t PSF 0.35 ± 0.10, k-t INTERP 0.57 ± 0.10 and k-t FASTER

0.73 ± 0.06. These results also demonstrate a highly significant mean difference for k-t FASTER compared with the other methods (paired t-tests, $p < 0.05$ corrected), indicating that k-t FASTER reconstructions are superior at retaining the dataset-specific detail.

One feature of these results is that low variance components produced lower correlations than high variance components in the k-t FASTER reconstruction (not shown), indicating that reconstruction performance is dependent on component strength. Supporting Figures S5 and S6 provide more detailed analysis of the average k-t FASTER dual regression results in context of varying amounts of rank information.

Figures 7 and 8 show example z-statistic maps from representative datasets for each method, alongside the ground truth maps, for an occipital visual RSN and bilateral parietal RSN, respectively. In each of the examples, characteristics common to each reconstruction method can be observed. The k-t SLOW maps resemble sparser versions of the truth maps, lacking the temporal DOF necessary to capture all the detail in each map. The k-t PSF data show relatively poor recovery of RSNs, while both k-t INTERP and k-t FASTER produced maps with good agreement with the ground-truth data. From both quantitative and qualitative evaluations of the data,
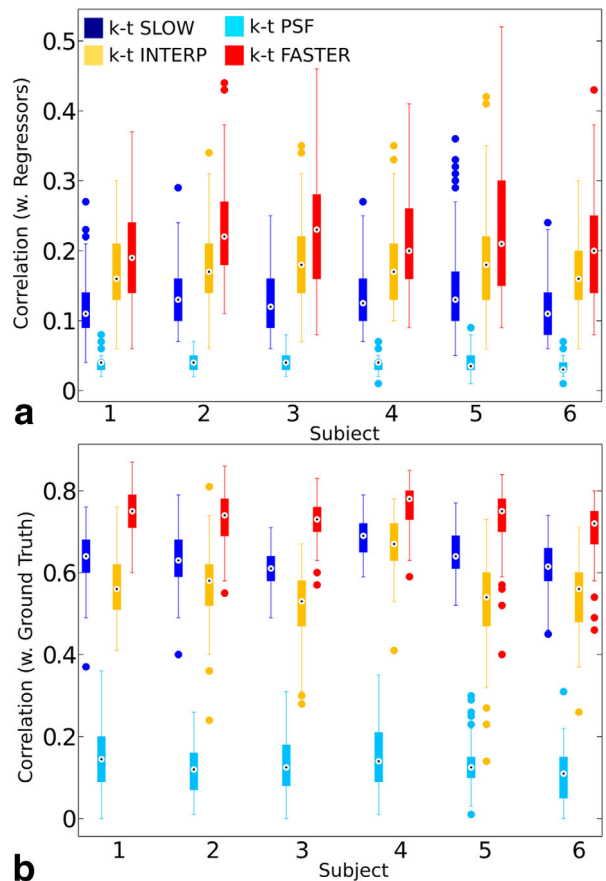

a

b

FIG. 6. **a**: Mean z-statistic map correlations for the reconstructed data compared to the canonical input regressors. **b:** Direct correlation of the z-statistic maps for each method with the truth maps, which show that the k-t FASTER maps match the truth maps best. In both cases, error bars denote standard deviation.
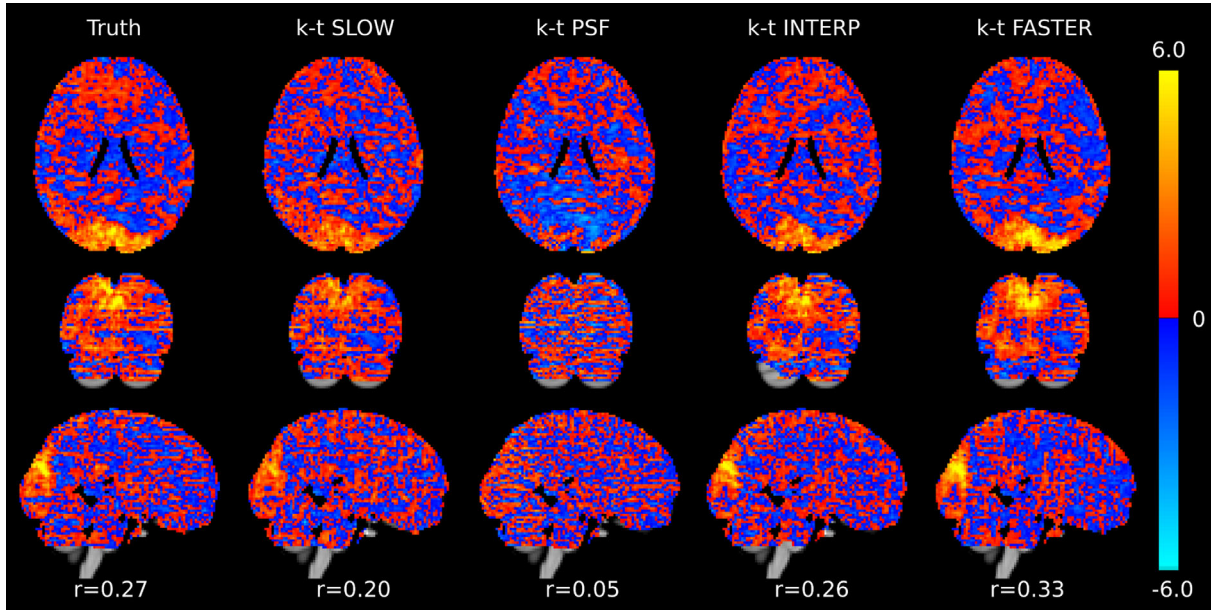
FIG. 7. Z-stat image dual regression output for a single dataset for a regressor corresponding to an occipital (visual) resting state network. All images are un-thresholded with color scale mapped between $0 < |z| < 6$.

however, the k-t FASTER method reproduces RSNs with the highest fidelity.

Prospective Under-Sampling Results

The prospectively under-sampled data were acquired such that the fully sampled k-t SLOW data had only 72 time points, whereas the k-t FASTER data produced 304 time points in the same duration. To accommodate the reduced temporal DOF in the k-t SLOW data, only 64 regressors were used for the dual regression analysis. As an additional test, the set of regressors was further reduced to 32, to further examine the effect of the DOF on these analyses. Figure 9 shows the correlation coefficients for the first 32 regressors (ordered by variance), by performing 2 separate dual regression analyses. The first dual regression used 32 regressors only, while the second dual regression used 64 regressors, including the same 32 used in the first analysis. There is a striking increase in correlation in k-t SLOW when halving the number of regressors, indicating insufficient DOF when using 64 regressors in a dataset with 72 time points. In
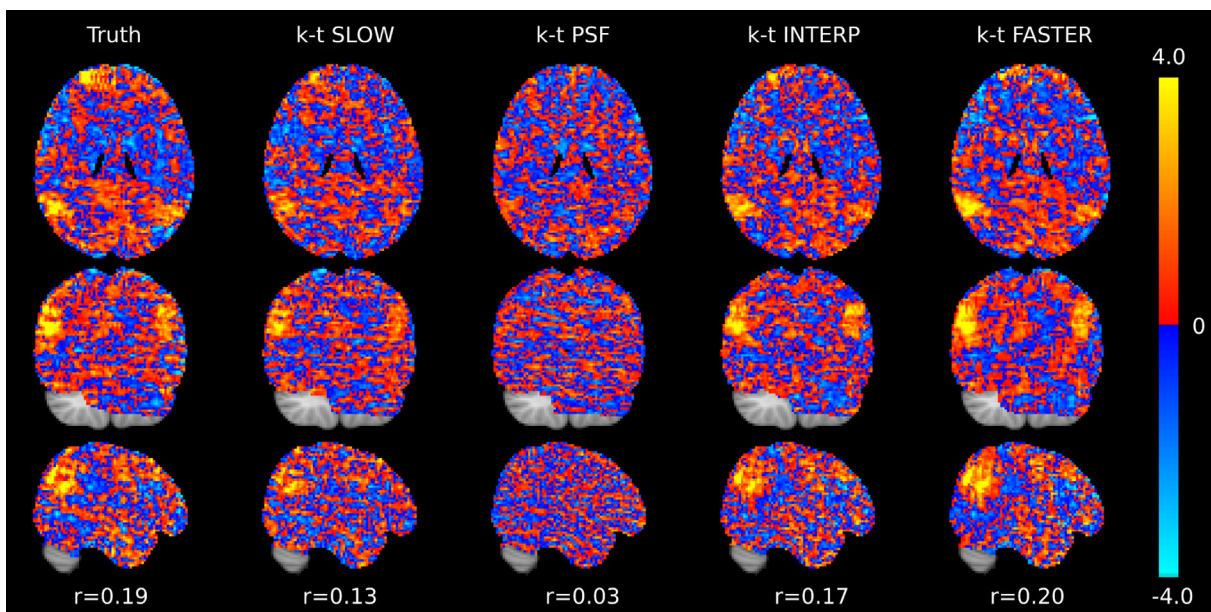

FIG. 8. Z-stat image dual regression output for a single dataset for a regressor corresponding bilateral parietal resting state network. All images are un-thresholded with color scale mapped between $0 < |z| < 4$.
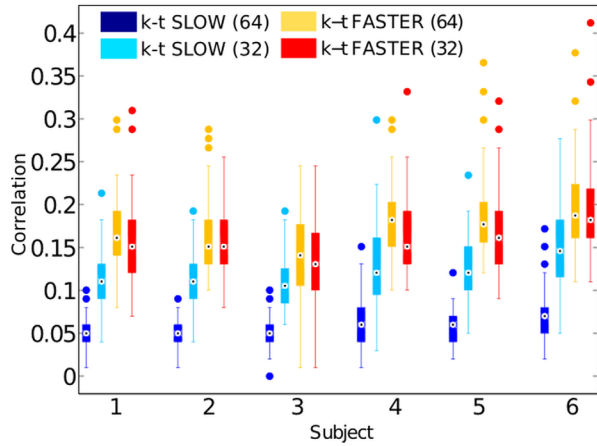
FIG. 9. Boxplot of correlation coefficients from the dual regression averaged across scans and 32 maps for each of the six subjects, using both 32 and 64 regressors.

comparison, k-t FASTER has consistently higher correlations than k-t SLOW, with essentially no change with increased number of regressors. These results demonstrate that, despite identical total acquisition durations for the two approaches, the k-t FASTER data contain greater temporal DOF than the k-t SLOW acquisition. This provides further evidence that k-t FASTER is not performing a trivial operation (e.g., interpolation) with the sampled points, but is actually identifying meaningful spatial-temporal structure in the data. Again, relatively low correlations are expected in individual subject maps.

Finally, Figure 10 shows an example z-statistic map corresponding to a left somatosensory RSN in a representative dataset, compared with the corresponding spatial regressor map. These maps were chosen to be representative of the mean correlations for each case. The k-t FASTER maps clearly show a greater extent of high (and spatially well localised) z-statistics, most evident in the axial and sagittal views, and are a good reflection of the aggregate correlation differences between the k-t SLOW and k-t FASTER data expressed in the boxplots of Figure 9.

DISCUSSION

Following the recent success of CS-based acceleration strategies, several new techniques and applications have emerged in MR exploiting the properties of sparsity or matrix rank for image reconstruction. Functional neuroimaging, however, has not seen wide adoption or exploration of any of these methods, and fMRI acceleration techniques have instead focused primarily on multicoil parallel imaging methods. k-t FASTER takes a different approach by exploiting the spatio-temporal structure of fMRI data, following similar strategies previously applied to dynamic contrast enhanced and cardiac cine imaging (29,30). While application to fMRI has been mentioned previously (29), and recently presented for task-fMRI (31), the k-t FASTER approach (32) makes the explicit link between low-rank acceleration strategies and resting state fMRI data models.

We were able to recover under-sampled image time-series, resulting in a higher temporal sampling rate than can be achieved with full sampling. The low-rank
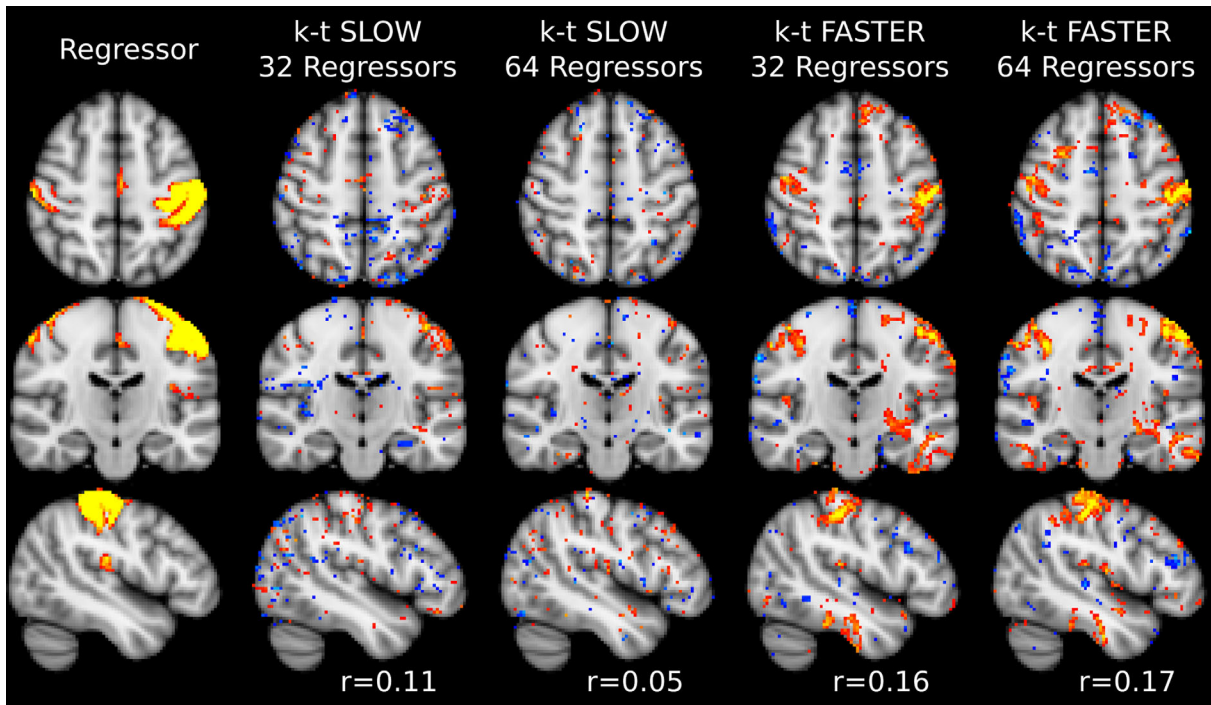

FIG. 10. Z-stat image dual regression output for a single subject for a regressor corresponding to left somatosensory resting state network. Correlation coefficients with the canonical regressors are shown at the bottom of the images. All maps are thresholded at |z|>2 and with color scale mapped between 2 < |z| < 6, except for the regressor maps, which are scaled and thresholded at 10×.

structure allows k-t FASTER to generate the missing points in an intelligent way, by preferentially recovering high-variance (i.e., signal) components using an iterative, SVD-based algorithm. Furthermore, the 3D segmented EPI sampling used here does not suffer from the jittery TE or irregular off-resonance artifacts that can be present with randomly undersampling phase-encoding lines of a traditional 2D EPI sequence (21).

Evaluation of a range of image reconstruction methods with similar sampling densities found the lowest Frobenius matrix norm errors using k-t FASTER. The k-space time-series reconstructed by k-t FASTER was further observed to capture true temporal features that would be missed by methods that impose temporal smoothing, such as keyhole or interpolation methods. More importantly, the dual regression analysis found the highest RSN reproduction fidelity for the k-t FASTER method, in both retrospectively and prospectively sampled evaluations. It should be noted that in this evaluation, the k-t FASTER results benefitted from the final data replacement step, which re-introduces higher rank information to the image estimate (beyond the rank constraint of 128). In contrast, the k-t PSF reconstructions at the comparatively low rank of 26 performed poorly in the dual regression analysis, which likely reflects the high complexity of the fMRI datasets studied.

The differences between the k-t SLOW and k-t FASTER results in the prospective sampling experiment indicate that more temporal DOF are available to the k-t FASTER reconstruction, despite having identical sampling durations. Moreover, information about higher temporal frequencies is retained using k-t FASTER. Even in the presence of the smooth hemodynamic response, this signal contains meaningful information about neural fluctuations (52). These features support the use of k-t FASTER as a method for generating fMRI datasets with higher temporal DOF than standard sampling approaches for use in applications like temporal ICA (5).

In the prospectively sampled data, reconstruction fidelity was evaluated using only correlations with group-averaged RSN spatial maps. With this metric, it is difficult to ensure that subject-specific detail is retained in the k-t FASTER reconstructions. While the retrospective sampling comparison may not have fully captured artifacts and phase effects that are present in real sampling conditions, it provided a validation where ground truth data were available, indicating that subject-specific spatial features are largely preserved. Also, while the SNR of the 3D sampling strategy used here (both retrospectively simulated and prospectively acquired) is independent of volume TR, it is important to note that in a standard 2D multislice acquisition strategy, k-t SLOW data would be expected to have increased SNR due to reduced saturation effects from longer TRs.

Although the low-rank reconstruction is largely independent of TR, consideration of SNR and physiological noise can be important. Imaging in low SNR conditions may not be feasible, because the low-rank signal components need to be distinguishable from additive thermal noise. Generally speaking, however, physiological noise can be viewed as additional signal components, given that it exhibits coherent signal fluctuations across space and time (i.e., not thermal noise). Large amounts of physiological noise can therefore increase the effective rank of data matrix, which would have to be taken into account in the reconstruction. Furthermore, the effects of motion and field inhomogeneity on the image phase and reconstruction fidelity should not be ignored, as the reconstruction is performed using fully complex data. Modeling of these effects and incorporating phase constraints in the reconstruction process is a topic of future consideration.

The achievable under-sampling factors in MC are intrinsically linked to the underlying rank of the matrix to be recovered. There would be a clear penalty in reconstruction fidelity associated with under-sampling beyond limits supported by the data, and improving acceleration factors in k-t FASTER will likely require the use of additional constraints, information (e.g., coil sensitivities) or longer sampling durations. One advantage of exploring coil-independent acceleration methods is that overall acceleration factors can be increased by the combination of both methods. In fact, the prospectively sampled data presented here already included a $2\times$ in-plane acceleration in addition to the $4.27\times$ acceleration in the partition encoding direction recovered using MC, demonstrating the compatibility of k-t FASTER acceleration with standard parallel imaging.

In this initial work, we used a hard-thresholding approach to MC, which required a fixed rank input as a hard thresholding parameter. This method was found to be more robust in the fMRI datasets studied than soft-thresholding approaches. One drawback of the k-t FASTER iterative reconstruction is that it is significantly slower than the noniterative approaches of k-t INTERP and k-t PSF, requiring hours of computation time for large k-t matrices. Future development of the IHT+MS algorithm used in k-t FASTER will investigate selection criteria for the hard rank thresholds, and optimization for speed and computational efficiency using parallel computation techniques and GPU-acceleration. Also, exploration of joint coil sensitivity constraints to improve reconstruction fidelity and achievable acceleration factors is a topic of future interest, and has already been explored in alternative dynamic low-rank imaging methods (53–55). Furthermore, many rank-constrained imaging methods are now incorporating joint sparsity constraints (31,44,56,57), or decomposition of sparse and low-rank components to an image time-series (54). These features can be integrated into k-t FASTER in future refinements of the technique using penalized matrix decomposition methods (58) or the robust PCA formalism (59). Finally, extension of the approach to include non-Cartesian sampling trajectories that take advantage of more distributed k-space coverage is easily facilitated through use of the NUFFT (60) instead of the FFT in the measurement operator.

ACKNOWLEDGMENTS

"multiband" approach (accelerated by $\times 8$, giving us $2 \times 2 \times 2$ mm data with $TR = 0.836$ s). We are very grateful to Kamil Ugurbil, Essa Yacoub, Steen Moeller and Eddie Auerbach for providing the multiband software. The authors would also like to thank Lynn Ossher for help with editing of the manuscript. This work is supported by the EPSRC (MC, TB) and the Wellcome Trust (S.S., K.M., P.K.).

APPENDIX A1: IHT+MS RECONSTRUCTION ALGORITHM

- Let the matrix Y contain the sampled data, M be a sampling matrix with 1s in sampled positions, and 0s in unsampled positions, and X^i be the i^{th} matrix estimate, where all matrices are m × n
- The · operator denotes element-wise multiplication
- Choose a rank r, step size $0 < \mu \leq 1$, and soft-thresholding parameter $0 \leq c \leq 1$
- Initialise matrix estimate X^0 to the zero matrix with appropriate dimensions (or some alternative initial guess)
- While $\|X^i - X^{i-1}\|_F > min\ update\ criterion$ or $i < max\ iterations$
  $Z = M \cdot (Y - X^i)$
  $U\Sigma V^* = SVD(X^i + \mu \cdot Z)$
  For $j = 1 : n$
    If $j \leq r$: $\tilde{\Sigma}(j, j) = \Sigma(j, j) - c \cdot \sigma_{r+1}$
    Else : $\tilde{\Sigma}(j, j) = 0$
  $X^{i+1} = U\tilde{\Sigma}V^*$

REFERENCES

1. Feinberg DA, Moeller S, Smith SM, Auerbach E, Ramanna S, Gunther M, Glasser MF, Miller KL, Ugurbil K, Yacoub E. Multiplexed echo planar imaging for sub-second whole brain FMRI and fast diffusion imaging. PLoS One 2010;5:e15710.
2. Frank LR, Buxton RB, Wong EC. Estimation of respiration-induced noise fluctuations from undersampled multislice fMRI data. Magn Reson Med 2001;45:635–644.
3. Biswal B, Yetkin FZ, Haughton VM, Hyde JS. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn Reson Med 1995;34:537–541.
4. Beckmann CF, DeLuca M, Devlin JT, Smith SM. Investigations into resting-state connectivity using independent component analysis. Philos Trans R Soc Lond B Biol Sci 2005;360:1001–1013.
5. Smith SM, Miller KL, Moeller S, et al. Temporally-independent functional modes of spontaneous brain activity. Proc Natl Acad Sci U S A 2012;109:3131–3136.
6. Lee H-L, Zahneisen B, Hugger T, LeVan P, Hennig J. Tracking dynamic resting-state networks at higher frequencies using MR-encephalography. Neuroimage 2013;65:216–222.
7. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: sensitivity encoding for fast MRI. Magn Reson Med 1999;42:952–962.
8. Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, Haase A. Generalized autocalibrating partially parallel acquisitions (GRAPPA). Magn Reson Med 2002;47:1202–1210.
9. Poser BA, Koopmans PJ, Witzel T, Wald LL, Barth M. Three dimensional echo-planar imaging at 7 Tesla. Neuroimage 2010;51:261–266.
10. Lin F-H, Wald LL, Ahlfors SP, Hämäläinen MS, Kwong KK, Belliveau JW. Dynamic magnetic resonance inverse imaging of human brain function. Magn Reson Med 2006;56:787–802.
11. Hennig J, Zhong K, Speck O. MR-encephalography: fast multi-channel monitoring of brain physiology with magnetic resonance. Neuroimage 2007;34:212–219.
12. Moeller S, Yacoub E, Olman CA, Auerbach E, Strupp J, Harel N, Ugurbil K. Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial

13. and temporal whole-brain fMRI. Magn Reson Med 2010;63:1144–1153.
13. Setsompop K, Gagoski BA, Polimeni JR, Witzel T, Wedeen VJ, Wald LL. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. Magn Reson Med 2012;67:1210–1224.
14. Madore B, Glover GH, Pelc NJ. Unaliasing by Fourier-encoding the overlaps using the temporal dimension (UNFOLD), applied to cardiac imaging and fMRI. Magn Reson Med 1999;42:813–828.
15. Tsao J, Boesiger P, Pruessmann KP. k-t BLAST and k-t SENSE: dynamic MRI with high frame rate exploiting spatiotemporal correlations. Magn Reson Med 2003;50:1031–1042.
16. Zarahn E, Aguirre GK, D'Esposito M. Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. Neuroimage 1997;5:179–197.
17. Lustig M, Donoho D, Pauly JM. Sparse MRI: the application of compressed sensing for rapid MR imaging. Magn Reson Med 2007;58:1182–1195.
18. Gamper U, Boesiger P, Kozerke S. Compressed sensing in dynamic MRI. Magn Reson Med 2008;59:365–373.
19. Jung H, Ye JC, Kim EY. Improved k-t BLAST and k-t SENSE using FOCUSS. Phys Med Biol 2007;52:3201–3226.
20. Jung H, Ye JC. Performance evaluation of accelerated functional MRI acquisition using compressed sensing. In Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI), Boston, Massachusetts, USA, 2009. p 702–705.
21. Zong X, Lee J, John Poplawsky A, Kim S-G, Ye JC. Compressed sensing fMRI using gradient-recalled echo and EPI sequences. Neuroimage 2014;92C:312–321.
22. Jeromin O, Pattichis MS, Calhoun VD. Optimal compressed sensing reconstructions of fMRI using 2D deterministic and stochastic sampling geometries. Biomed Eng Online 2012;11:25.
23. Holland DJ, Liu C, Song X, Mazerolle EL, Stevens MT, Sederman AJ, Gladden LF, D'Arcy RCN, Bowen CV, Beyea SD. Compressed sensing reconstruction improves sensitivity of variable density spiral fMRI. Magn Reson Med 2013;70:1634–1643.
24. Liang Z-P. Spatiotemporal imaging with partially separable functions. In Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI), Arlington, Virginia, USA, 2007. p 988–991.
25. Pedersen H, Kozerke S, Ringgaard S, Nehrke K, Kim WY. k-t PCA: temporally constrained k-t BLAST reconstruction using principal component analysis. Magn Reson Med 2009;62:706–716.
26. Candes EJ, Recht B. Exact matrix completion via convex optimization. Found Comput Math 2009;9:717–772.
27. Zhao B, Haldar JP, Brinegar C, Liang Z-P. Low rank matrix recovery for real-time cardiac MRI. In Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI), Rotterdam, 2010. p 996–999.
28. Candes EJ, Plan Y. Matrix completion with noise. Proc IEEE 2010;98:925–936.
29. Haldar JP, Liang Z-P. Spatiotemporal imaging with partially separable functions: a matrix recovery approach. In Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI), Rotterdam, 2010. p 716–719.
30. Haldar JP, Liang Z-P. Low-rank approximations for dynamic imaging. In Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI), Chicago, Illinois, USA, 2011. p 1052–1055.
31. Lam F, Zhao B, Liu Y, Liang Z-P, Weiner M, Schuff N. Accelerated fMRI using low-rank model and sparsity constraints. In Proceedings of the 21st Annual Meeting of ISMRM, 2013, Salt Lake City, Utah, USA. Abstract 2620.
32. Chiew M, Smith SM, Koopmans PJ, Blumensath T, Miller KL. k-t FASTER: a new method for the acceleration of resting state FMRI Data Acquisition. In Proceedings of the 21st Annual Meeting of ISMRM, Salt Lake City, Utah, USA, 2013. Abstract 1752.
33. Hyvärinen A. Fast and robust fixed-point algorithms for independent component analysis. IEEE Trans Neural Netw 1999;10:626–634.
34. Muckley MJ, Peltier SJ, Noll DC, Fessler JA. Model-based reconstruction for physiological noise correction in functional MRI. In Proceedings of the 21st Annual Meeting of ISMRM, Salt Lake City, Utah, USA, 2013. Abstract 2623.
35. Beckmann CF, Smith SM. Probabilistic independent component analysis for functional magnetic resonance imaging. IEEE Trans Med Imaging 2004;23:137–152.

36. Smith SM, Fox PT, Miller KL, et al. Correspondence of the brain's functional architecture during activation and rest. Proc Natl Acad Sci U S A 2009;106:13040–13045.

37. Candes EJ, Tao T. The power of convex relaxation: near-optimal matrix completion. IEEE Trans Inf Theory 2010;56:2053–2080.

38. Ma S, Goldfarb D, Chen L. Fixed point and Bregman iterative methods for matrix rank minimization. Math Program 2011;128:321–353.

39. Smith SM, Beckmann CF, Andersson J, et al. Resting-state fMRI in the Human Connectome Project. Neuroimage 2013;80:144–168.

40. Blumensath T, Davies ME. Iterative hard thresholding for compressed sensing. Appl Comput Harmon Anal 2009;27:265–274.

41. Blumensath T. Sampling and reconstructing signals from a union of linear subspaces. IEEE Trans Inf Theory 2011;57:4660–4671.

42. Recht B, Fazel M, Parrilo PA. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. Siam Rev 2010;52:471–501.

43. Goldfarb D, Ma S. Convergence of fixed-point continuation algorithms for matrix rank minimization. Found Comput Math 2011;11:183–210.

44. Lingala SG, Hu Y, DiBella E, Jacob M. Accelerated dynamic MRI exploiting sparsity and low-rank structure: k-t SLR. IEEE Trans Med Imaging 2011;30:1042–1054.

45. Van Vaals JJ, Brummer ME, Dixon WT, Tuithof HH, Engels H, Nelson RC, Gerety BM, Chezmar JL, Boer Den JA. "Keyhole" method for accelerating imaging of contrast agent uptake. J Magn Reson Imaging 1993;3:671–675.

46. Yun SD, Reske M, Vahedipour K, Warbrick T, Shah NJ. Parallel imaging acceleration of EPIK for reduced image distortions in fMRI. Neuroimage 2013;73:135–143.

47. Jenkinson M, Bannister P, Brady M, Smith S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. Neuroimage 2002;17:825–841.

48. Beckmann CF, Mackay CE, Filippini N, Smith SM. Group comparison of resting-state FMRI data using multi-subject ICA and dual regression. OHBM, 2009. doi: 10.1016/S1053-8119(09)71511-3.

49. Filippini N, Macintosh BJ, Hough MG, Goodwin GM, Frisoni GB, Smith SM, Matthews PM, Beckmann CF, Mackay CE. Distinct patterns of brain activity in young carriers of the APOE-epsilon4 allele. Proc Natl Acad Sci U S A 2009;106:7209–7214.

50. Van Essen DC, Smith SM, Barch DM, Behrens TEJ, Yacoub E, Ugurbil K, WU-Minn HCP Consortium. The WU-Minn Human Connectome Project: an overview. Neuroimage 2013;80:62–79.

51. Ugurbil K. Development of functional imaging in the human brain (fMRI); the University of Minnesota experience. Neuroimage 2012;62:613–619.

52. Niazy RK, Xie J, Miller K, Beckmann CF, Smith SM. Spectral characteristics of resting state networks. Prog Brain Res 2011;193:259–276.

53. Lingala SG, DiBella E, Adluru G, McGann C, Jacob M. Accelerating free breathing myocardial perfusion MRI using multi coil radial k-t SLR. Phys Med Biol 2013;58:7309–7327.

54. Otazo R, Candes E, Sodickson DK. Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components. Magn Reson Med 2015;73:1125–1136.

55. Trzasko JD, Manduca A. A unified tensor regression framework for calibrationless dynamic, multi-channel MRI reconstruction. In Proceedings of the 21st Annual Meeting of ISMRM, Salt Lake City, Utah, USA, 2013. Abstract 0603.

56. Zhao B, Haldar JP, Christodoulou AG, Liang ZP. Image reconstruction from highly undersampled (k, t)-space data with joint partial separability and sparsity constraints. IEEE Trans Med Imaging 2012;31:1809–1820.

57. Majumdar A. Improved dynamic MRI reconstruction by exploiting sparsity and rank-deficiency. Magn Reson Imaging 2013;31:789–795.

58. Witten DM, Tibshirani R, Hastie T. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. Biostatistics 2009;10:515–534.

59. Candés EJ, Li X, Ma Y, Wright J. Robust principal component analysis? J ACM 2011;58: Article 11.

60. Fessler JA, Sutton BP. Nonuniform fast Fourier transforms using min-max interpolation. IEEE Trans Signal Process 2003;51:560–574.