Magnetic Resonance in Medicine 50:1031–1042 (2003)

k-t BLAST and *k-t* SENSE: Dynamic MRI With High Frame Rate Exploiting Spatiotemporal Correlations

Jeffrey Tsao,* Peter Boesiger, and Klaas P. Pruessmann

Dynamic images of natural objects exhibit significant correlations in *k*-space and time. Thus, it is feasible to acquire only a reduced amount of data and recover the missing portion afterwards. This leads to an improved temporal resolution, or an improved spatial resolution for a given amount of acquisition. Based on this approach, two methods were developed to significantly improve the performance of dynamic imaging, named *k-t* BLAST (Broad-use Linear Acquisition Speed-up Technique) and *k-t* SENSE (SENSitivity Encoding) for use with a single or multiple receiver coils, respectively. Signal correlations were learned from a small set of training data and the missing data were recovered using all available information in a consistent and integral manner. The general theory of *k-t* BLAST and *k-t* SENSE is applicable to arbitrary *k*-space trajectories, time-varying coil sensitivities, and under- and overdetermined reconstruction problems. Examples from ungated cardiac imaging demonstrate a 4-fold acceleration (voxel size 2.42 × 2.52 mm^2, 38.4 fps) with either one or six receiver coils. *k-t* BLAST and *k-t* SENSE are applicable to many areas, especially those exhibiting quasiperiodic motion, such as imaging of the heart, the lungs, the abdomen, and the brain under periodic stimulation. Magn Reson Med 50:1031–1042, 2003. © 2003 Wiley-Liss, Inc.

Key words: MRI; fast dynamic imaging; *k-t* BLAST; *k-t* SENSE; prior-information-driven parallel imaging

Dynamic MRI captures an object in motion by acquiring a series of images at a high frame rate. Conceptually, the straightforward approach would be to acquire the full data for reconstructing each time frame separately. This requires the acquisition of each time frame to be short relative to the object motion in order to effectively obtain an instantaneous snapshot. However, this approach is limited by physical (e.g., gradient strength and slew rate) and physiological (e.g., nerve stimulation) constraints on the speed of data acquisition. Over the years a number of strategies have been proposed to further increase the acquisition rate by reducing the amount of acquired data by a given factor, referred to as the acceleration factor hereafter. These strategies are able to reduce data acquisition without compromising image quality significantly because typical image series exhibit a high degree of spatiotemporal correlations, either by nature or by design. Therefore, there is a certain amount of redundancy within the data. In general, such strategies for reducing data acquisition can be divided into three approaches, based on exploiting correlations in *k*-space, in time, or in both *k*-space and time.

The first approach, based on exploiting correlations in *k*-space, encompasses a wide variety of methods, including partial Fourier methods (1,2), reduced-field-of-view methods (3), parallel imaging (4,5), and prior-information-driven methods (1,6). These methods speed up acquisition by collecting only a fraction of *k*-space at each time frame. The missing data are then recovered based on the measured *k*-space points from the same time frame. Although the methods differ significantly in their technique of data recovery, they are all based on the principle that each point in *k*-space contains some information about other points in *k*-space. This correlation can be used to recover the missing information. In all of these methods, the image at each time frame is reconstructed independently from the images at other time frames.

The second approach, based on exploiting temporal correlations, encompasses methods such as keyhole (7,8) and various view-sharing strategies (9–13). In these methods the full *k*-space is successively updated, but the rate of update may vary for different portions of *k*-space. The measured data at a given position in *k*-space form a series of samples in time. The missing data at all other time points can then be interpolated or extrapolated from these samples, using a variety of schemes, including zero-order hold, nearest-neighbor interpolation, and linear interpolation, among others. More sophisticated extensions of these methods rely on interpolators with a frequency response that is tailored to the temporal-frequency contents of the object, e.g., the method of Xiang and Henkelman (14), Dynamic Imaging by Model Estimation (DIME) (15), or the fMRI example of UNaliasing by Fourier-encoding the Overlaps using the temporaL Dimension (UNFOLD) (16,17). Regardless of the choice of the interpolator, each point in *k*-space is reconstructed separately from all other points.

The third approach, based on exploiting correlations in both *k*-space and time (16,18–22), is a combination of the two above approaches. Thus, a missing data point is estimated based on other available points, typically within its vicinity in both *k*-space and time. The advantage of this approach is that it exploits more of the relevant correlations, thus improving the estimation of missing data. This improvement may translate to either reduced estimation error or higher achievable reduction in data acquisition. Examples of this approach are the cardiac example of UNFOLD (16) and a related method developed independently by Willis and Bresler (18–21). Both of these methods (16,18–21) share the interesting property that the reconstruction only involves a filtering process, which does

Institute for Biomedical Engineering, University of Zurich, and Swiss Federal Institute of Technology (ETH) Zurich, Zurich, Switzerland.

Grant sponsor: EUREKA; Grant number: EU1353; Grant sponsor: KTI; Grant number: 3030.2; Grant sponsors: Philips Medical Systems, Best, the Netherlands; Canadian Institutes of Health Research (postdoctoral fellowship to J.T.).

*Correspondence to: Jeffrey Tsao, Ph.D., Institute for Biomedical Engineering, ETH Zurich, Building ETF, Room C108, Sternwartstrasse 7, 8092 Zurich, Switzerland. E-mail: jtsao2@hotmail.com

Received 10 December 2002; revised 3 July 2003; accepted 7 July 2003.

DOI 10.1002/mrm.10611

Published online in Wiley InterScience (www.interscience.wiley.com).

not amplify noise. However, to achieve this desirable property the data have to be acquired according to a sampling pattern that solves a related Euclidean packing problem (17–21). This requirement places a stringent restriction on the achievable acceleration factor. In practice, a 2-fold acceleration is typically achievable (16,17), while a 4-fold acceleration is possible only under restrictive conditions (see fig. 14 in Ref. 18).

In the present work, two new methods that exploit correlations in both *k*-space and time are proposed. They allow the use of higher acceleration factors by relaxing the sampling restriction of the previous methods. Higher acceleration is achieved at the expense of some noise amplification during reconstruction. Nevertheless, this trade-off between the achievable acceleration and image quality is gentle and nonabrupt. Also, it is adjusted to match each imaging situation automatically, since the methods learn the actual spatiotemporal distribution of the object signals, rather than make a priori assumptions about it (17). Improved accuracy is achieved if multiple receiver coils are available to provide additional data for reconstruction. The proposed methods are named *k-t* BLAST (Broad-use Linear Acquisition Speed-up Technique (6)) and *k-t* SENSE (SENSitivity Encoding (5)) for use with a single or multiple receiver coils, respectively.

THEORY

Dynamic MRI acquires the raw data in *k*-space at different time points, *t*. Using the formalism of Xiang and Henkelman (23), the raw data can be equivalently viewed as being acquired in a higher dimensional *k-t* space. The arrangement of these discrete samples in *k-t* space is referred to as the *k-t* sampling pattern.

Reconstruction of a dynamic image series involves determining the object signals in *k-t* space from the discretely sampled data. According to the properties of Fourier transformation, discrete sampling in *k-t* space leads to a convolution of the object signals in the reciprocal *x-f* space with a point spread function. "*x*" and "*f*" are the conjugate axes of "*k*" and "*t*" after inverse Fourier transform. Here, they refer to spatial axis (along the phase-encoding direction) and temporal frequency, respectively.

To simplify the subsequent description, we focus on 2D dynamic imaging only, although the proposed methods are applicable to any number of spatial dimensions. We divide the remainder of this section into three parts. The first part considers the situation where the *k-t* sampling pattern conforms to a lattice structure (i.e., a grid that may be sheared or rotated) (20,21), such as by using interleaved echo-planar imaging (EPI) (24). If multiple receiver coils are used, it is assumed in this part that the coil sensitivity information is time-invariant. These two conditions lead to tremendous simplifications in the reconstruction, allowing for efficient computation. The second part describes the general situation where the *k-t* sampling pattern may not conform to a lattice structure and the coil sensitivity may be time-varying. Finally, the third part describes how to obtain the signal estimates needed for the reconstruction described in the first two parts.

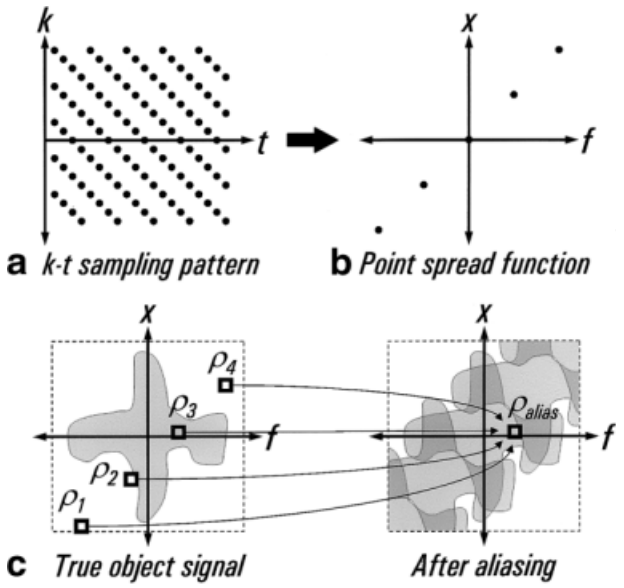

FIG. 1. **a:** *k-t* space sampling pattern with a 4-fold acceleration. "*k*" and "*t*" refer to the phase-encode index and time, respectively. The sampling pattern is equivalent to sampling on a sheared grid. **b:** Resulting point spread function in *x-f* space. "*x*" and "*f*" refer to the spatial position along the phase-encoding direction and temporal frequency, respectively. **c:** Convolution of object signals in *x-f* space with point spread function, resulting in aliasing which maps ρ_1, ρ_2, ρ_3, and ρ_4 onto a single aliased voxel ρ_{alias}. To avoid clutter, not all the signal replicates are shown.

Lattice Sampling in *k-t* Space

Figure 1a shows a *k-t* sampling pattern with a lattice structure, which can be viewed as a sheared grid, or in this special case, a rotated grid. At each time point, *k*-space is sampled in a Cartesian manner. The frequency-encoding direction is omitted for simplicity. It is oriented perpendicular to the page and it is reconstructed by inverse Fourier transform. "*k*" refers to the index of the phase-encode line. At each time point *t*, several regularly spaced phase-encode lines are acquired. Different sets of phase-encode lines are acquired at successive time points. This sampling pattern is referred to as "sparse," since one or more phase-encode lines are skipped between the sampled ones. This sampling pattern is, in reality, slightly skewed, since the phase-encode lines of each time point are not acquired instantaneously. Nevertheless, this skew is relatively small, so it can be neglected in practice.

As a result of the lattice structure of the *k-t* sampling pattern, the point spread function in *x-f* space also has a lattice structure (Fig. 1b). Since the sampling in *k-t* space is finite, the point spread function is strictly speaking a superposition of shifted sinc functions. In high-resolution imaging, the sinc functions have narrow widths, so the point spread function is adequately approximated as a set of delta functions arranged in a lattice, which simplifies the subsequent treatment significantly. Convolution with such a point spread function leads to periodic replication of the object signals (17). Figure 1c illustrates the effects on the object signals in *x-f* space. The convolution maps the signals from several locations (e.g., from *x-f* space voxels

labeled ρ_1, ρ_2, ρ_3, and ρ_4) onto the same location ρ_{alias}, resulting in aliasing. It should be noted that the aliasing occurs among voxels at different spatial positions x as well as different temporal frequencies f. Mathematically, this aliasing is described as:

$$\begin{bmatrix} 1 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} \rho_1 \\ \rho_2 \\ \rho_3 \\ \rho_4 \end{bmatrix} = \mathbf{1}\boldsymbol{\rho} = \rho_{\text{alias}}. \quad [1]$$

Eq. [1] is underdetermined. Thus, there are an infinite number of feasible solutions that satisfy Eq. [1]. The task of reconstruction is to find an appropriate solution out of the solution set, which can be handled in a number of ways. For example, in a typical time series of images most of the signal energy is contained in the lower temporal frequencies. Therefore, one solution is to assign all of the signal in ρ_{alias} to the voxel at the lowest temporal frequency (i.e., $\rho_2 = \rho_{alias}$ in this example since ρ_2 is closest to the vertical axis at $f = 0$), while all other voxels are set to zero (i.e., $\rho_1 = \rho_3 = \rho_4 = 0$). This strategy is sound if the Nyquist sampling criterion is fulfilled. Otherwise, it can lead to incorrect signal assignment, as in the present example.

A second solution is to make use of prior information, which can be obtained either from a priori assumptions (16,17) or measured explicitly, as described later. If it is known that only one of the voxels contains a significant amount of signal (i.e., ρ_3 in this example), then all of the signal in ρ_{alias} can be attributed to it, while the remaining voxels can be set to zero (i.e., $\rho_1 = \rho_2 = \rho_4 = 0$). In other words, the aliasing is resolved through an n-way decision that selects one appropriate voxel out of n voxels involved in the aliasing (16,18–21). The advantage of this strategy is that it does not lead to any error amplification, since it only involves assigning signals to the appropriate voxels. However, it is unable to resolve the aliasing if more than one voxel is known to contain signal. One way to circumvent this limitation is to alter the *k-t* sampling pattern in order to further separate the peaks of the point spread function and to reduce the amount of aliasing. However, this may lead to lower achievable acceleration as well. If the aliasing is confined to voxels with low signals, one may also circumvent the limitation by setting all the voxels involved to zero (18,19) (i.e., $\rho_1 = \rho_2 = \rho_3 = \rho_4 = 0$). This leads to some filtering of the object signals, but avoids any aliasing.

The purpose of this work is to present a more flexible solution to this aliasing problem by making use of additional information that can be obtained easily in practice. Suppose that an estimate of the signal magnitudes (m_1, m_2, m_3, and m_4) is available for the voxels involved (ρ_1, ρ_2, ρ_3, and ρ_4), one can obtain a feasible solution to Eq. [1] that minimizes the following weighted norm:

$$\sqrt{\sum_i |\rho_i / m_i|^2}. \quad [2]$$

This solution is desirable since it selects the feasible solution that matches both the data as well as the expected relative signal magnitudes. If the magnitude estimates are unavailable (i.e., $m_1 = m_2 = m_3 = m_4 =$ arbitrary value), this solution reduces exactly to the minimum-norm solution. The weighted-minimum-norm solution to Eq. [1] is (25,26):

$$\boldsymbol{\rho} = \mathbf{M}(\mathbf{1M})^{+}\rho_{alias} \quad [3]$$

where $\mathbf{M}$ is a diagonal matrix with m_1 to m_4 along the diagonal. The superscript + indicates a Moore-Penrose pseudoinverse. Eq. [3] can be written as:

$$\boldsymbol{\rho} = \mathbf{M}^2\mathbf{1}^H(\mathbf{1M}^2\mathbf{1}^H)^{-1}\rho_{alias} = \frac{[|m_1|^2|m_2|^2|m_3|^2|m_4|^2]^H}{|m_1|^2 + |m_2|^2 + |m_3|^2 + |m_4|^2}\,\rho_{alias}. \quad [4]$$

It can be seen from Eq. [4] that the weighted-minimum-norm solution simply distributes the signal in ρ_{alias} to the voxels involved according to their expected signal powers. In essence, $\mathbf{M}$ specifies where signal changes are likely to occur in x-f space, in a similar fashion to how the dynamic reference image in the original BLAST method (6) specifies where signal changes occur in image space. If only a single voxel is expected to have signal (i.e., $m_i = 0$ for all others), the solution becomes identical to the n-way decision described before. In general, $\mathbf{M}^2$ in Eq. [4] corresponds to the signal covariance matrix (i.e., $\mathbf{M}^2 = \langle \boldsymbol{\rho}, \boldsymbol{\rho}^{\mathrm{H}} \rangle$), so it is not restricted to being a diagonal matrix only. In this work, we treat $\mathbf{M}^2$ simply as a diagonal matrix, since the diagonal elements (i.e., the signal variance) can be estimated rather easily, as described below.

Equation [4] can be expanded in two ways. First, using the approach of Wiener filtering, Eq. [4] can be regularized to reduce the sensitivity to and to absorb some measurement noise. Second, from a statistical point of view, the signals in ρ can be considered to fluctuate around a baseline, which can be incorporated into the reconstruction to aid the determination of ρ. A regularized version of Eq. [4], incorporating the baseline estimate, is given by (see Ref. 27 for a similar formulation or eq. 16 in Ref. 28 from a Bayesian perspective):

$$\boldsymbol{\rho} = \underline{\boldsymbol{\rho}} + \mathbf{M}^2\mathbf{1}^{\mathrm{H}}(\mathbf{1M}^2\mathbf{1}^H + \psi)^{-1}(\rho_{alias} - \mathbf{1}\underline{\boldsymbol{\rho}}) \quad [5]$$

where ψ is the noise variance. $\underline{\boldsymbol{\rho}}$ is a complex-valued baseline estimate of ρ. In this expression, $\mathbf{M}^2$ is redefined to be the covariance matrix of the signal deviation from baseline $\underline{\boldsymbol{\rho}}$.

By solving Eq. [5] for every set of aliased voxels, one obtains the reconstructed object signals in x-f space. Finally, applying a Fourier transform along f yields the object signals in x-t space, which corresponds to a series of images over time t. This is the reconstruction procedure of *k-t* BLAST.

If multiple receiver coils are available, the data from them can be incorporated into the above equations in a consistent manner. Here, it is assumed that the sensitivity of the coils is time-invariant. The case of time-varying sensitivity is described below.

The use of multiple coils can be viewed as providing additional observations (i.e., additional matrix rows) to Eq.

[1]. Specifically, by incorporating the sensitivity information, Eq. [1] generalizes to:

$$\mathbf{S}\boldsymbol{\rho} = \boldsymbol{\rho}_{\mathbf{alias}} = \begin{bmatrix} \rho_{alias;1} \\ \rho_{alias;2} \\ \vdots \\ \rho_{alias;n_C} \end{bmatrix} \quad [6]$$

where $\mathbf{S}$ is the sensitivity matrix, containing the complex sensitivity values of the n_C receiver coils, as described in Ref. 5; $\rho_{\text{alias};\gamma}$ is the aliased signal from the γ-th coil. The regularized solution to Eq. [6], incorporating the baseline signals, the expected deviation from baseline, and the noise characteristics, is given by:

$$\begin{aligned} \boldsymbol{\rho} &= \underline{\boldsymbol{\rho}} + \mathbf{M}^2\mathbf{S}^{\mathrm{H}}(\mathbf{S}\mathbf{M}^2\mathbf{S}^{\mathrm{H}} + \boldsymbol{\Psi})^{-1}(\boldsymbol{\rho}_{\mathbf{alias}} - \mathbf{S}\underline{\boldsymbol{\rho}}) \\ &= \underline{\boldsymbol{\rho}} + (\mathbf{S}^{\mathrm{H}}\boldsymbol{\Psi}^{-1}\mathbf{S} + \mathbf{M}^{-2})^{-1}\mathbf{S}^{\mathrm{H}}\boldsymbol{\Psi}^{-1}(\boldsymbol{\rho}_{\mathbf{alias}} - \mathbf{S}\underline{\boldsymbol{\rho}}) \end{aligned} \quad [7]$$

where $\boldsymbol{\Psi}$ is the noise covariance matrix of the receiver coils (29). The two expressions in Eq. [7] are mathematically equivalent formulations (27). The first expression is more efficient to evaluate if $\mathbf{S}$ has more columns than rows (i.e., when the acceleration factor is higher than the number of receiver coils), and vice versa. This is the reconstruction formula of *k-t* SENSE.

It can be seen that the reconstruction formula is identical for *k-t* BLAST (Eq. [5]) and *k-t* SENSE (Eq. [7]), except for the interchange of $\mathbf{1}$ and $\mathbf{S}$ in the respective formulae. The only difference is that the equation is always underdetermined for *k-t* BLAST, whereas it can be overdetermined for *k-t* SENSE if the number of receiver coils is more than the acceleration factor. The transition between the under- and overdetermination regimes is handled in a smooth and consistent manner by the equations above.

To summarize, the raw data are acquired according to the specified *k-t* sampling pattern. The frequency-encoding direction is reconstructed by inverse Fourier transform. Then the data at each frequency-encoded position undergo separate reconstruction as follows. The data are arranged in a *k-t* array, with the dimensions being "the image size along the phase-encode direction" by "the number of time points." The unacquired portions of this array are filled with zeroes. For *k-t* SENSE, the data from each coil is arranged in a separate *k-t* array. Inverse Fourier transform is applied along *k* and *t* to yield the object signals in *x-f* space, which may suffer from aliasing. Equations [5] or [7] are solved for each set of aliased voxels for *k-t* BLAST or *k-t* SENSE, respectively. The reconstructed *x-f* array undergoes Fourier transform along *f* to yield an *x-t* array. Collecting the *x-t* arrays from all frequency-encoded positions yields the reconstructed images at all time points.

General Sampling in *k-t* Space

The reconstruction is considerably more demanding if the *k-t* sampling pattern does not conform to a lattice structure. To keep the formulation as general as possible, it is assumed that the data are acquired with multiple coils and that the coil sensitivities vary with time, such as due to subject motion. The *k*-space data acquired from the γ-th coil at time *t* is given by:

$$d_\gamma(\vec{k},\, t) = \int \rho(\vec{x},\, t) \cdot s_\gamma(\vec{x},\, t) \cdot e^{-i\vec{k}\cdot\vec{x}} d\vec{x} \quad [8]$$

where $\rho(\vec{x},\, t)$ and $s_\gamma(\vec{x},\, t)$ are the object signal and the sensitivity of the γ-th coil at position $\vec{x}$ and time *t*. If there is only a single coil, the factor $s_\gamma(\vec{x},\, t)$ can be removed, as it can be incorporated into $\rho(\vec{x},\, t)$. In matrix form, the problem is expressed as:

$$\mathbf{E}\boldsymbol{\rho}_{x,f} = \begin{bmatrix} \mathbf{d}_{k,t;1} \\ \mathbf{d}_{k,t;2} \\ \vdots \\ \mathbf{d}_{k,t;n_C} \end{bmatrix} \quad [9]$$

where $\boldsymbol{\rho}_{\mathbf{x,f}}$ represents the discrete object signals in *x-f* space, expressed as a vector. $\mathbf{d}_{k,t;\gamma}$ is the raw data in *k-t* space from the γ-th coil, expressed as a vector. $\mathbf{E}$ is an encoding matrix representing a discretization of the linear mapping between $\boldsymbol{\rho}_{\mathbf{x,f}}$ and $\mathbf{d}_{k,t;\gamma}$. In general, Eq. [9] also holds for the continuous domain, with the product $\mathbf{E}\,\boldsymbol{\rho}_{\mathbf{x,f}}$ representing a matrix where each entry represents the integral of the object signals multiplied with an encoding function.

The encoding matrix $\mathbf{E}$ includes the following operations: Fourier transform along *f* ($\mathrm{FT}_{f\to t}$), multiplication with the time-varying sensitivity of the γ-th coil ($\mathrm{S}_{x,t;\gamma}$), Fourier transform along *x* ($\mathrm{FT}_{x\to k}$), and retaining only the sampled points in *k-t* space ($\Xi_{k,t}$):

$$\mathbf{E} = \begin{bmatrix} \Xi_{k,t} & \mathrm{FT}_{x\to k} & \mathbf{S}_{x,t;1} & \mathrm{FT}_{f\to t} \\ \Xi_{k,t} & \mathrm{FT}_{x\to k} & \mathbf{S}_{x,t;2} & \mathrm{FT}_{f\to t} \\ & & \vdots & \\ \Xi_{k,t} & \mathrm{FT}_{x\to k} & \mathbf{S}_{x,t;n_C} & \mathrm{FT}_{f\to t} \end{bmatrix} \quad [10]$$

The regularized weighted-minimum-norm solution to Eq. [9] is obtained in an analogous fashion as before (Eq. [7]):

$$\begin{aligned} \boldsymbol{\rho}_{\mathbf{x,f}} &= \underline{\boldsymbol{\rho}_{\mathbf{x,f}}} + \mathbf{M}^2_{x,f}\mathbf{E}^{\mathrm{H}}(\mathbf{E}\mathbf{M}^2_{x,f}\mathbf{E}^{\mathrm{H}} + \boldsymbol{\Psi}_{k,t})^{+}(\mathbf{d}_{k,t} - \mathbf{E}\underline{\boldsymbol{\rho}_{\mathbf{x,f}}}) \\ &= \underline{\boldsymbol{\rho}_{\mathbf{x,f}}} + (\mathbf{E}^{\mathrm{H}}\boldsymbol{\Psi}^{-1}_{k,t}\mathbf{E} + \mathbf{M}^{-2}_{x,f})^{+}\mathbf{E}^{\mathrm{H}}\boldsymbol{\Psi}^{-1}_{k,t}(\mathbf{d}_{k,t} - \mathbf{E}\underline{\boldsymbol{\rho}_{\mathbf{x,f}}}) \end{aligned} \quad [11]$$

where $\underline{\rho_{x,f}}$ is an estimate of the baseline signals in *x-f* space. $\mathbf{M}^2_{x,f}$ is the signal covariance matrix. $\Xi_{k,t}$ is the noise covariance matrix of the coils in *k-t* space. $\Xi_{k,t}$ has a nested structure, and the entries corresponding to the covariance between two time points are zero, since the noise at different time points should be uncorrelated.

In general, it is computationally demanding to solve Eq. [11]. Nevertheless, considerable simplification can be achieved if the coil sensitivities are time-invariant or the *k-t* sampling pattern conforms to a lattice structure. Otherwise, direct analytic solution of the inverse problem may be infeasible even for moderate image sizes and number of time points and iterative approaches are required (29).

Obtaining Signal Estimates

The reconstruction formulae outlined above require an estimate of the baseline signal $\underline{\boldsymbol{\rho}_{\mathbf{x,f}}}$ and an estimate of the

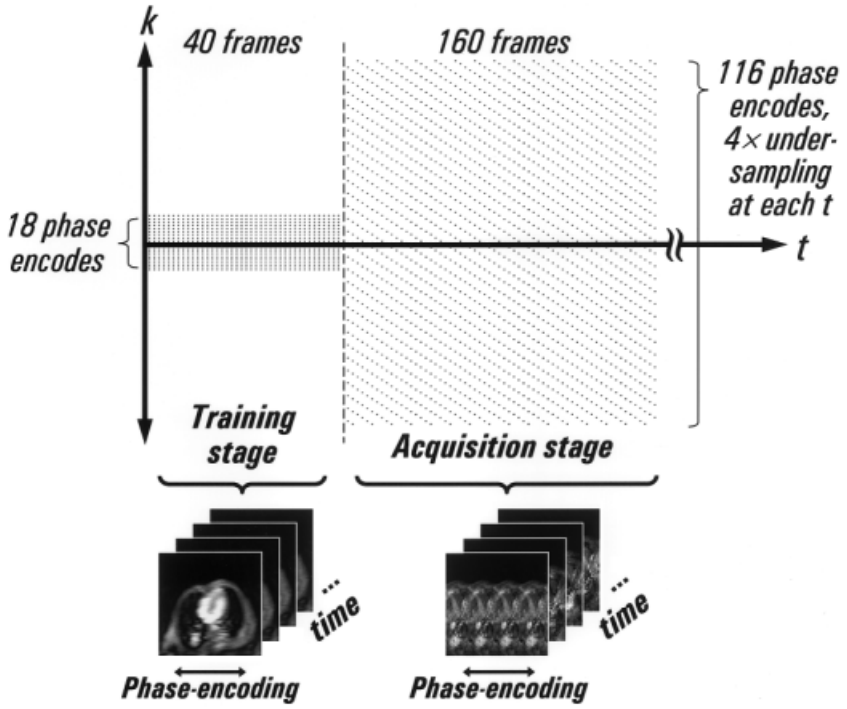

FIG. 2. *k-t* sampling pattern used in simulation, consisting of 40 and 160 frames for the training and acquisition stages, respectively. In the training stage, only 18 central phase-encode lines were used at each time frame, resulting in a series of low-resolution training images. In the acquisition stage, only every fourth phase-encode line was used at each time frame, resulting in a series of aliased images. The phase-encoding direction was horizontal in the images shown at the bottom.

squared deviation from baseline $\mathbf{M}^2_{x,f}$. The corresponding ρ and $\mathbf{M}^2$ from Eqs. [5] and [7] are simply truncated versions of $\underline{\boldsymbol{\rho}}_{\mathbf{x,f}}$ and $\mathbf{M}^2_{x,f}$ in Eq. [11], incorporating only the entries pertaining to each set of aliased *x-f* voxels. To simplify notation, the more general $\underline{\boldsymbol{\rho}}_{\mathbf{x,f}}$ and $\mathbf{M}^2_{x,f}$, corresponding to the full data, will be used to refer to both hereafter.

We obtained $\underline{\boldsymbol{\rho}}_{\mathbf{x,f}}$ directly from the sparsely sampled data and introduced a separate training period for obtaining $\mathbf{M}^2_{x,f}$. Specifically, each dynamic MRI experiment is divided into a short training stage and the usual acquisition stage (Fig. 2). The scanner parameters are identical between the two stages, except for the *k*-space trajectory. In the training stage, *k*-space is densely sampled, but only in the low spatial frequencies. This yields a series of low-resolution images at a high frame rate. In general, the intensities of the training images should be scaled similarly to those in the final reconstructed images, since the former will be used to estimate $\mathbf{M}^2_{x,f}$ of the latter. Special attention is needed to correct for normalizations of any applied filters, or of the inverse Fourier transform (e.g., when the training and acquisition data have different number of time points).

The raw training data in *k*-space undergo the following steps to yield $\mathbf{M}^2_{x,f}$ (Fig. 3). First (step 3.1 in Fig. 3), filtering is optionally applied to the raw data in *k*-space to reduce ringing artifacts from truncation. The choice of filtering is subject to the usual considerations regarding the tradeoff between blurring and artifact reduction. Second (step 3.2), inverse Fourier transform is applied to *k*-space. For *k-t* SENSE, the data from the multiple coils are combined using SENSE (5). From this point onwards, each frequency-encoded position is processed separately. Then (step 3.3) inverse Fourier transform is applied along *t*, and we denote the result as $\rho_{\text{train}}(x,f)$. The temporally invariant (i.e., direct-current or DC) term (i.e., $\rho_{\text{train}}(x,0)$) is set to zero (step 3.4). Filtering is optionally applied along *f* to reduce noise with high temporal frequencies (step 3.5). Since the training images have low resolution, they may potentially underrepresent certain fine image features. A simple pragmatic solution is to multiply $\rho_{\text{train}}(x,f)$ with a scaling factor (e.g., 2×, step 3.6), which provides a safety margin by attenuating the regularization. A higher value for this safety margin allows more image features to be reconstructed at the expense of some noise increase. Since its effect on image quality is gradual, much like the progressive effects introduced by common filtering techniques, its precise value is not critical and can be adjusted by the user based on the signal-to-noise ratio. Taking the squared magnitude of the result (i.e., $|\rho_{\text{train}}(x,f)|^2$) yields the diagonal elements of $\mathbf{M}^2_{x,f}$ (step 3.7). This step may require interpolation if $|\rho_{\text{train}}(x,f)|^2$ and $\mathbf{M}^2_{x,f}$ have different sizes. Linear interpolation can be used for computational efficiency, since $\mathbf{M}^2_{x,f}$ is only an estimate of the true covariance matrix, so it is unnecessary to achieve very high accuracy in the interpolation.

To obtain the baseline estimate $\underline{\boldsymbol{\rho}}_{\mathbf{x,f}}$, the sparsely sampled *k-t* space data from the acquisition stage undergo the following steps (Fig. 4). The sparsely sampled data are averaged along *t* to yield a temporal average (step 4.1 in Fig. 4). The temporally averaged *k*-space is first subtracted from the corresponding data at each time frame (step 4.2). The temporally averaged *k*-space and the subtracted data at each time frame were separately inverse Fourier transformed to yield a temporally averaged image and a series of aliased images, respectively (step 4.3). Each frequency-encoded position (e.g., see vertical dashed line) is processed separately hereafter as follows. The image column from the temporally averaged image is used as the DC term (at $f = 0$) of the baseline estimate $\underline{\boldsymbol{\rho}}_{\mathbf{x,f}}$ (step 4.4). The value of $\underline{\boldsymbol{\rho}}_{\mathbf{x,f}}$ is set to zero at all other *f*. The image columns from

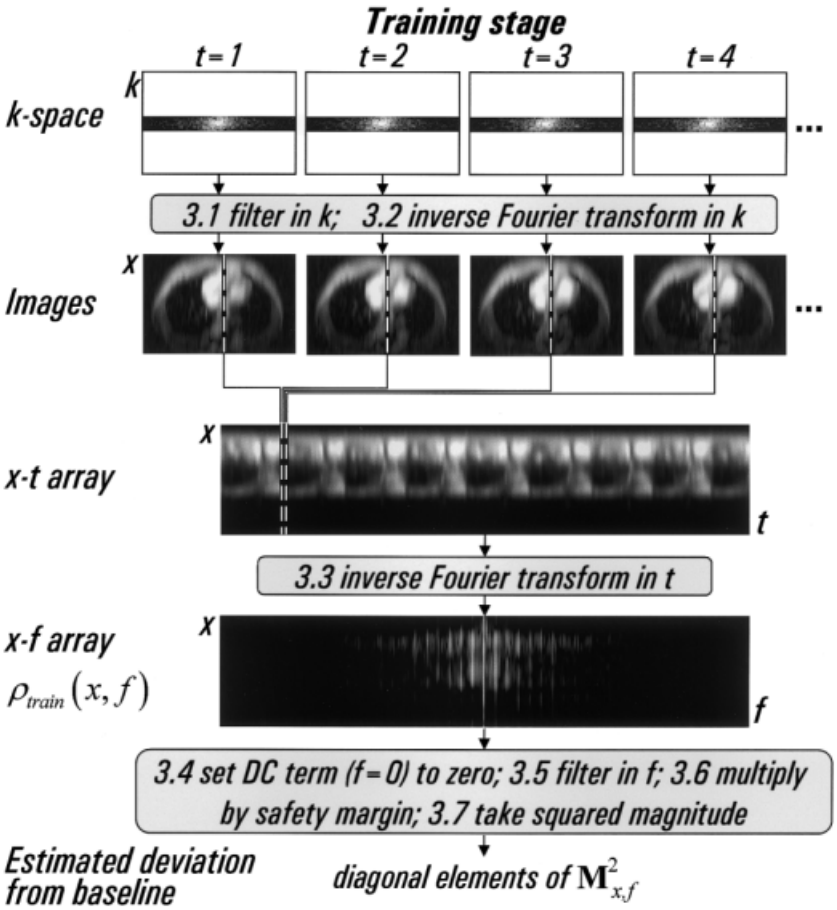

FIG. 3. Main processing steps for the training stage of *k-t* BLAST. The processing steps were similar for *k-t* SENSE, with an additional step to combine the data from all coils using SENSE. Only 18 central phase-encode lines were acquired at each time frame and the unacquired *k*-space data were zero-filled. The data were filtered along the phase-encoding direction to reduce truncation artifacts (3.1) and inverse Fourier transformed (3.2) to yield a series of low-resolution training images. Each frequency-encoding position (e.g., see vertical dashed line) was processed separately hereafter as follows. The image columns from all time frames were gathered into an *x-t* array and inverse Fourier transformed along *t* (3.3) to yield an *x-f* array. The central vertical strip at $f = 0$, corresponding to the DC term, was set to zero (3.4). Filtering was applied to attenuate the high temporal frequencies (3.5). The intensity magnitude in the final *x-f* array was multiplied by a safety margin (3.6). Taking the squared magnitude (3.7) yielded the diagonal elements of $\mathbf{M}^2_{x,f}$, representing the estimated squared deviation from baseline.

the aliased images are gathered and inverse Fourier transformed along *t* (step 4.5) to yield an *x-f* array, which represents the difference data for reconstruction, corresponding to the term in the rightmost parentheses in Eqs. [5] and [7] and the equivalent term in *k-t* space from Eq. [11].

Two features can be noticed about the difference data. First, the central vertical strip corresponding to the DC term is zero, since the DC component has been subtracted already. Second, the signals are replicated (i.e., aliased) 4-fold along the diagonal from the upper right to the lower left, according to the alias pattern shown in Fig. 1. The aliasing in the difference data is subsequently resolved (see bottom panel of Fig. 4) by reconstructing with Eqs. [5], [7], or [11] for lattice sampling with one coil, lattice sampling with multiple coils, or arbitrary sampling with an arbitrary number of coils, respectively.

MATERIALS AND METHODS

Simulation

k-t BLAST and *k-t* SENSE were validated by simulation using a previously reconstructed ungated cardiac image series (30), which was obtained with SENSE at a frame rate of 31.8 fps and a matrix size of 128 × 116 (frequency × phase). The image series consisted of 200 complex-valued frames. From these, 40 frames were used in low resolution (18 phase-encode lines) for the training stage, while only 25% of the data in the remaining 160 frames were used for the acquisition stage, simulating a 4-fold acceleration. A sequential interleaved *k-t* sampling pattern was used, as shown in Fig. 2. The simulation assumed a uniform body coil for *k-t* BLAST, and a six-element phased array placed around the body for *k-t* SENSE. The complex-valued sensitivities of the phased array coils were calculated using Biot-Savart's law. The sensitivity maps used in generating the simulation data were also used in the reconstruction. Normally distributed complex-valued noise (SD = 10% of mean intensity) was added to the data.

The *k-t* BLAST and *k-t* SENSE reconstructions were compared with the original images. The reconstruction errors are reported as the relative artifact power, which was calculated as the mean squared absolute difference between the reconstructed and the true images at each time point, normalized by the mean squared absolute intensity of the true image. For comparison, the results are also shown for sliding-window reconstruction using the same data as for *k-t* BLAST. The sliding window was symmetric in time with an effective length of four time points. Thus, the image at time *t* was reconstructed using the data from $t-1$, t, and $t+1$ and the average of the data from $t-2$ and $t+2$.

Experiments

Ungated cardiac imaging experiments were conducted on a healthy volunteer in a Philips 1.5 T Intera whole body

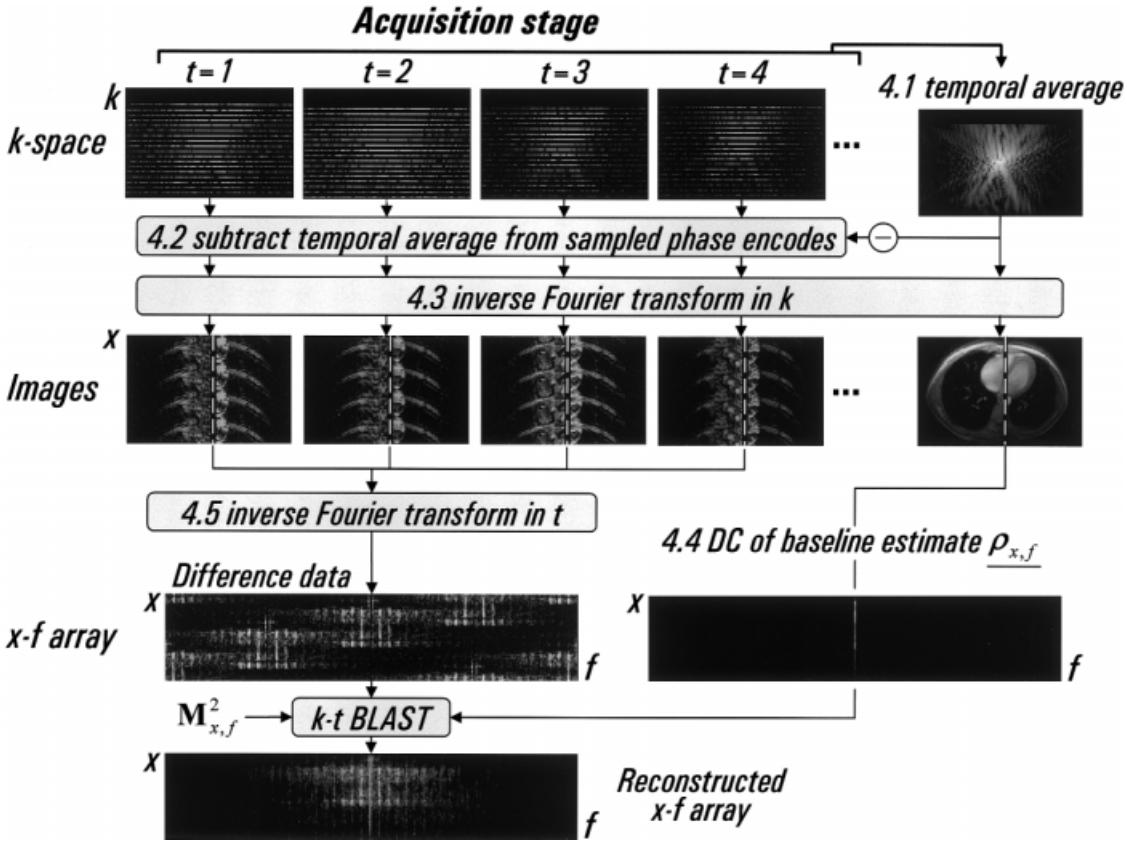

FIG. 4. Main processing steps for the acquisition stage of *k-t* BLAST. The processing steps were similar for *k-t* SENSE and the data from all coils were combined during the final reconstruction step (using *k-t* SENSE instead of *k-t* BLAST). Only every fourth phase-encode line was acquired at each time frame and the unacquired *k*-space data were zero-filled. The black band across the top of each *k*-space reflects the slightly asymmetric *k*-space coverage. These data were averaged to yield a temporally averaged *k*-space (4.1), which was first subtracted from the corresponding data at each time frame (4.2). The temporally averaged *k*-space and the subtracted data at each time frame were separately inverse Fourier transformed to yield a temporally averaged image and a series of aliased images, respectively (4.3). Each frequency-encoding position (e.g., see vertical dashed line) was processed separately hereafter as follows. The image column from the temporally averaged image was used as the DC term (at $f = 0$) of the baseline estimate (4.4). The rest of the baseline estimate was set to zero. The image columns from the aliased images were gathered and inverse Fourier transformed along *t* (4.5) to yield an *x-f* array, which can be seen to suffer from aliasing. This array was reconstructed with *k-t* BLAST, incorporating information from the baseline estimate and the estimated deviation from baseline. The aliasing was effectively resolved after reconstruction.

scanner (Philips Medical Systems, Best, the Netherlands). Informed consent was obtained before imaging. For the *k-t* BLAST experiment, a single 20 cm-diameter surface coil was taped to the subject's chest. For *k-t* SENSE experiment, a six-element phased array was mounted to a rigid frame around the subject's body, as described by Weiger et al. (30). The imaging slice was placed transversely at the level of the heart.

Data were acquired in a free-running mode using a segmented gradient-echo EPI sequence (24), with an echo train length of 9, an FOV of 310 × 201.5 mm^2, a slice thickness of 10 mm, a matrix size of 128 × 80 (frequency × phase), a repetition time (TR) of 13 ms, an echo time of 6.2 ms, and echo time shifting (24).

For the training stage, phase-encode lines −9 to 8 were acquired by two EPI segments in an interleaved manner. For the acquisition stage, phase-encode lines −30 to 41 were acquired by eight EPI segments in an interleaved manner, with a segment order of 0, 4, 1, 5, 2, 6, 3, 7. The data from every two EPI segments were combined into a single time point, resulting in a frame duration of 2 TR = 26 ms. This allowed an increased number of phase-encode lines for each training image without an excessively long echo train, and it retained a lattice sampling pattern for the undersampled data. Generally, lumping adjacent data to the same time point has little impact on image quality, as long as the combined duration is short relative to the object motion. In all, 400 frames were acquired for the training stage and for the acquisition stage.

To construct the coil sensitivity maps for *k-t* SENSE reconstruction, a temporally averaged image was obtained separately for each coil in the training stage and in the acquisition stage. These temporally averaged images were divided by the mean absolute image of all coils and smoothed to reduce noise. Therefore, a separate set of sensitivity maps was obtained for the training and acquisition stages. If the coil sensitivity is time-varying, this procedure can be modified to that of Kellman et al. (22), in which time-varying sensitivity maps are obtained from sliding-window (i.e., moving-average) reconstructions,

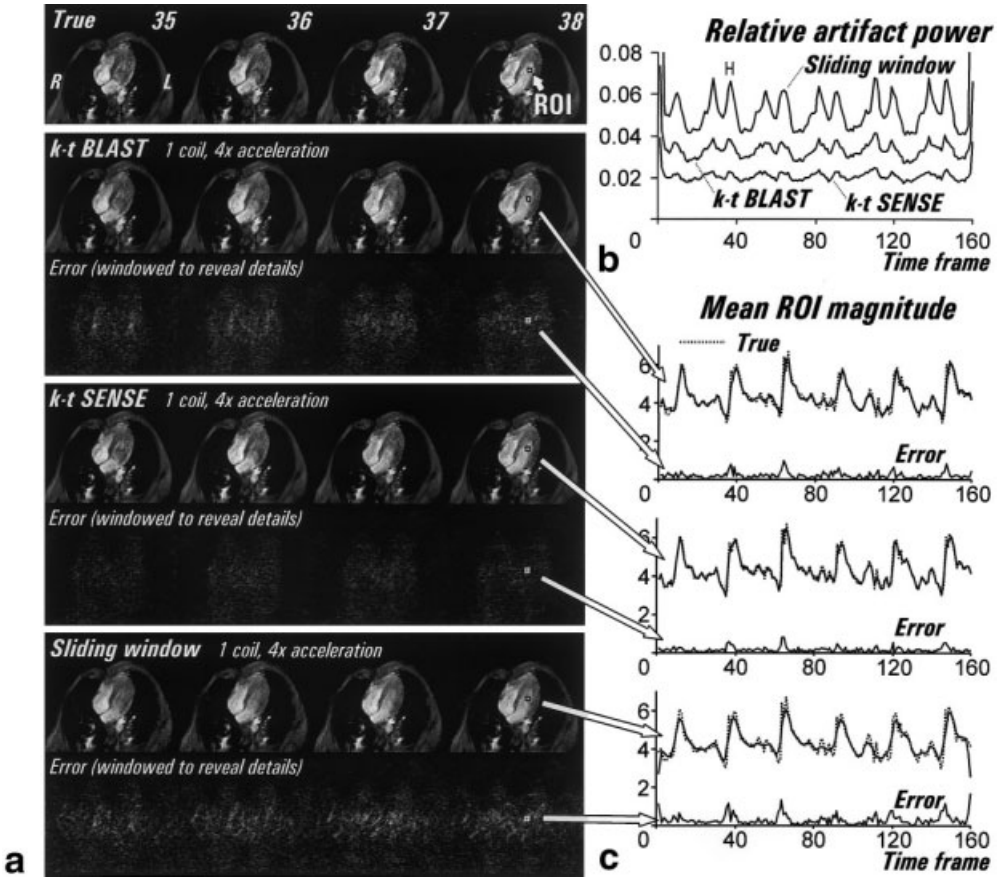

FIG. 5. **a:** Simulation results, comparing the true images and reconstructed images from *k-t* BLAST, *k-t* SENSE, and sliding window over four representative and consecutive frames. **b:** Relative artifact power for all methods over time. The H-shaped marker indicates the time frames shown in **a**. **c:** Absolute mean intensity over time for an ROI placed in the left ventricle of the reconstructed images. The corresponding curve from the true images is shown in dotted line for comparison and the magnitude of the complex difference between each reconstructed curve and the true curve is shown at the bottom of each panel. Both *k-t* BLAST and *k-t* SENSE captured the overall dynamics well, but sudden intensity variations led to increased errors. In general, the errors were lower for *k-t* SENSE. Sliding window captured a temporally low-pass version of the intensity time curve, so it was unable to depict the sudden time-varying features.

rather than the globally time-averaged reconstructions used in the present implementation. For the training stage, the data from all coils at each time frame were combined during reconstruction with SENSE.

General Reconstruction Issues

The noise variance (for *k-t* BLAST) or the noise covariance matrix (for *k-t* SENSE) was estimated using background voxels. For the training stage, filtering was applied to *k*-space along the phase-encoding direction and to the temporal frequency direction. The *k*-space filter was a Hamming filter to reduce truncation artifacts and the temporal frequency filter was a low-pass filter, with a central pass band of 50% bandwidth and half-Hanning-shaped transition bands of 20% bandwidth on either side. A safety margin of 2 was used in the determination of $\mathbf{M}^2_{x,f}$ (see step 3.6 in Fig. 3).

RESULTS

Simulation

Figure 5a shows representative images over four consecutive time frames. The true images are shown on the top row for comparison. The reconstructed images from *k-t* BLAST, *k-t* SENSE, and sliding window and the corresponding error images are shown below. For *k-t* BLAST, the reconstructed images captured the overall dynamics well, although a slight loss of image details is discernible for fast-moving signals. For *k-t* SENSE, the reconstructed images captured the overall dynamics even better, and virtually all image details were represented faithfully. The corresponding error images contained mostly random noise patterns for both methods. The sliding-window reconstructions show more temporal blurring. For example, the moving front of fresh blood in the left ventricle is visualized less crisply. The corresponding error images contained noise as well as an increased amount of residual aliasing for fast-moving signals, indicating that it was less apt at reconstructing the higher temporal frequencies. The error images show that the reconstructed images from *k-t* BLAST and *k-t* SENSE contain less noise than those from sliding window, due to regularization of the reconstruction formulae, which absorbs some of the noise.

Figure 5b shows the reconstruction error as the relative artifact power at different time frames. The error was highest for sliding window reconstruction. Between *k-t* BLAST and *k-t* SENSE, the error was lower for the latter, as expected, since *k-t* SENSE had data from additional receiver coils. In the present study the reconstruction was underdetermined for *k-t* BLAST, while it was overdetermined for *k-t* SENSE. The error for *k-t* SENSE was nonzero because of the added noise in the simulation and the specific treatment of the DC term, as described in the Discussion. The error was higher for all methods at the beginning and end of the image series. For *k-t* BLAST and *k-t* SENSE, this was due to the temporal discontinuity between the initial and last frames of the image series, while the Fourier transform implicitly assumed periodic boundaries. For

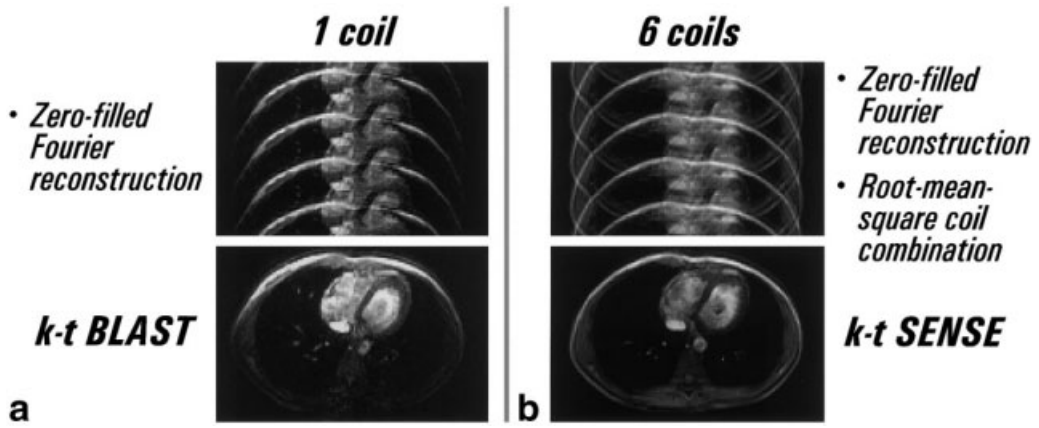

FIG. 6. Ungated cardiac imaging experiment with 4-fold acceleration, comparing representative (**a**) *k-t* BLAST and (**b**) *k-t* SENSE reconstructions (bottom) with their respective aliased images (top). The aliased images were reconstructed from the undersampled data using zero-filled Fourier reconstruction, followed by root-mean-square coil combination for the multiple-coil data. The reconstructed image from *k-t* BLAST exhibits an intensity drop off towards the posterior due to the use of a single surface coil placed on the chest. The aliasing artifacts from undersampling were eliminated in the *k-t* BLAST and *k-t* SENSE reconstructions.

sliding window, this was due to the sliding window extending beyond the acquired time frames. In practice, the increased error at the beginning and the end of the acquisition stage does not pose a problem, since it is easily overcome by acquiring several additional image frames in the beginning and the end of the acquisition stage and discarding the extra frames after reconstruction. The reconstruction error was also seen to exhibit temporal periodicity, which was related to the periodic changes in the image contents from cardiac motion.

Figure 5c shows the temporal variation of the mean intensity magnitude for a 3×3 region of interest (ROI) placed in the left ventricle. The corresponding curve from the true images is overlaid as a dotted line for comparison and the magnitude of the complex difference between each reconstructed curve and the true curve is shown at the bottom. The results confirm those from visual inspection. Compared to the intensity time curve of the true images, *k-t* BLAST and *k-t* SENSE were able to reconstruct the overall temporal dynamics well. However, sudden changes in intensity, represented by sharp corners in the intensity time curve, were somewhat temporally smoothed, and the smoothing was less for *k-t* SENSE. The temporal smoothing was noticeably worse for sliding-window reconstruction and the corresponding error was higher. This smoothing was expected for sliding-window reconstruction, which is effectively a low-pass temporal filter, resulting in attenuation of higher temporal frequencies.

Experiments

Figure 6 shows a representative frame from the *k-t* BLAST (Fig. 6a) and *k-t* SENSE (Fig. 6b) reconstructions (bottom), compared with zero-filled Fourier reconstructions (top) followed by root-mean-square coil combination for the multiple-coil data. It should be noted that the data for *k-t* BLAST and *k-t* SENSE were acquired in separate experiments, so the images cannot be directly compared in a voxel-by-voxel manner. The zero-filled Fourier reconstructions exhibited aliasing artifacts due to the 4-fold undersampling along the phase-encoding direction (vertical direction). This aliasing was eliminated by *k-t* BLAST and *k-t* SENSE reconstruction. The reconstructed image from *k-t* BLAST exhibits an intensity dropoff towards the posterior due to the use of a single surface coil placed on the chest. In comparison, the reconstructed image from *k-t* SENSE had a more uniform intensity distribution, due to the use of a phased array placed around the body.

Figure 7 shows additional reconstructed images from *k-t* BLAST (Fig. 7a) and *k-t* SENSE (Fig. 7b) in consecutive frames. Both image series revealed the rapid dynamics of the blood inside the myocardium and in the aorta. The complex flow patterns, the myocardial motion, and the free-breathing motion of the lungs can be more readily appreciated when viewed in cine.

DISCUSSION

We have presented the theory and experimental verification of *k-t* BLAST and *k-t* SENSE. Both methods are based on the premise that the raw data in dynamic imaging exhibit correlations in *k*-space and in time. Thus, it is sufficient to acquire only a reduced amount of data and recover the missing portion afterwards, thereby leading to increased spatiotemporal resolution. We showed that the correlations in *k*-space and time could be determined from a set of training images, which allowed the signal distribution in *x-f* space to be learned. By exploiting this prior information, it was feasible to push the spatial and temporal resolutions of ungated imaging beyond current capabilities (30) during the acquisition stage. In this work, we presented experimental results for ungated cardiac imaging, demonstrating a 4-fold acceleration with a voxel size of 2.42×2.52 mm^2 at 38.4 fps. The concept of *k-t* BLAST and *k-t* SENSE can be combined with imaging sequences other than EPI, and they can be used with partial Fourier acquisition to allow asymmetric *k*-space coverage in order to achieve short TE and TR and a higher temporal resolution.

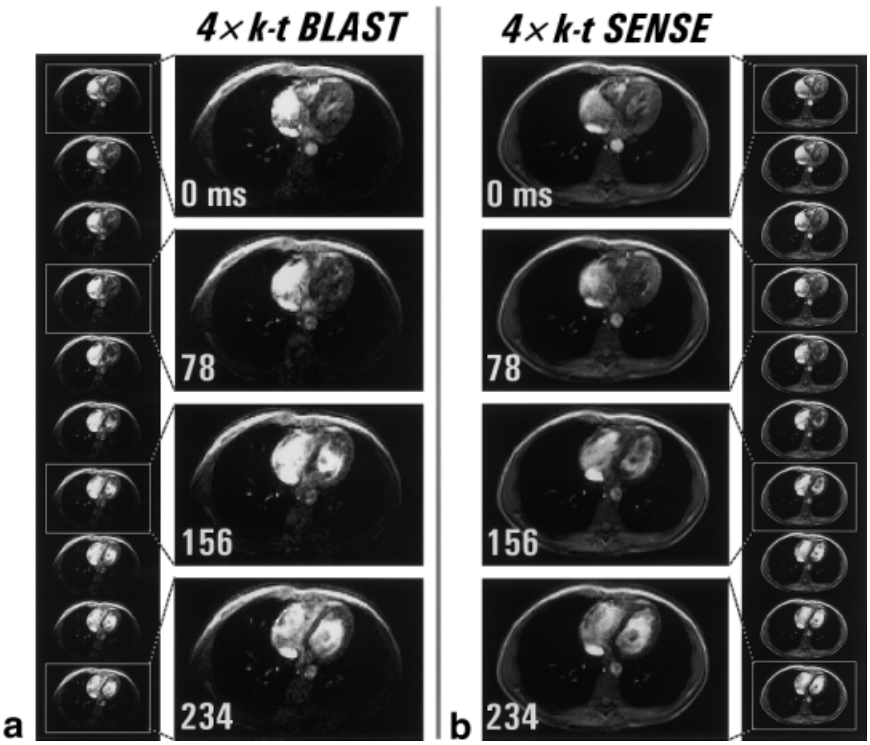

FIG. 7. Ungated cardiac imaging experiment with 4-fold acceleration, showing (**a**) *k-t* BLAST and (**b**) *k-t* SENSE reconstructions at consecutive time frames with a temporal resolution of 26 ms. The *k-t* BLAST images exhibit an intensity drop-off due to the use of a single surface coil placed on the chest. The images from *k-t* BLAST and *k-t* SENSE were acquired in separate experiments, so they cannot be compared directly in a frame-by-frame manner.

k-t BLAST and *k-t* SENSE are quite general in the sense that they can handle a wide range of object motions. The only requirement is that the motion learned from the training stage is representative of the motion in the acquisition stage, so the training images need to have sufficient resolution and duration to observe changes in image contents (e.g., ~5–10 sec in duration and ~10–20 phase-encode lines for real-time cardiac imaging). Nevertheless, the actual form of the motion can be arbitrary, although motions that are more localized in *x-f* space (i.e., more localized in space (26) and more periodic) favor the reconstruction, since they allow denser signal packing in *x-f* space.

Treatment of DC Term

In the present methods, the DC term (i.e., data for $f = 0$) was reconstructed differently from all other temporal frequencies. This special treatment of the DC term makes use of the knowledge that in typical image series the DC term has by far the strongest signal intensities, while the signal intensities of the aliased frequencies are likely to be insignificant in comparison. Since the reconstruction formula contains quantities that may be inexact (e.g., baseline signal, signal covariance, and coil sensitivities), one may commit more errors by trying to resolve this aliasing than by overlooking it. Hence, in the present methods the DC term was obtained simply as the temporal average of all of the data from the acquisition phase. The result of this procedure was that any temporal frequencies that were aliased to the DC term were not resolved. Therefore, in the reconstructed images the signal intensities of these aliased frequencies were added to the DC term, while their own reconstructed intensities were set to zero. In practice, this has negligible impact on the reconstructed images, since it only affects the narrow bands of frequencies that alias exactly onto the DC term. If the aliased frequencies are known to have strong signals, one may also alter the sequence timing slightly such that those frequencies no long alias exactly onto the DC term.

k-t Sampling Issues

Reconstruction accuracy can be improved further by optimizing the *k-t* sampling pattern. For example, the data from the training stage can be used as an estimate of the signal distribution. Then, the effects of sampling can be assessed by decimating the data and trying to recover them afterwards. The optimal sampling pattern is the one that leads to minimal signal degradation. However, this optimization is impractical in general due to high computational demand and no analytic solutions are known (31). Nevertheless, the problem can be made tractable if certain simplifying assumptions are made about the signal distribution. In those cases, the optimal sampling pattern can be obtained efficiently from either a precalculated library (32) or on-the-fly calculation (18,19).

In the present study, the *k-t* sampling pattern was predetermined, so it was not optimized for the actual distribution of object signals in *x-f* space. Additionally, the signal distribution was not assumed to adopt any particular shape (17). Nevertheless, *k-t* BLAST and *k-t* SENSE were able to produce good reconstruction results even without sampling optimization. This robustness stemmed from the fact that the present methods did not need to avoid all signal aliasing with the choice of *k-t* sampling pattern only (17). The weighted-minimum-norm approach also contributed to resolving the signal aliasing by using knowledge of the signal distribution in *x-f* space and of the coil sensitivities. The aliasing was completely resolved when the reconstruction problem was overdetermined and

well-conditioned, and it was partially resolved otherwise. As a result, the reconstruction quality was not entirely dependent on the sampling pattern and sampling optimization was therefore less critical. This flexibility of using arbitrary sampling patterns simplifies the practical use of the present methods.

Comparison of *k-t* BLAST and *k-t* SENSE

k-t BLAST and *k-t* SENSE are closely related, except that *k-t* SENSE additionally accounts for the coil sensitivities. This difference underlies the respective advantages of these methods. When the coil sensitivities are not taken into account (i.e., *k-t* BLAST) and the *k-t* sampling pattern conforms to a lattice structure, the computation becomes exceedingly straightforward, involving no matrix inversion. Also, the modulation of the coil sensitivity can be incorporated into the object function $\rho(\bar{x},t)$ itself, thus eliminating the need for sensitivity assessment. These advantages allow for a simple experimental set-up and a short reconstruction time, which is especially important for high-dimensional (e.g., time resolved 3D imaging) and/or real-time applications. On the other hand, by accounting for the coil sensitivities (i.e., *k-t* SENSE), signal overlaps in *x-f* space are easier to resolve with the complementary information from multiple coils. This potentially allows tighter signal packing in *x-f* space, thereby further improving acquisition efficiency. With the multicoil data the reconstruction becomes more exact due to the incorporation of additional observations. As expected, the accuracy is improved when the acceleration factor is relatively low compared to the number of receiver coils. However, one has the freedom to increase the acceleration factor beyond the number of receiver coils, since the proposed methods make a smooth transition from the overdetermined to the underdetermined case. The degradation in image quality through this transition is gentle, as seen by comparing the presented results from *k-t* BLAST (i.e., underdetermined case) and *k-t* SENSE (i.e., overdetermined case).

CONCLUSIONS

The performance of dynamic imaging can be significantly enhanced by exploiting signal correlations in *k*-space and in time. These correlations are represented as the expected signal distribution in *x-f* space, which can be obtained, for example, by acquiring a small set of training data. *k-t* BLAST and *k-t* SENSE were developed to exploit these correlations as prior information to recover unacquired *k*-space data in a flexible and adaptive manner. These methods provide an integral approach to recovering the missing data based on all of the acquired data and available information. The general concept is applicable to arbitrary *k*-space trajectory and time-varying sensitivity information.

The present approach is fundamentally different from the conventional approach for dynamic imaging, which repeats the same scan over time to acquire a series of separate images. This "scan-and-repeat" approach treats every voxel at every time point as a separate unknown. Such a treatment is somewhat unnecessary unless every voxel in an image is completely random and its variation over time is entirely unpredictable. We have shown that the high degree of spatiotemporal correlations in natural image series can be exploited to improve the efficiency of the imaging process, which in turn allows significant improvements in the achievable spatial and temporal resolutions.

The present methods should not be confused with data-sharing methods, which apply a fixed interpolator to estimate any unacquired data in order to increase the apparent frame rate. This interpolation process does not resolve any potential signal aliasing and may lead to degradation of the true signals at higher temporal frequencies. In contrast, the present methods learn and adapt to the expected signal distribution. Thus, the estimation process is tuned to the available image contents, resulting in improved quality of the reconstructed images and achieving a high frame rate without additional temporal interpolation. *k-t* BLAST and *k-t* SENSE are applicable to a wide range of areas, particularly those with object motion exhibiting quasi-periodic behaviors. These include most anatomies of interest, such as the heart, the lungs, the abdomen, and the brain under periodic stimulation, among others.

ACKNOWLEDGMENTS

J.T. thanks Dr. Sebastian Kozerke for expert advice on implementation issues, and Prof. Paul C. Lauterbur for valuable comments and encouragement that initiated an effort culminating in the developments described here.

REFERENCES

1. Liang ZP, Constable T, Haacke EM, Boada F, Smith M, Lauterbur PC. Constrained reconstruction methods in MR imaging. Rev Magn Reson Med 1992;4:67–185.
2. McGibney G, Smith MR, Nichols ST, Crawley A. Quantitative evaluation of several partial Fourier reconstruction algorithms used in MRI. Magn Reson Med 1993;30:51–59.
3. Hu X, Parrish T. Reduction of field of view for dynamic imaging. Magn Reson Med 1994;31:691–694.
4. Sodickson DK, Manning WJ. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. Magn Reson Med 1997;38:591–603.
5. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: sensitivity encoding for fast MRI. Magn Reson Med 1999;42:952–962.
6. Tsao J, Behnia B, Webb AG. Unifying linear prior-information-driven methods for accelerated image acquisition. Magn Reson Med 2001;46: 652–660.
7. Jones RA, Haraldseth O, Muller TB, Rinck PA, Oksendal AN. *K*-space substitution: a novel dynamic imaging technique. Magn Reson Med 1993;29:830–834.
8. van Vaals JJ, Brummer ME, Dixon WT, Tuithof HH, Engels H, Nelson RC, Gerety BM, Chezmar JL, den Boer JA. "Keyhole" method for accelerating imaging of contrast agent uptake. J Magn Reson Imag 1993;3: 671–675.
9. Riederer SJ, Tasciyan T, Farzaneh F, Lee JN, Wright RC, Herfkens RJ. MR fluoroscopy: technical feasibility. Magn Reson Med 1988;8:1–15.
10. Parrish T, Hu X. Continuous update with random encoding (CURE): a new strategy for dynamic imaging. Magn Reson Med 1995;33:326–336.
11. Doyle M, Walsh EG, Blackwell GG, Pohost GM. Block regional interpolation scheme for *k*-space (BRISK): a rapid cardiac imaging technique. Magn Reson Med 1995;33:163–170.
12. Korosec FR, Frayne R, Grist TM, Mistretta CA. Time-resolved contrast-enhanced 3D MR angiography. Magn Reson Med 1996;36:345–351.
13. Zaitsev M, Zilles K, Shah NJ. Shared *k*-space echo planar imaging with keyhole. Magn Reson Med 2001;45:109–117.

14. Xiang QS, Henkelman RM. Dynamic image reconstruction: MR movies from motion ghosts. J Magn Reson Imag 1992;2:679–685.
15. Liang ZP, Jiang H, Hess CP, Lauterbur PC. Dynamic imaging by model estimation. Int J Imag Syst Technol 1997;8:551–557.
16. 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.
17. Tsao J. On the UNFOLD method. Magn Reson Med 2002;47:202–207.
18. Willis NP, Bresler Y. Optimal scan for time-varying tomography. I. Theoretical analysis and fundamental limitations. IEEE Trans Image Process 1995;4:642–653.
19. Willis NP, Bresler Y. Optimal scan for time-varying tomography. II. Efficient design and experimental validation. IEEE Trans Image Process 1995;4:654–666.
20. Willis NP, Bresler Y. Lattice-theoretic analysis of time-sequential sampling of spatiotemporal signals. I. IEEE Trans Info Theory 1997;43:190–207.
21. Willis NP, Bresler Y. Lattice-theoretic analysis of time-sequential sampling of spatiotemporal signals. II. Large space-bandwidth product asymptotics. IEEE Trans Info Theory 1997;13:208–220.
22. Kellman P, Epstein FH, McVeigh ER. Adaptive sensitivity encoding incorporating temporal filtering (TSENSE). Magn Reson Med 2001;45:846–852.
23. Xiang QS, Henkelman RM. *K*-space description for MR imaging of dynamic objects. Magn Reson Med 1993;29:422–428.
24. McKinnon GC. Ultrafast interleaved gradient-echo-planar imaging on a standard scanner. Magn Reson Med 1993;30:609–616.
25. Golub GH, Van Loan CF. Matrix computations, 3rd ed. Baltimore: Johns Hopkins University Press; 1996.
26. Tsao J. Extension of finite-support extrapolation using the generalized series model for MR spectroscopic imaging. IEEE Trans Med Imag 2001;20:1178–1183.
27. Pruessmann KP, Tsao J, Boesiger P. Feedback regularization for SENSE reconstruction. In: Proc 10th Annual Meeting ISMRM, 2002, Hawaii. p 739.
28. Rosenfeld D. New approach to gridding using regularization and estimation theory. Magn Reson Med 2002;48:193–202.
29. Pruessmann KP, Weiger M, Bornert P, Boesiger P. Advances in sensitivity encoding with arbitrary *k*-space trajectories. Magn Reson Med 2001;46:638–651.
30. Weiger M, Pruessmann KP, Boesiger P. Cardiac real-time imaging using SENSE. Magn Reson Med 2000;43:177–184.
31. Reeves SJ, Zhe Z. Sequential algorithms for observation selection. IEEE Trans Signal Proc 1999;47:123–132.
32. Nagle SK, Levin DN. Creating a library of generalized Fourier sampling patterns for irregular 2D regions of support. Magn Reson Med 2001;46:624–627.