Available online at www.sciencedirect.com

SCIENCE DIRECT

Magnetic Resonance Imaging 24 (2006) 133–154

MAGNETIC RESONANCE IMAGING

Resolution enhancement in MRI

Eyal Carmi[a,*], Siuyan Liu[b], Noga Alon[a], Amos Fiat[a,*], Daniel Fiat[c,*]

[a]School of Computer Science, Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel 39040
[b]Department of Bioengineering, University of Illinois at Chicago, Chicago, IL 60607, USA
[c]Magnetic Resonance Imaging Laboratory, Department of Physiology and Biophysics, University of Illinois at Chicago, Chicago, IL 60612, USA

Received 6 December 2004; revised 30 September 2005; accepted 30 September 2005

Abstract

We consider the problem of super-resolution reconstruction (SRR) in MRI. Subpixel-shifted MR images were taken in several fields of view (FOVs) to reconstruct a high-resolution image. A novel algorithm is presented. The algorithm can be applied locally and guarantees perfect reconstruction in the absence of noise. Results that demonstrate resolution improvement are given for phantom studies (mathematical model) as well as for MRI studies of a phantom carried out with a GE clinical scanner. The method raises questions that are discussed in the last section of the paper. Open questions should be answered in order to apply this method for clinical purposes.

Keywords: Resolution enhancement; Super-resolution reconstruction; Phantom

1. Introduction

In this article, we introduce a novel method for resolution enhancement in MR images. The problem of image resolution, whether in regular cameras or in MRI, is a well-studied problem. There are many algorithms and methods [1–14] that deal with the problem of having low-resolution images limited by the available hardware and methods of acquisition. This article introduces a simplified model of the low-resolution images. Based on this model and the assumption that we can achieve subpixel shifts of the images, we present a novel approach and algorithm for the super-resolution reconstruction (SRR) problem and show that under some conditions perfect high-resolution reconstruction is definite (in the absence of noise).

The main contributions of this paper are as follows:

- We give a new novel algorithm for subpixel super-resolution images. This algorithm has unique error propagation properties not known for other algorithms and suggests many new directions of study not previously considered.
- Many experiments have been performed on both simulated data and on specially constructed phantoms devised for super-resolution studies. The simulated experiments are very good. The actual MRI results are less satisfactory; we discuss further directions of study to validate the practicality of our schemes. The construction of special phantoms for the study of resolution enhancement is unique to our work. Clearly, a rigorous method of study is required.
- Despite previous arguments that SRR in MRI is impossible in various circumstances [1], we argue this is not true.

We deal with SRR and seek to avoid error propagation and unnecessary assumptions (see Section 3). Our main theoretical results appear in (Sections 4 and 5, Lemmas 4.1 and 5.1). In those, we give a localized reconstruction of a high-resolution image. We attain a resolution enhancement of factor 3, but our technique can give any desired improvement at the cost of additional noise.

Structure of the paper. In Section 2, the general SRR problem is discussed and a review of previous works on SRR in general and on SRR in MRI is presented. Sections 3-5 present a simplified model and algorithm for one-, two- as well as for higher dimensional cases. Section 6 includes a description of the experiment conducted in order to verify the algorithm. Analysis of the input data is provided in order to justify the use of the proposed method in the case of MRI. Experimental results are provided in Section 7. This paper

* Corresponding authors.
E-mail addresses: eyal_carmi1@recanati-alum.tau.ac.il (E. Carmi), fiat@tau.ac.il (A. Fiat), fiat@uic.edu (D. Fiat).

0730-725X/$ – see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.mri.2005.09.011

presents results of model-simulated images as well as images acquired with a GE MRI scanner of a phantom designed and constructed for this purpose. Section 8 discusses unanswered open issues.

1.1. MRI spatial resolution

MRI provides intensities for each voxel. Those intensities are proportional to the number of nuclei in each voxel and are affected by the nuclear relaxation times and the pulse sequence used. Those effects affect the image contrast [15]. MRI spatial resolution is determined by gradients' intensity, digital imaging filter bandwidth, the number of "readout" points and phase encoding steps. MRI resolution along the third dimension (Z) in 2-D pulse sequences is determined by the slice selection pulse. Enhancement of the spatial resolution may be achieved by (a) shifting the frame of reference in steps smaller than the pixel or voxel size (this can be done along one, two and three dimensions) and (b) carrying out complementary measurements at several fields of view (FOVs). A more complete description will be given in the paper.

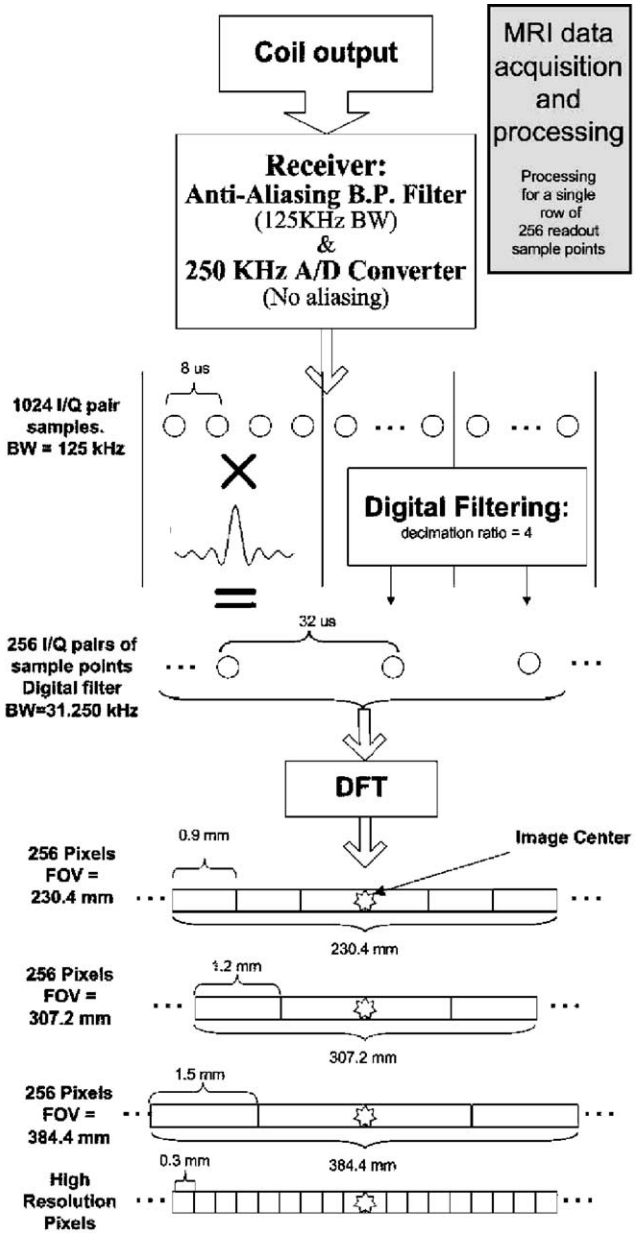

Fig. 1. MRI data acquisition and processing. The figure demonstrates the steps in generating a single 1×256 line of the MR image. The process begins from the raw analog input of the MR device and ends with 256 readout points; each represents a pixel in the final image. This process is repeated 256 times (256 phase encoding steps) to generate the final 256×256 MR image.

It is reasonable to assume that the high-resolution reconstruction procedure to be described does not modify signal intensity and image contrast, at least in first-order approximation.

1.2. Definitions and notations

MR Images can be shifted along three orthogonal axes. Along positive and negative directions: right and left (X, $-X$), anterior and posterior (Y, $-Y$), and superior and inferior (Z, $-Z$).

X and Y directions consist of the frequency and phase encoding directions (or vice versa — operator's option). Shifts in the frequency (readout) direction are achieved by shifting the receiver local oscillator frequency. (The GE scanner has a 10-MHz clock that enables frequency modulations in integer multiples of 0.596 Hz.) Shifts in the phase encoding direction are achieved by changing the receiver local oscillator phase. Shifts in the Z direction are achieved by moving the subject along the Z direction or by varying the transmitter oscillator frequency.

Enhancement of the spatial resolution can also be accomplished by carrying out measurements at (1) several bandwidths as FOV can be modified either by modifying the intensity of the magnetic field gradient or by changing the bandwidth and (2) by carrying out the measurements at several digitization rates and several phase encoding steps. A 2-D spin echo pulse sequence was used. For MR image creation process (based on GE specifications), see Fig. 1.

1.3. Resolution limitations

Instrumental limitations, signal to noise, and nuclear relaxation times considerations impose limitations on the maximum feasible spatial resolution.

Spatial resolution can be enhanced by:

(1) Decreasing the FOV.
(2) Increasing the number of readout points.
(3) Increasing the number of phase encoding steps.

Where:

(1) FOV is limited by (a) gradient strength and (b) subject dimension in the readout direction.
(2) The number of readout points is limited by the transverse nuclear relaxation time (T_2). Extending the readout period significantly beyond the transverse relaxation time decreases the signal-to-noise

ratio (SNR) significantly. The maximum theoretical number of readout points is limited by the local oscillator frequency divided by the exciter/receiver register bit. This limit is impractical since the long acquisition time decreases the SNR to an unfeasible degree. Though T_2 decay is usually the only important limit, the amount of available memory for storing the data also comes up occasionally as a limiting factor.

(3) The number of phase encoding steps is limited by the acquisition time. Increasing the number of phase encoding steps increases the acquisition time, proportionally.

In this paper, we demonstrate that planar spatial resolution may be enhanced by (a) carrying out MRI measurements at three FOVs and (b) shifting the center of each FOV by a predetermined subpixel distance as will be further explained in Section 6.2.

The merit of the proposed method is that it circumvents the limitations of Points 1 and 2 mentioned above.[1]

MRI applications of the proposed method are as follows:

(1) It has the potential to extend spatial resolution to microscopic levels for all nuclei.
(2) It enhances the resolution of nuclei with low gyromagnetic ratio as ^{17}O.
(3) It enhances the spatial resolution in studies of nuclei with short nuclear relaxation times where the fast decay limits effective readout time, e.g., ^{17}O.

A more complete description and the mathematical basis for the method will be presented in the body of this paper.

2. Background

2.1. Super resolution and the MRI

2.1.1. General

Enhancing the resolution of images is a well-studied field of research. Specifically, for about two decades scientists have been attempting to enhance the spatial resolution based on data collected from several low-resolution images taken from the same scene. This problem is called the *super-resolution problem*.

Super-resolution reconstruction The process of combining several low-resolution images to create a high-resolution image

Initial image resolution is based on the properties of the sensor. The sensor can vary from common cameras, satellites, SAAR radars, MR devices, etc. Each sensor has its own characteristics that affect the images it produces. In order to solve the SRR problem, we must first model the imaging process. We now describe the common model (see Refs. [1,5–7,9,10,13,14] and others) for the SRR problem. Given a set $\{Y_k\}_N^{k=1}$ of low-resolution images, one can write the following equation:

$$Y_k = D_k B_k G_k X + E_k, k = \{1...N\} \quad (2.1.1)$$

where Y_k is the kth low-resolution input image, D_k is the decimation operator for the kth image, B_k is the blur operator of the kth image, G_k is the geometric transformation operator for the kth image and E_k is the white additive noise.

The decimation operator, D_k, defines the "sampling rate" of the high-resolution scene used to create the low-resolution input images.

The blur operator, B_k, sometimes referred as the point spread function (PSF), is defined by the physical properties of the imaging device and differs from one image to the next. Usually, the blur is modeled as a low-pass filtering over the "real-world" scenario.

The geometric transformation, G_k, operator brings all the input images to the same base point (reference grid) so they could be combined. In cases where G_k is unknown, registration algorithms are applied to compute the geometric transformation of the image.

White additive noise, E_k, always exists due to the nature of the imaging device.

2.1.2. Simplifying the model

In this paper, we use a simplified model to solve the SRR problem. In Section 6, we present justification for using the simplified model based on the process done by the MR device. To compare the model to previous works, one can think of the following adjustments on the general model:

(1) Our analysis is based on three types of decimation operators. The input consists of low-resolution images, with pixels of dimension (also noted as pixel resolution): 3×3, 4×4 and 5×5. We will use these images to reconstruct a 1×1 high-resolution image.
(2) Based on Ref. [1], we can assume the PSF is space invariant and the same for all images (taken with the same resolution). Moreover, in Ref. [1] the PSF was modeled using box and Gaussian functions. Our basic model assumes box-type blur (uniformly). In Section 7.3, we discuss the result of applying other types of blur operators.
(3) In this paper, we assume we have exact knowledge as to the offset of each given image (based on the MRI device being used); thus, there is no need for registration in our model. In our experiment, subpixel shifts were implemented by changing the receiver oscillator frequency. However, see discussion in Section 2.2.3.
(4) We will show that in the absence of noise our algorithm gives perfect reconstruction of the desired

[1] It has the potential to further enhance spatial resolution beyond the practical limits of (1) increasing the readout and phase encoding steps, (2) decreasing the FOV and (3) increasing the gradient strength.

high-resolution image. We also discuss the effect of noise on the algorithm and demand certain properties to minimize the effect of noise on the result. One might also consider applying denoising filtering on the images; such filtering may be applied on the input LR images as preprocessing action or on the final HR image. Such techniques may give better results when applied over the final image due to SNR considerations.

2.2. Previous work

2.2.1. General super resolution

The super-resolution problem was extensively addressed in the literature during the last two decades. In this chapter, we will discuss the main reconstruction techniques that are in use. These methods can be classified as follows:

(1) Frequency domain techniques.
(2) Iterative algorithms.
(3) Statistical methods.

Frequency domain techniques are based on the generalized sampling theorem laid by Papoulis [16] and Yen [17]. The first SRR algorithm was suggested by Tsai and Huang [2]; they assumed nonblurred and nonnoisy images. The technique is based on utilizing aliasing effects of band-limited signals. Their work was followed by Kim et al. [3] using least-squares minimization on noisy and blurred images. Ur and Gross [4] proposed a spatial domain method where the high-resolution image was created using interpolation over the low-resolution images; they assumed known 2-D translations and uniform invariant blur. Most frequency domain methods are based on transforming the input images to the frequency domain [using 2-D discrete Fourier transform (DFT)], combining the spectral data and returning the output image (after applying 2-D IDFT). Because the images are reconstructed in the frequency domain, these methods have "recursive" properties, and new samples can be combined into the final reconstructed image as they arrive.

The main SRR iterative algorithm is iterative back projection (IBP) proposed by Irani and Peleg [5,6]. In this algorithm, the high-resolution output image, $\overline{X}$, is built iteratively to best describe the input sample images, Y_k. In every step of the algorithm they generate (using the common model, see Section 2.1.1) a set of low-resolution images (back projected), $Y_{\bar{k}}$. The back projection is done using the current best guess, $\overline{X}$, as the high-resolution scene. In each step, the algorithm refines the best guess $\overline{X}$ such that $Y_{\bar{k}}$ better describes Y_k.

Other iterative methods were suggested such as the projection onto convex sets (POCS) algorithm, proposed by Patti et al. [7]. The POCS algorithm resembles IBP and assumes known convex constraints on the solution so the iterative algorithm updates the current best guess according to these constraints. Elad and Feuer [8,9] presented a generalized model for the super-resolution problem and analyzed it under the maximum likelihood estimator (ML), maximum a posteriori probability estimator and POCS methods. The analysis assumes prior knowledge of the model operators (i.e., space-varying blur, geometric transformation, additive gaussian noise and measurements resolutions). Elad and Feuer [8] proposed a hybrid algorithm that combines the simple ML and the POCS methods. Later, Elad and Hel-Or [10] proposed a new method that separates the treatment to deblurring and measurement fusion to create an efficient SRR algorithm for the case of space-invariant blur and pure translation motions.

Statistical methods seek the high-resolution image with maximal probability to "create" the low-resolution input images (according to the imaging model). Such algorithms were presented by Cheeseman et al. [11] who used the Bayes' theorem and by Shekarforoush and Chellappa [12] who used the Markov random field to model and solve the problem.

2.2.2. Super-resolution in MRI

To the best of our knowledge, the idea of attaining subpixel resolution in MRI first appeared in 1997 (see Ref. [18]). The techniques mentioned in Ref. [18] include using various pixel shifts and varying the pixel sizes, using a variety of physical techniques. Thus, our paper should be viewed as giving a model and algorithm within the framework of Ref. [18]. Super resolution in MRI is a new field of study, Peled and Yeshurun [13] suggested an applicable SRR algorithm for MRI. They used spatially subpixel-shifted images of diffusion-weighted imaging and diffusion tensor imaging in vivo and created new images with improved resolution. Results are given after applying IBP [5] and state that optimal algorithm for MRI should be selected from the available ones based on comparison analysis. Shortly after the publication of the above paper, doubts were raised as to the possibilities of applying SRR for MRI. All images were taken with the same FOV, and the subpixel shifts were most likely[2] generated by means of a postprocessing step. The Fourier-encoded data given by the MR device are inherently band limited, thus seemingly eliminating possible SRR. The improvement in the image resolution was ascribed to the increase in the reconstructed image's SNR. Also noted was that super resolution is probably possible when non-Fourier encoding is used (i.e., conventional slice selection or line-scan imaging).

Similar corollaries as to applying SRR in 2-D MRI were given by Greenspan et al. [1] who suggested using SRR for 3-D MRI in the slice-select direction using several 2-D low-resolution images. They showed results that in-plain (2-D) SRR of Fourier-based MRI is not possible and that every result made can be replicated using interpolation via zero padding. In the slice selection, on the other hand, they

[2] Based on the response to the article that claimed that the shifts are usually done as a postprocessing step. The authors stated that at least one MR manufacturer guaranteed that this step was not done postprocessing.

showed that SRR is possible and present an algorithm, also based on the Irani and Peleg [5] IBP, to create high-resolution 3-D images.

2.2.3. Fundamental problems with SRR in MRI

We note the following problems with SRR in MRI:

- Shifted images attained by postprocessing steps cannot add new information, thus eliminating possible resolution enhancement [13].
- It seemingly follows from the generalized sampling theorem [16] that subpixel shifts via changing the receiver oscillator frequency do not add new information needed for SRR either.

The argument about image shifts via frequency change is not quite precise because of the finite sampling (see Section 6.2). However, infinite sampling is not possible, and longer sampling periods are limited by the relaxation rate of the nuclear and SNR considerations.

Nonetheless, it is strongly suggested by our work that physical shifts allow SRR in MRI with the limitation of the PSF function (see Sections 6.2 and 7.3). We remark that this point was not clear to us initially and that the MRI experiments make use of frequency shifts.

3. Modelling the problem

We introduce the following definitions and model:

(1) The *subject area* is a 1×1 rectilinearly aligned square whose bottom right corner is at the (x, y) origin.
(2) The *true image* is represented as a matrix of real values associated with squares of a rectilinearly aligned grid of arbitrarily high resolution. These values are called the *true high-resolution values*.
(3) A scan of the image at some arbitrary $m\times m$ *pixel resolution*, at offset (δ_x, δ_y), means that the (i, j) entry of the scanned image contains the sum of the values of all high-resolution grid squares enclosed within the rectilinearly aligned square of side length $1/m$ and whose bottom left corner is at the point $(\delta_x+i/m, \delta_y+j/m)$. We will say that the *physical dimensions* of the pixel are $1/m\times1/m$.
(4) Define the maximal resolution to be n, i.e., scans can only be performed at pixel resolutions $m\times m$, where $m\le n$.
(5) Define $1/\delta>n$ to be the *maximal offset resolution*. We can perform scans at offsets (δ_x, δ_y), where δ_x and δ_y are integral (not necessarily positive) multiples of δ.
(6) We will refer to pixels of an $m\times m$ pixel resolution scan, for $m=1/(c\delta)$, $c\in Z^+$, as being a scan of pixels of dimension $c\times c$. Note that this is not the same as the physical dimensions of the pixels, which is $1/m\times1/m$. The definition of dimensions $c\times c$ follows because we can consider every such pixel as having a value equal to the sum of c^2 component higher resolution pixels (of physical dimensions $\delta\times\delta$) which it overlays.

Our goal is to compute an image of the subject area with pixel resolution $1/\delta\times1/\delta$, although the maximal pixel resolution than can be measured during a scan is $n\times n$ and $n<1/\delta$.

The assumptions above require justification. One possible reason that the pixel resolution and the offset resolution are different could be that the underlying physical mechanisms underlying both types of resolution are entirely different.

In the context of MRI, it seems (as discussed in Sections 2 and 6.2) that the pixel resolution is limited by the maximal magnetic gradient one can impose and by the maximum feasible readout points and phase encoding steps as discussed in Section 1.3. The offset resolution can be determined with much higher resolution in (1) the main magnetic field (Z direction), (2) readout and (3) phase encoding directions by varying the frequencies of the transmitter phase, respectively. The technique described in this work was instigated by the desire to enhance ^{17}O spatial resolution; however, it is applicable to other nuclei.

Another case where this model may possibly be applicable is in the context of satellite imaging where the three-dimensional location in space of the satellite may be known with higher precision than the underlying resolution of the image. (In the satellite example, the lower resolution images required by the model could be obtained simply by changing the angle at which the camera points to the target).

It also follows from the underlying physical explanation that varying the gradient of the magnetic field intensity can be done continuously, justifying the $m\times m$, $m<n$, lower resolution scans mentioned in the model.

We introduce two error measures in this paper, and we seek to minimize both or find an appropriate tradeoff between them

(1) A reasonable assumption is that the errors are proportional to the size of the pixel. It therefore follows that when representing a high-resolution pixel as a linear combination of low-resolution pixels, the total area used in the linear combination (and the coefficients) is small.
(2) Additionally, there may be errors limited to some physical location (e.g., motion of the subject or vibration of the scanner), so one additional (but related) goal is that localized errors should have only localized effects.

3.1. Multiple offsets of a single-resolution scan

Consider the scan of pixels of dimensions 2×2 seen in Fig. 2 and at the top of Fig. 3. Imagine that the pattern continues throughout the infinite grid. Consider any offset of this scan by (δ_x, δ_y), where δ_x and δ_y are integral

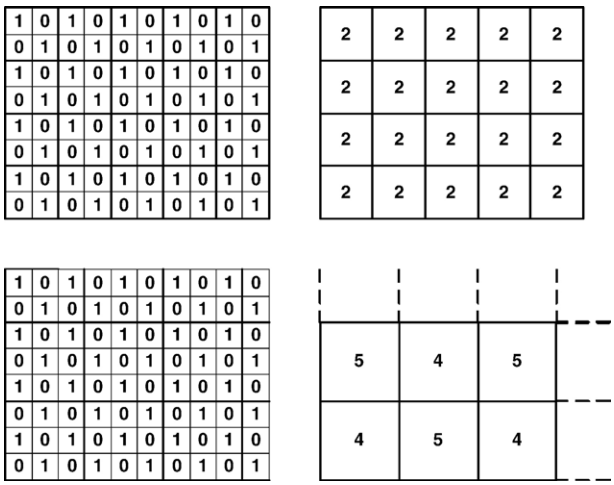

Fig. 2. Pixels of dimensions 2×2 and 3×3.

multiples of δ. The scan will assign all pixels a value of 2, irrespective of the offset (δ_x, δ_y).

Likewise, consider the bottom high-resolution image in Fig. 3. The pixel scan of this image also has all pixels with value 2, and so does any scan with pixels of dimensions 2×2 with offsets (δ_x, δ_y), δ_x and δ_y integral multiples of δ.

This means that there is no way we can distinguish between the top and bottom high-resolution images of Fig. 3 using pixels of dimensions 2×2 and any set of permissible (δ_x, δ_y) offsets. Similar examples can be constructed for pixels of arbitrary dimensions $c \times c$.

3.2. Making use of boundary value conditions

If we assume that all high-resolution pixels outside the subject area have a value of 0 (or some other known value), then we can make use of multiple scans with the same $m \times m$ pixel resolution to reconstruct the image.

Given that $m=1/(c\delta)$, $c \in Z^+$, we can perform c^2 scans at pixel resolution $m \times m$ [at offsets $(\delta_x, \delta_y) \in \{(i\delta, j\delta) | -1 \le i < c-1, -1 \le j < c-1\}$], from which one can write a set of linear equations where the variables are the values of

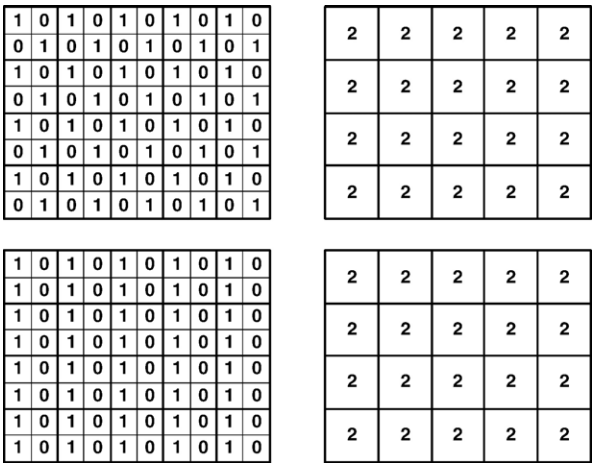

Fig. 3. Two high-resolution patterns that cannot be distinguished by any shift of a scan with pixels of dimension 2×2.

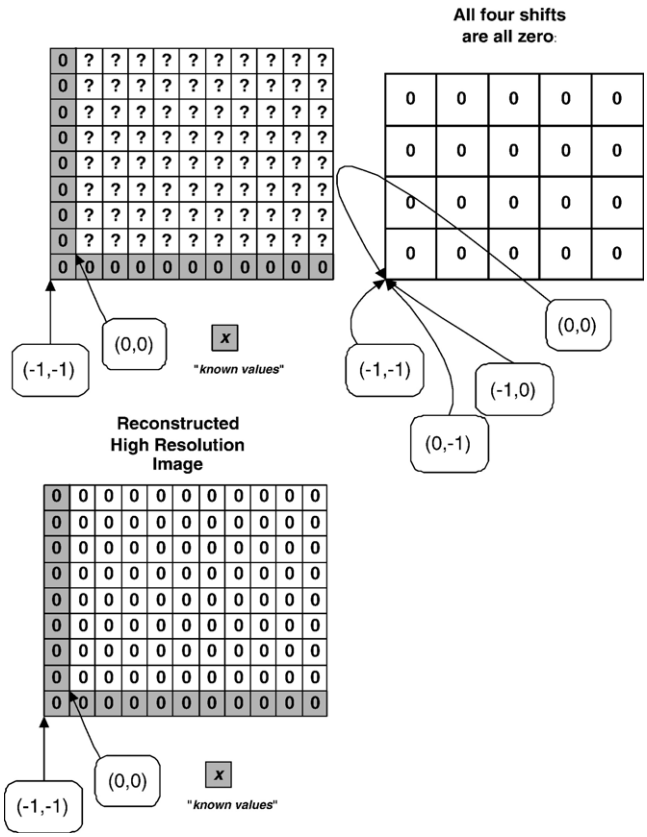

Fig. 4. Reconstructing the high-resolution image under assumptions on the boundary pixels. To simplify notation in this and subsequent figures, we have scaled everything so that offsets are integers rather than integral multiples of δ.

the underlying high-resolution pixels. We introduce a variable for every high-resolution pixel of physical dimensions $\delta \times \delta$ within the subject area. We have a total of $(1/\delta)^2$ linear equations, each of which contains up to c^2 variables. We also need to argue that these are indeed linearly independent, but this is easy to see — for example, one argument is that it is easy to perform Gaussian elimination for such a matrix. See Fig. 4 for a simple example of how four scans with pixels of dimensions 2×2 allow us to reconstruct the high-resolution image under the assumption that the boundary high-resolution pixels are known.

This solution is problematic for the following two reasons:

(1) It may be unclear that the "true value" of the high-resolution pixels outside the subject area is really zero.
(2) We have error propagation throughout the image. Any single error will propagate throughout the image. See Fig. 5 for an example of how a single error propagates throughout the image.

Given the assumption that we do know the values of the high-resolution pixels outside the study area, we could make use of additional information available by scanning the images. What we have not done above is to make use of the

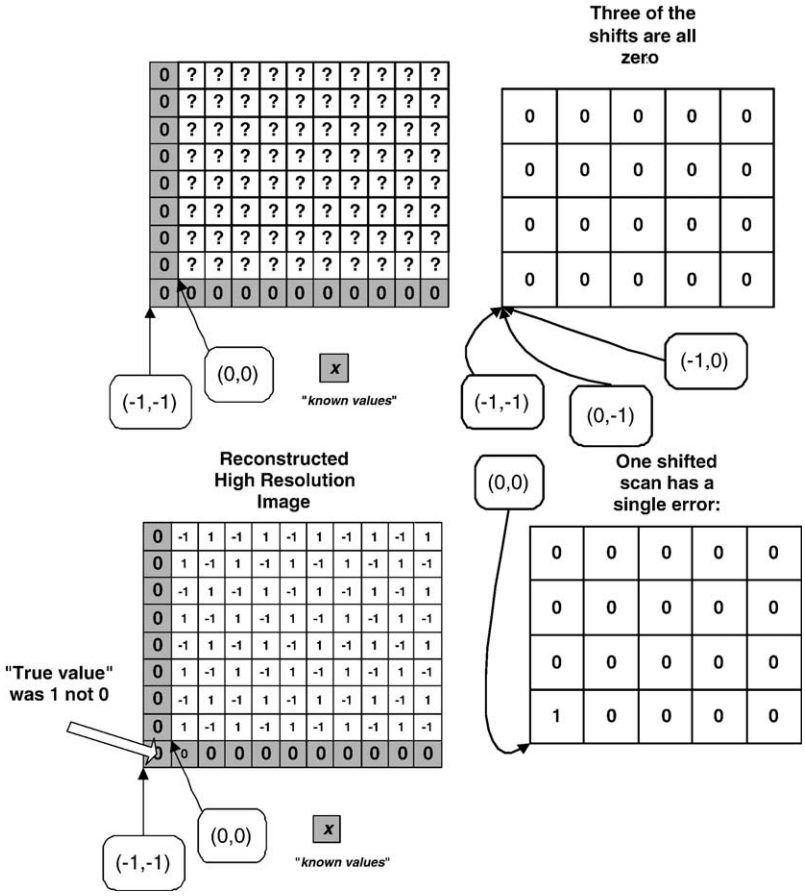

Fig. 5. A single error in one pixel in a single scan changes the reconstruction completely.

values of $\delta \times \delta$ physical dimension pixels that partially overlap the study area on the right and top (we have made use of those $\delta \times \delta$ physical dimension pixels that partially overlap the study area on the left and on the bottom). The potential advantage of using these additional pixels (and their associated linear equations) is that we can potentially reduce the errors arising from a single wrong high-resolution pixel.

This creates what is known as an *overdetermined system of equations* [19]. The linear equalities we have added clearly create linear dependencies between the rows of the matrix (as the number of rows is greater than the number of columns). One tool one can use to overcome the issues of overdetermined systems of equations is to use least-squares techniques (see Ref. [19]): minimize $\|Ax\text{-}b\|_2$, where $A \in R^{l \times k}$ with $l \geq k$ and $b \in R^l$.

Unfortunately, while it is true that a single error will be reduced in size, the least-squares solution to the overdetermined set of linear equations will not prevent the error from propagating throughout the image and the error will be reduced in size by no more than a constant factor (due to the nature of the linear dependencies created).[3]

[3] This is not entirely trivial to see, but if one tries out small examples this is easy to see.

What we seek therefore is a high-resolution image construction method that does not suffer from either of the problems above:

(1) We need not assume anything about the value of any underlying high-resolution image.
(2) Errors will not propagate throughout the image, but will remain localized in scope.

We will describe how to compute the value of a single high-resolution pixel as a function of spatially close low-resolution pixels (at various offsets). Again, as above, there will be several different linear combinations of different low-resolution pixels that will compute the high-resolution pixel. We can use the least-squares techniques to reduce the error terms here as well. Unlike the previous solution, error propagation will be localized.

4. The one-dimensional version of the problem

To simplify notation from this point on, rather than use offsets that are integral multiples of δ, we simply scale everything so that all offsets are integral. For clarity of exposition, we introduce a one-dimensional version of the reconstruction problem. We later build upon this

one-dimensional reconstruction to obtain two- and three-dimensional reconstructions. Consider pixels of dimensions $1\times x$ and $1\times y$ for positive integers x and y. Let $\gcd(x, y)$ denote the greatest common divisor of x and y. It follows from the extended Euclidean algorithm that there exist integer values a and b such that $ax+by=\gcd(x, y)$.

This means that if we are given all x possible (different) offsets of the $1\times x$ pixels and all y possible offsets of the $1\times y$ pixels, then we can compute the values of pixels of dimensions $1\times\gcd(x, y)$, at any offset. In particular, this implies that if $\gcd(x, y)=1$, then we can compute the actual high-resolution image.

Notation: let $v(z, i)$, $i\in Z$, denote the value of the $1\times z$ pixel at an offset of i from the origin. We are given the values of the form $v(x, y)$, $0\le j\le m-x$, and $v(y, j)$, $0\le j\le m-y$. As either $a>0$ and $b<0$ or $a<0$ and $b>0$, assume without loss of generality that $a>0$. To compute the $1\times\gcd(x, y)$ pixel at offset i compute

$$\sum_{j=0}^{a-1} v(x, i+xj) - \sum_{j=0}^{|b|-1} v(y, i+1+yj).$$

An example of this reconstruction is given in Fig. 6. There may be a problem in that the required pixels are missing, but note that the signs can be reversed and the pixel reconstructed by taking any combination of $1\times x$ rectangles and any combination of $1\times y$ rectangles so that their difference gives the required $1\times\gcd(x, y)$ result.

An alternative interpretation of the expression $ax+by=1$, $a>0$, $b<0$, is $ax=1 \bmod y$, i.e., add multiples of x and reduce modulo y. This leads to a reconstruction algorithm that is entirely localized in that errors that appear further away than $x+y$ high-resolution pixels away from the pixel

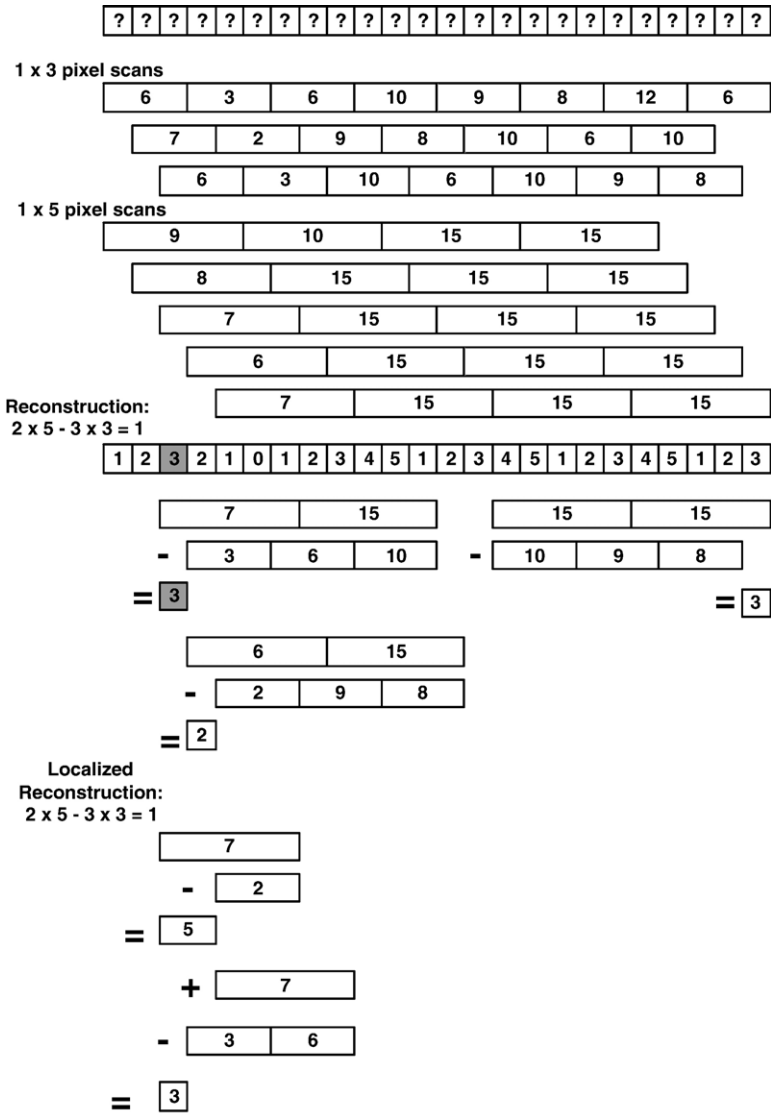

Fig. 6. One-dimensional reconstruction, pixels of dimensions 1×3 and 1×5.

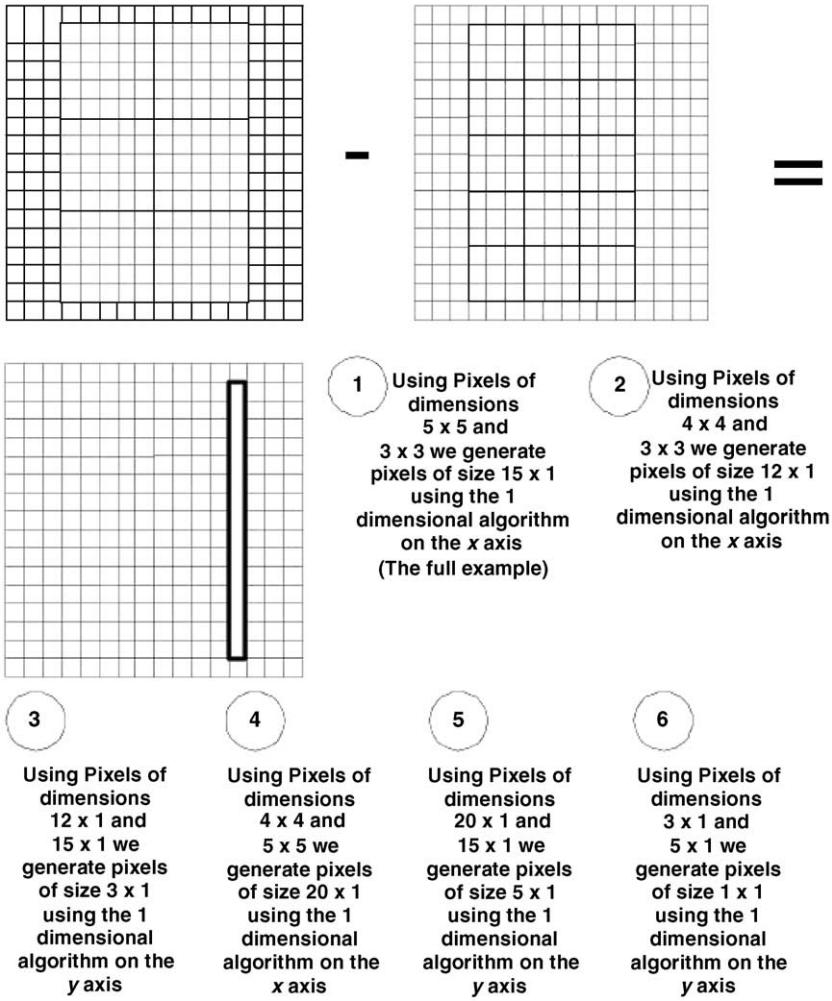

Fig. 7. Two-dimensional reconstruction, pixels of dimensions 3×3, 4×4 and 5×5.

being reconstructed do not influence the outcome at all. The localized reconstruction is also given in Fig. 6.[4] The same can be done in two or more dimensions.

A consequence of this algorithm is the following lemma:

Lemma 4.1. Let $p_i \equiv v(1, i)$ be the true value of the ith high-resolution pixel. Given a set of equations for high-resolution pixels derived from pixels of dimensions $1\times x$ and $1\times y$ at all possible offsets, where x and y are relatively prime, it is possible to reconstruct p_i locally, i.e., using only linear equations involving $p_i, \ldots, p_i+x+y-1$. The coefficients in the linear equations are limited in size by $\max\{x,y\}$.

5. Two and more dimensions

Given pixels of size $x\times x$, $y\times y$ and $z\times z$, where x, y and z are pairwise relatively prime, we can likewise reconstruct all the high-resolution pixels. This can be done while ensuring locality of error propagation.

The error propagation is limited to an area of $O(xyz)$ high-resolution pixels and that the linear combinations are all with small (constant) coefficients. That is, if the ratio of value to error is $o(1/xyz)$, then we can reconstruct meaningful data in the high-resolution reconstruction.

Additionally, we can compute the same high-resolution pixel using at least four nonintersecting sets of low-resolution images, giving some degree of error control.

Given pixels of dimensions $x\times x$ and $y\times y$, we can easily construct pixels of dimensions $x\times xy$ and $y\times xy$, simply by stacking pixels of appropriate sizes one atop the other.

Using the one-dimensional version of the problem, we can construct pixels of dimensions $1\times xy$. Similarly, this can be done for pixels of dimensions $1\times xz$ and $1\times yz$. Given pixels of dimensions $1\times xy$ and $1\times xz$, we can compute pixels of dimensions $1\times x$ as $\gcd(xy, xz)=x$. Similarly, we can compute pixels of dimensions $1\times y$ from the $1\times xy$ and $1\times zy$ pixels. From the $1\times x$ and $1\times y$ pixels, we can compute the underlying high-resolution pixels.

[4] There must be some way to do so for the higher dimensional versions as well (aside from the obvious — use the 1-D version where appropriate).

A worked out example using pixels of dimensions 3×3, 4×4 and 5×5 is given in Fig. 7.

This can be generalized to any number k dimensions using k+1 low-resolution pixels whose dimensions are relatively prime.

A consequence of the above algorithm is the following lemma:

Lemma 5.1. Let $p_{i,j}$ be the true value of the high-resolution pixel with index (i, j). Given a set of equations for high-resolution pixels derived from pixels of dimensions x×x, y×y and z×z at all possible offsets, where x, y and z are relatively prime, it is possible to reconstruct $p_{i,j}$ locally, i.e., using only equations involving

$$\begin{pmatrix} p_{i+d,j} & \cdots & p_{i+d,j+d} \\ \vdots & \cdots & \vdots \\ p_{i,j} & \cdots & p_{i,j+d} \end{pmatrix} \text{where } d = O(xyz).$$

The coefficients in the linear equations are limited in size by O(max{x, y, z}).

6. Experimental design

Experiments were performed using a GE clinical 1.5-T MRI scanner, Advantage, version 5.4 (GE Medical System, Milwaukee, WI). The scanner is equipped with gradient coils generating gradients with a maximum strength of 10^{-2} T/m.

Planar and longitudinal spatial resolution enhancement studies were carried out. In this paper, planar resolution of phantoms with well-defined regular and simple structure will be described.

The phantoms (shown in Fig. 8) consisted of plastic frames made of polycarbonate of thickness of 2.54 mm that served as spacers between sheets. Ten sheets of identical thickness were glued between 11 frames to form a cluster. Five clusters each containing 10 sheets of identical thickness (from left to right) 1.27, 1.01, 0.50, 0.375 and 0.63 mm were placed in a container. The 0.375-mm sheet was made of polycarbonate material, and the other sheets were made of amorphous polyester material. The size of the sheets was 50×50 mm. The size of the spacers was 53×53 mm. Each frame that served as a spacer between the sheets had a square hole in it of size of 20×20 mm. The square holes in the spacers as well as the container were filled with water. The length of the container was such as to provide a tight fit of the clusters to the container.

The container was placed horizontally in the x direction of the magnet. ^{1}H images of the container were obtained using GE head coil and spin echo pulse sequence, bandwidth of 32 kHz, echo time (the time period between the 90° pulse and the echo) of 20 ms. And the interval between consecutive pulse sequence, TR=1000 ms. Averages of four acquisitions (NEX=4) were obtained, resulting in acquisition time of 17 min and 8 s. Spatial resolution of 256×256 was selected.

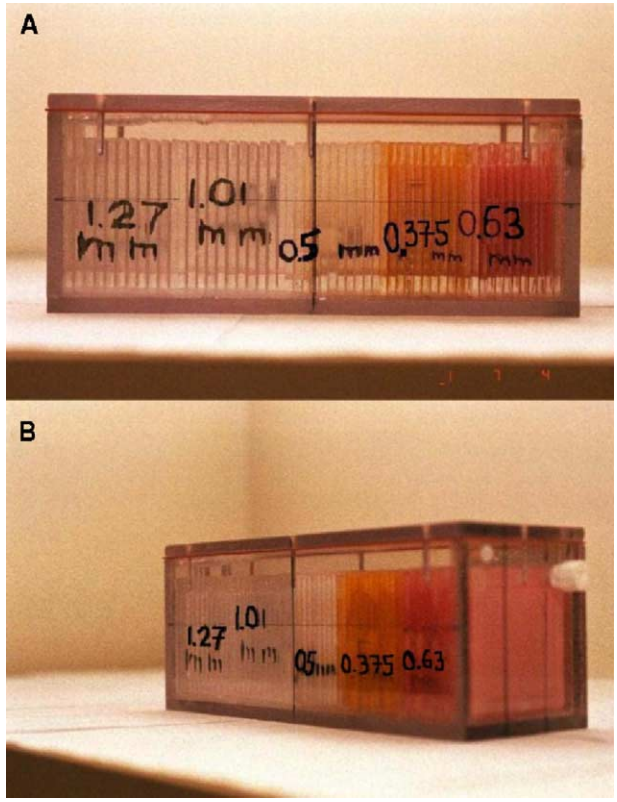

Fig. 8. Phantom image: (A) front view (in-plane), (B) side view. Spacers (sheets) width for each cluster is written on the image.

The selected mode was "rawdata"=0 and "autolock"=1. Those modes provide both the raw data that consist of the echo of free induction decay (FID) for NEX=4 and the images reconstructed by the GE software.

Axial slices were obtained at three FOVs: 230.4, 307.2 and 384 mm. Slice thickness was set at 3 mm. Only one slice was taken; however, the distance between fictitious slices was set at 10 mm (slice separation=10 mm). The position of the phantom was the same in all studies.

The center of the FOV was shifted in the right (R) and posterior (P) directions. Images of all combinations of steps of 0.3 mm were obtained. Three shifts in the R direction and three shifts in the P direction for FOV=230.4 mm, resulting in nine images. Four shifts in the R direction and four in the P direction for FOV=307.2 mm, resulting in 16 images. Five shifts in the R direction and five shifts in the P direction, resulting in 25 images.

To avoid updated scaling for each image, at the end of each scan the following sequence of commands was used: "cancel", followed by manual "prescan", followed by "scan". The FOV was modified through the modify "cv" option.

6.1. Input images correctness

The input images were divided into three FOVs corresponding to pixel sizes of 3×3, 4×4 and 5×5. To assure that the input images are correct, we created a mathematical model of the phantom. The designed model was based on the phantom known structure (see Section 6).

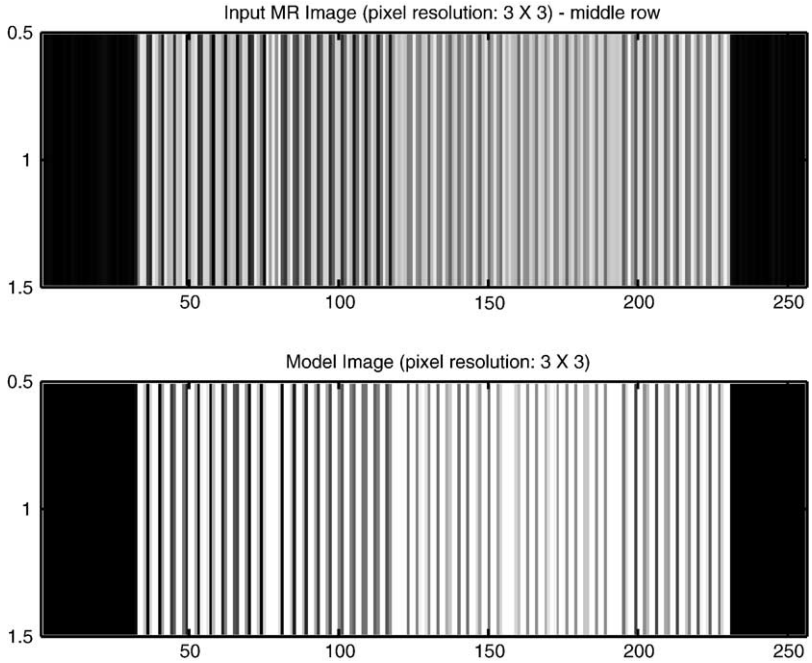

Fig. 9. Image reconstructed from MR scanner data at FOV=230.4 mm (top) vs. the mathematical model (bottom). Slice thickness: one voxel.

We first created a model of one row (the main parameters were the spacers width and the distance between spacers) in high-resolution (0.06×0.06 mm), we then applied the rectangular blur operator and decimated it to the desired resolution. The model images were compared to the MR images; the results are shown in Figs. 9–11.

Moreover, we checked for consistencies between the different FOVs by generating super images (or energy maps) for each pair of FOVs [at (0,0) offset] and compared them. The super images of FOVs $d1$ and $d2$ consisted of pixel resolution of $d1*d2$. The intensity of each such "energy pixel" is the sum of the underlying pixels of the original image. The images were compared and found to match each other (examples are given in Figs. 12–14). We note that when comparing the super images of two nonmatching pair of images (selecting two random offsets), the results were of inferior quality than that of matching pairs. This result raises a few questions as to the consistency of the input images. Feasible reasons for the inconsistency are discussed in Section 7.3.

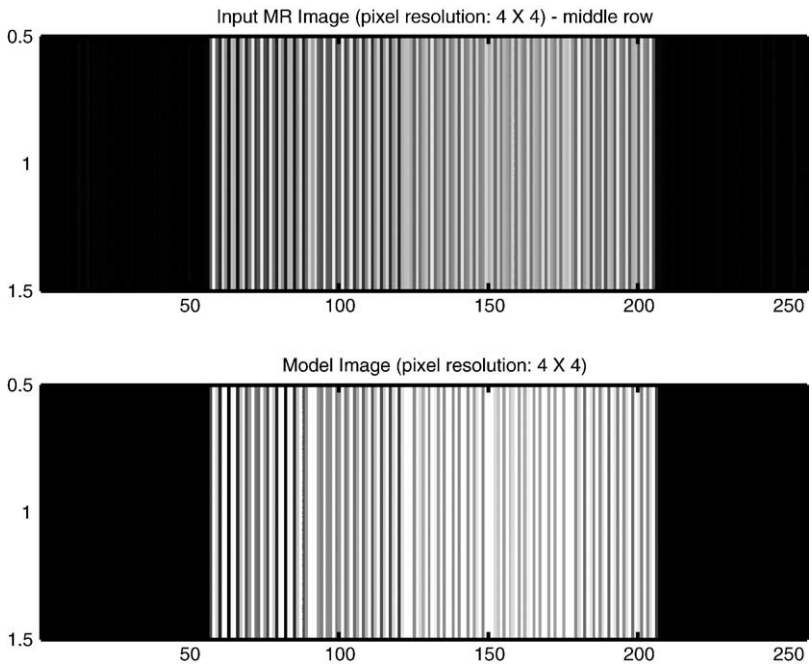

Fig. 10. Image reconstructed from MR scanner data at FOV=307.2 mm (top) vs. the mathematical model (bottom). Slice thickness: one voxel.

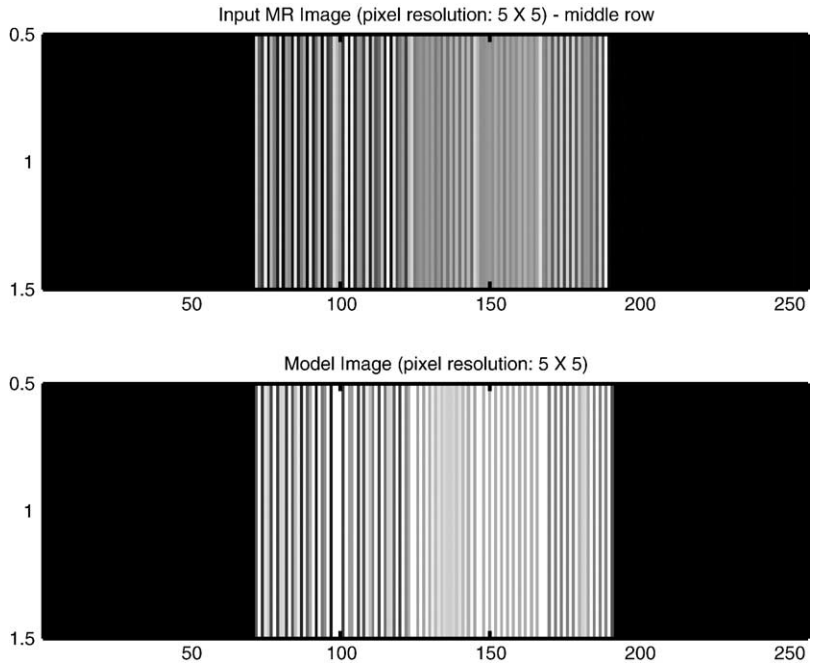

Fig. 11. Image reconstructed from MR scanner data at FOV=384 mm (top) vs. the mathematical model (bottom). Slice thickness: one voxel.

The images' SNR values were computed as the mean of a high-intensity region of interest divided by the standard deviation of a region of noise (see Ref. [1]). The measured SNRs were 45, 72 and 108 for FOVs 3, 4 and 5, respectively.

6.2. Bandwidth and resolution analysis

In Refs. [1,13], it was stated that SRR in MRI is not possible in the in-plane since the Fourier-based MRI is band limited and that subpixel shifts are done as a postprocessing step; thus, no new information is gathered. To address these issues, let us look at each claim and analyze it.

6.2.1. Acquisition of raw data

Raw data are always acquired in the time domain. An RF pulse is applied, and following the RF pulse the NMR signal (named: echo of the FID) is acquired as a function of time. By Fourier transform, the time dependence is converted to frequency dependence. The GE device acquires the signal as a function of time; we repeat the same measurement several times, in order to obtain an average of the signal and, consequently, enhance the SNR. Following the data acquisition, we set the instrument to provide the "raw" data (in the time domain, average of several pulses) and perform Fourier transform that provides the final product, the image (in the frequency domain).

6.2.2. Subpixel shifts

The *k*-space sampling locations are only determined by the gradient waveforms on the readout axis that are applied before and during the readout. The oscillator frequency change shifts the frequency of the acquired data, causing the phase of that data to change linearly during the readout and shifting the reconstructed object location by subpixel.

The shift is applied in *k*-space using frequency modulation in a way that cannot be replicated by a postprocessing mathematical manipulation as described in Ref. [13].

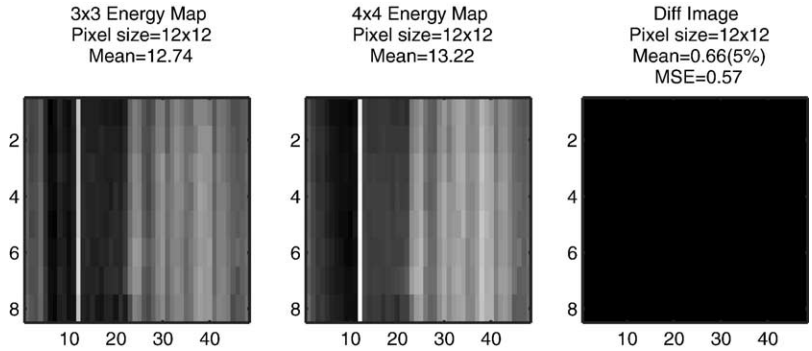

Fig. 12. Left image: Energy map of real MR image with FOV=3. Middle image: Energy map of real MR image with FOV=4. Right image: Energy map difference image; mean error is 5% of the original energy map mean intensity.

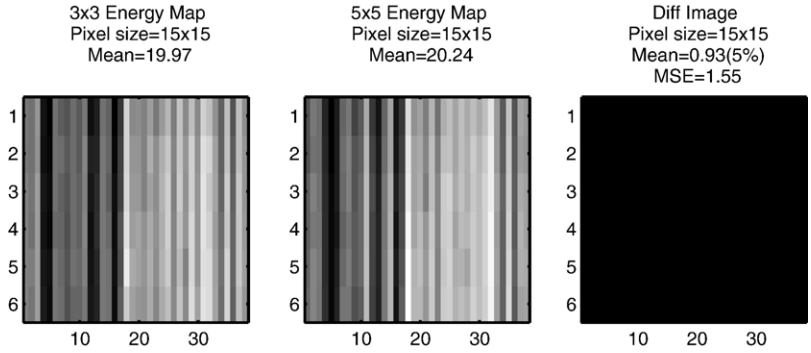

Fig. 13. Left image: Energy map of real MR image with FOV=3. Middle image: Energy map of real MR image with FOV=5. Right image: Energy map difference image; mean error is 5% of the original energy map mean intensity.

This is different than a postprocessing manipulation because the shift is applied before the band-limiting filter (The filter that happens before the A/D sampling and prevents aliasing; see Fig. 1.) The band-limiting filter "cuts off" the object, so applying the frequency shift before the filter shifts the object before it is cut off. Applying the frequency shift after the filter shifts the object after it is cut off. The latter is obviously different and is exactly equivalent to shifting the object with postprocessing.

We agree that postprocessing manipulation can only interpolate the signal but (as stated by the authors) cannot create new information.

6.2.3. k-Space limitations

k-Space is a commonly used presentation of MRI-acquired raw data. The raw data (in the time domain) have characteristics similar to those of a Fourier transform: $S(\mathbf{k})=\int_{\text{object}}\rho(\mathbf{r})e^{-i2\pi\mathbf{k}\cdot\mathbf{r}}d\mathbf{r}$. The meaning of this interpretation (as shown in Fig. 15) is that we collect the data in the "frequency" domain and use them to reconstruct the original signal. Let us analyze the characteristics of the two signals (*k*-space and image space): the image-space signal, $S(x)$, is finite (its dimension is the FOV); thus, its spectrum is infinite. The MR device acquires samples of this spectrum (named the *k*-space). *k*-Space signal is an infinite, band-limited signal: $K(t)$, $-\infty<t<+\infty$; we acquire N samples of $K(t)$ sampled at a rate of $F_s>2\cdot$Nyquist frequency from $0<t<T_{\text{scan}}$. Since we do not have samples of $K(t)$ from $t>T_{\text{scan}}$ and $t<0$, we cannot fully reconstruct the original signal $S(x)$.

A generalization of the Nyquist theorem states that under some general conditions for perfect reconstruction, it is enough to sample the signal at a rate that is "on the average" higher than the Nyquist rate. Namely, if we can sample the signal outside the interval at any nonzero rate, we would still be able to compensate by oversampling within the interval, such that $\lim_{T\to\text{inf}}\ 2T/N\le T_{\text{Nyquist}}$. However, since no samples were acquired outside the above time interval, the signal cannot be reconstructed from any numerable (even infinite) number of samples within the interval. It is interesting to remark, in this context, that the generalization also implies that if we knew the signal over a continuous interval (no matter how small that interval is), we would be able to reconstruct the entire signal (as long as it is band limited), since in this case we have an innumerably infinite number of samples, and the "average" sample rate in this case is infinite.

In the case of MRI, we do not have any samples outside the interval and cannot obtain innumerably infinite number of samples inside the interval; thus, perfect reconstruction is not possible. But still, increasing the number of samples inside the interval can (theoretically) improve the accuracy of the reconstruction, but the improvement is usually not overwhelming.

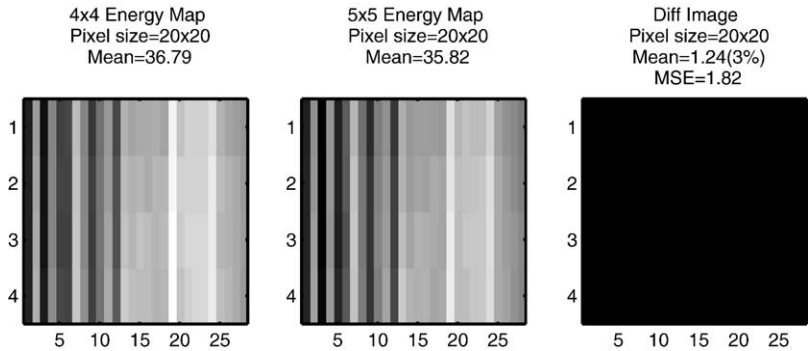

Fig. 14. Left image: Energy map of real MR image with FOV=4. Middle image: energy map of real MR image with FOV=5. Right image: Energy map difference image; mean error is 3% of the original energy map mean intensity.

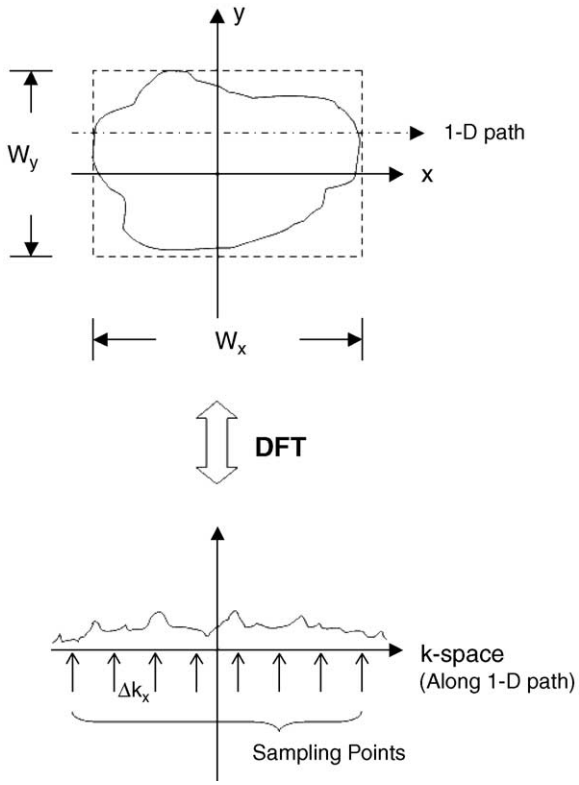

Fig. 15. Top: An object bounded by a rectangle of width W_x and W_y (FOV). The signal is finite; thus, its spectrum is infinite. Bottom: Sampling points of k-space along a 1-D path (one row of the image). k-Space is an infinite band-limited signal.

6.2.4. Pixel resolution

Analyzing the data process (see Fig. 1), we see that the input bandwidth is limited to 128 kHz due to the anti-aliasing filter. This bandwidth is sufficient since the effective output images' bandwidth is limited to 32 kHz. Thus, the resolution of each single image is limited by the Fourier pixel size [Δx_F] related to the DFT PSF [20]. Fourier pixel size is given by $\Delta x_f = 1/(N\Delta k)$, where $N=256$ points. This resolution changes when the FOV is changed according to $\Delta x_f = \mathrm{FOV}/N$. The algorithm presented in the next section models Δx_F, showing that combining several images with different resolutions at different subpixel shifts allows us to fully reconstruct a high-resolution image.

Scanner's frequency resolution is determined by the bandwidth of the exciter/receiver local oscillator and the register resolution. For GE scanner, the local oscillator is 10 MHz and the register resolution is 24 bits. The frequency resolution is therefore $10^7/2^{24}=0.596$ Hz. The local oscillator frequency must be changed in integer multiples of this frequency.

Field of view can be expressed in two ways: spatial units or frequency units. The FOVs we used were 384, 307.2 and 230.4 mm (also referred to as 5, 4 and 3). The FOVs correspond to the bandwidth, which is ±16,000 Hz= 32,000 Hz. In frequency units, the resolution is the same for all three FOVs. We used a resolution of 256×256, which means that the pixel planar resolution is 32,000/256= 125 Hz. In the current study, which serves as an example, we demonstrate a threefold planar resolution enhancement that corresponds to 125 Hz/3=41.7 Hz. The scanner's frequency step of 0.596 Hz is more than adequate to properly describe and achieve the desired resolution.

The above corollaries should suffice to justify the use of the resolution enhancement model and algorithm in the case of MR images. To support the above statements, we applied the resolution enhancement algorithm on the acquired phantom images.

In Fig. 16, we see an original low-resolution image (FOV=230.4 mm) and its power spectrum (top row). The high-resolution reconstructed image (times larger) is shown in the middle row, and on the bottom row we see a high-resolution image created using bi-cubic low-pass interpolation. As expected, the bilinear low-pass interpolation spectrum did not contain high frequencies (due to the low-pass property of the interpolation). Also, it is clearly observed that the power spectrum of the high-resolution reconstructed image contains high frequencies created in the resolution enhancement process.

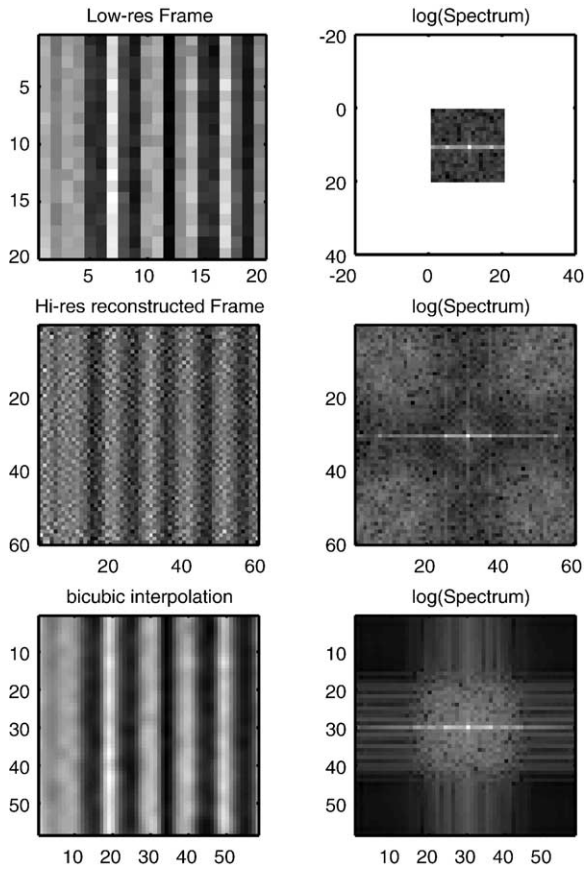

Fig. 16. Top row: Original image (pixel resolution 3×3) and its power spectrum. Middle row: High-resolution image generated by the algorithm and its power spectrum. Lower row: High-resolution image generated using bi-cubic low-pass interpolation.

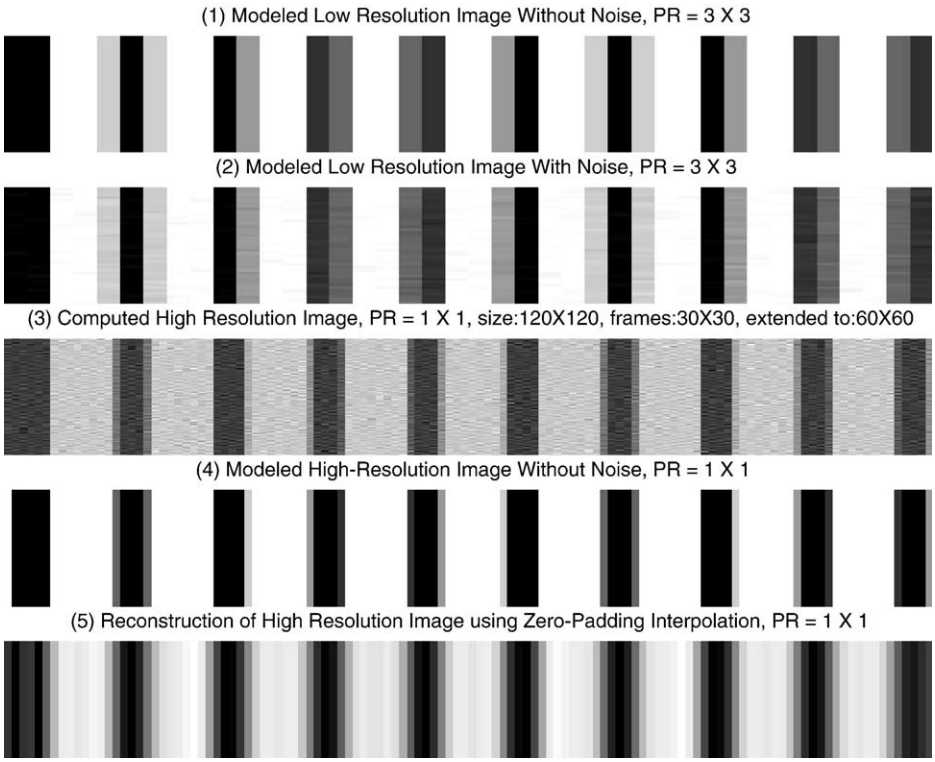

Fig. 17. Cluster 1: spacers width: 2.54 mm; distance between spacers: 1.27 mm. Image description (ordered top to bottom, same for Figs. 18–21): (1) modeled low-resolution image without noise, FOV=230.4 mm (3×3); (2) modeled low-resolution image with noise (SNR=45), FOV=230.4 mm (3×3); (3) high-resolution algorithm result (1×1); (4) modeled high-resolution image without noise (1×1); (5) reconstruction of high-resolution image using zero-padding interpolation (1×1).

7. Experimental results

The resolution enhancement algorithm was implemented by a Matlab program. The implemented algorithm differs from the one described in the previous sections for simplicity reasons. In the program, we directly used the least-squares method to derive the high-resolution image. The subject area was divided into small frames (30×30 pixels). Each frame was reconstructed independently by computing a larger high-resolution frame

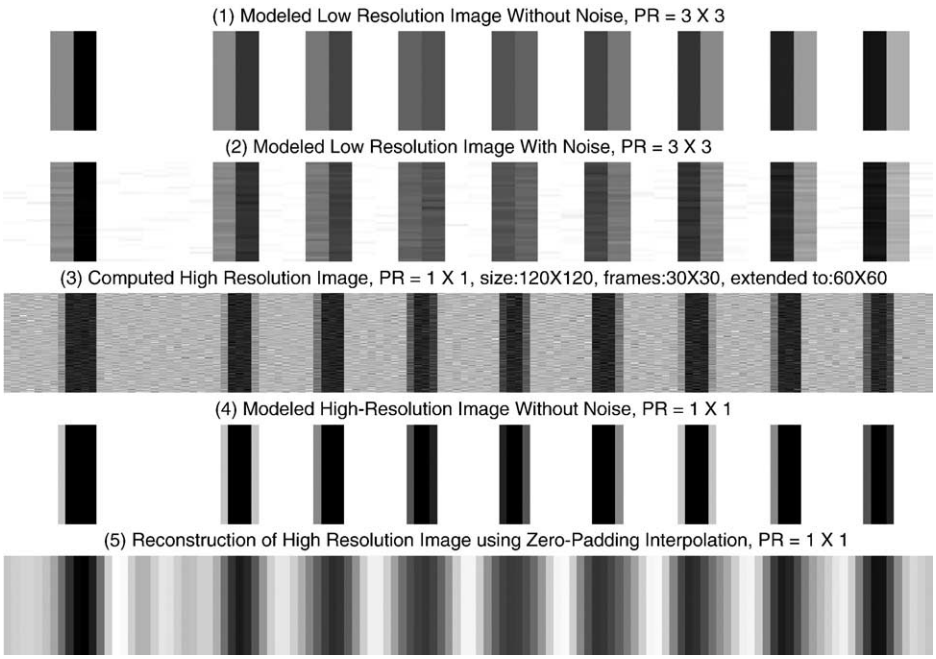

Fig. 18. Cluster 2: spacers width: 2.54 mm; distance between spacers: 1.01 mm.

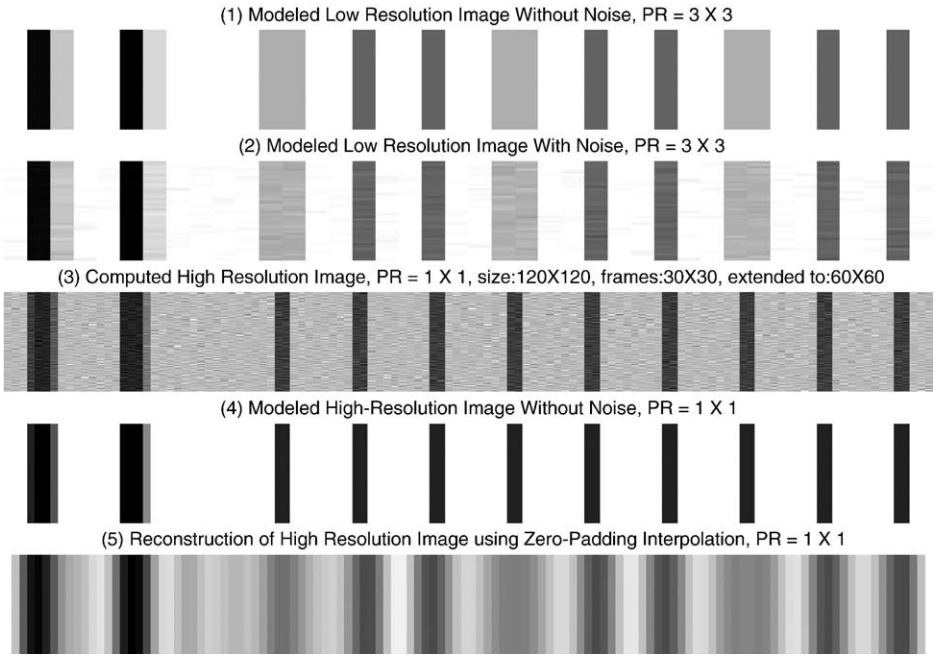

Fig. 19. Cluster 3: spacers width: 2.54 mm; distance between spacers: 0.5 mm.

(60×60 pixels), taking only the inner 30×30 pixels for the final image.

There is no known method to estimate the improvement of image quality by enhancing the image resolution. While some implementations would be interested in sharper edges in the image, others may wish to improve SNR at all cost. In the case of phantom studies, we suggested two measures of image quality: (1) sheet thickness (based on the phantom structure they should be equal) and (2) distance between spacers (based on the phantom structure they should be equal within each cluster and vary for each cluster; see Section 6 for details).

7.1. Model results

Experiments were made with a dataset generated by the mathematical model. The "raw" images were generated according to the phantom's specifications; after that they were blurred and sampled to the desired resolution and FOV.

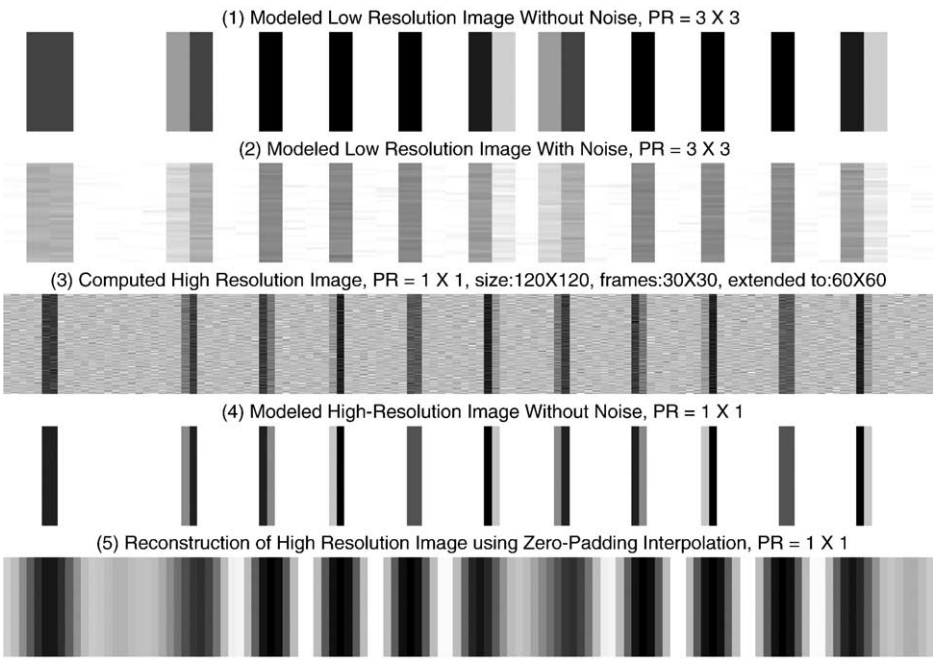

Fig. 20. Cluster 4: spacers width: 2.54 mm; distance between spacers: 0.375 mm.

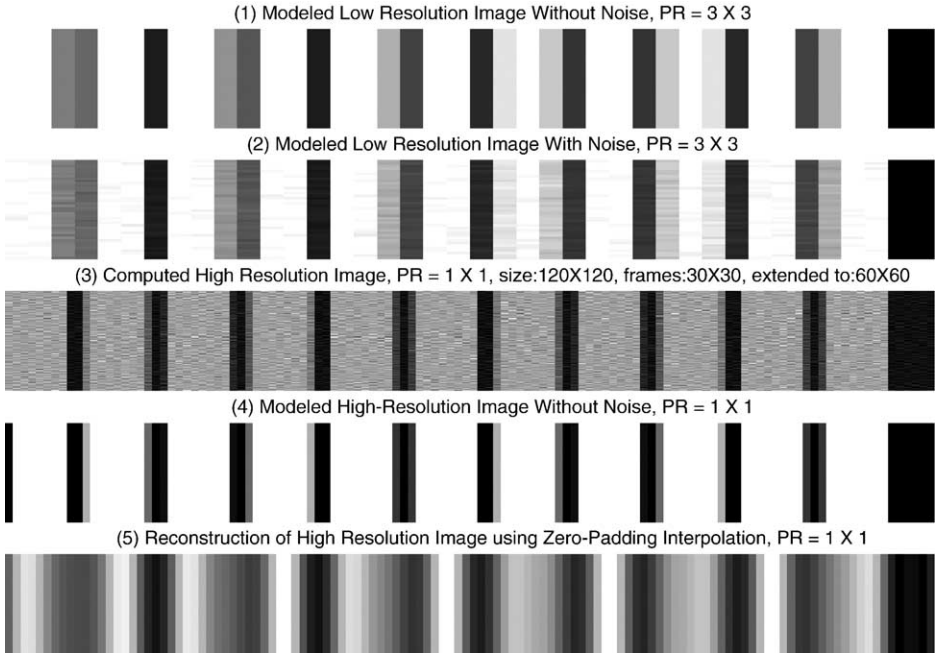

Fig. 21. Cluster 5: spacers width: 2.54 mm; distance between spacers: 0.63 mm.

White noise was also added to match the previously computed SNR in each FOV (see Section 6.1).

We present the resolution enhancement results for each cluster in Figs. 17–21; the high-resolution images are presented next to four reference images for comparison:

(1) Modeled low-resolution image (3×3) without noise.
(2) Modeled low-resolution image with noise (SNR= 45 — same as measured in the real MR image).
(3) Modeled high-resolution image without noise (see Section 6.1).
(4) Reconstruction of high-resolution image using zero-padding interpolation on the modeled low-resolution image without noise.

The high-resolution images reconstructed by the algorithm match the high-resolution image generated by the model (the model represents the best possible solution).

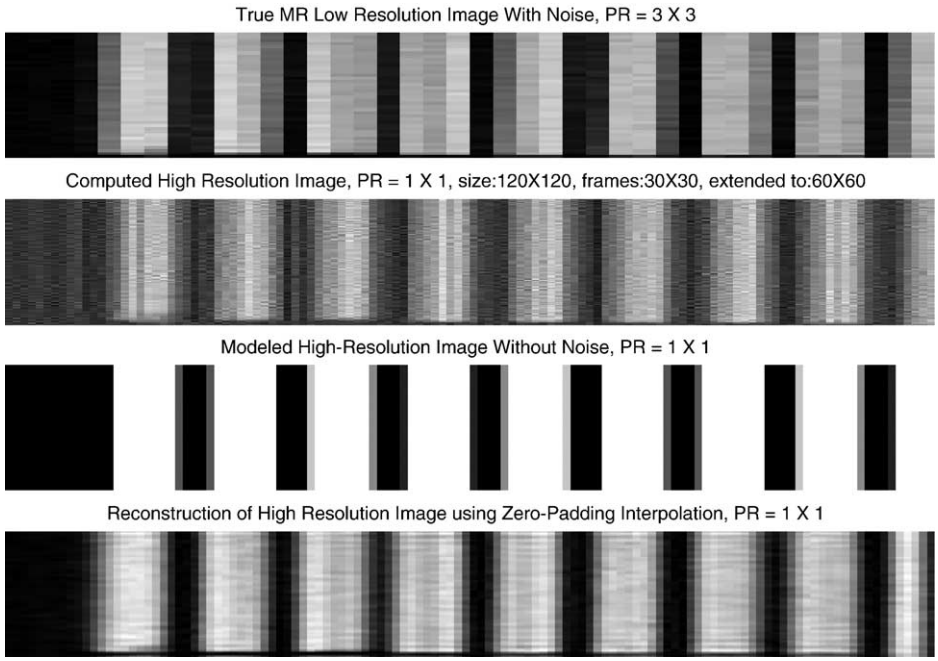

Fig. 22. Cluster 1: spacers width: 2.54 mm; distance between spacers: 1.27 mm. Image description (same for Figs. 23–26): top: low-resolution MR image, FOV=230.4 mm (3×3); middle up: high-resolution algorithm result (1×1); middle down: modeled high-resolution image (1×1); bottom: high-resolution image using zero-padding interpolation (1×1).

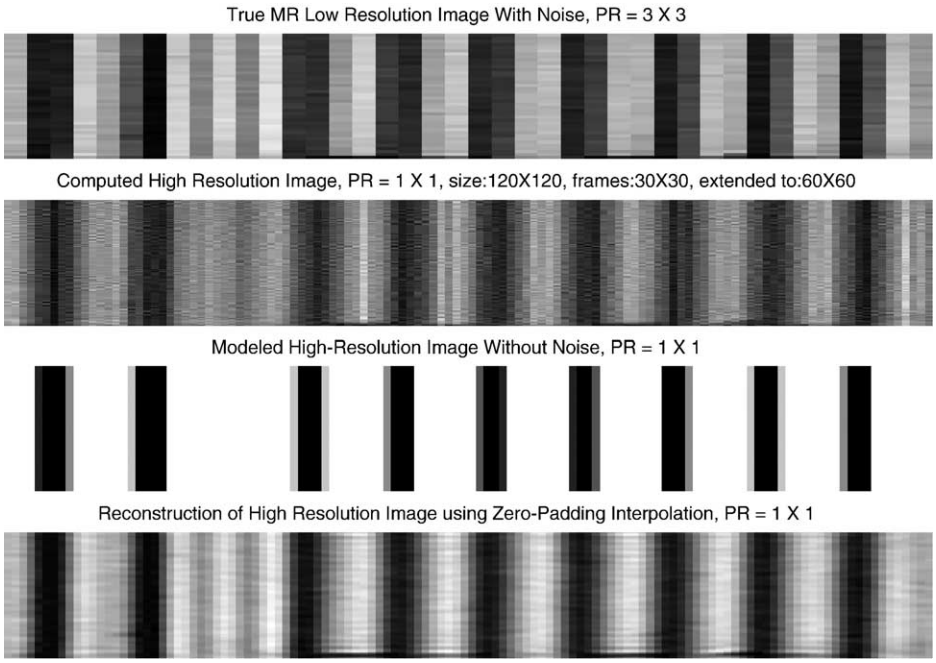

Fig. 23. Cluster 2: spacers width: 2.54 mm; distance between spacers: 1.01 mm.

Clearly, the algorithm results are superior to those of the zero-padding interpolation, both in line width and in the distance between lines.

7.2. Phantom results

High-resolution images of the real phantom images (created with GE MR device) were generated using the resolution enhancement algorithm.

Results are presented for each cluster in Figs. 22–26; the high-resolution images are presented next to three reference images for comparison: (1) true MR low-resolution image (3×3); (2) modeled high-resolution image (1×1) (see Section 6.1); and (3) high-resolution image generated using zero-padding interpolation.

The results of the algorithm on the MRI data were nondecisive. The algorithm results clearly show resolution enhancement. Nevertheless, no clear advantage to the algorithm reconstruction is observed over the zero-padding interpolation as observed on the model results. We view this as a major point for further study.

7.3. Discussion

The differences between the results of the modeled data and the results of the MRI data can be explained by the following.

Blur. The model assumes a rectangular (uniform) blur, while the true blur is unknown.[5] In Ref. [1], it was shown that using a Gaussian-shaped blur SRR algorithm gives better results. This result conforms with the following corollaries:

- The raw data (in *k*-space) is truncated by the MR device.
- Truncation means multiplying by a rectangular function.
- The image is convoluted with the Fourier transform of this rectangular function, resulting in a sync PSF function.

Experiments were carried out with other blur functions (Gaussian and sync functions taking only the central part of the sync with an additional lobe). We note that we observed a resolution enhancement using the box PSF and that no other function seemed superior over it.

Based on the model assumptions, each low-resolution pixel had an area of influence determined by the image resolution, d. The blur function was applied over $d \times d$ high-resolution pixels within this area. The resolution (and thus the area of influence) is defined as the width of the main lobe (or the 3-dB point). In reality, pixels outside this area also affect the intensity of the pixel, but they are not modeled.

Modifying the model in a way that would include the effects of pixels that are outside of the main lobe is possible. Experiments were carried out with different sizes of blur function (extending the area of influence beyond the 3-dB point), but no improvement in image resolution was observed.

[5] The true blur can be calculated or simulated by simply taking the 2-D Fourier transform of the window function that multiplies the *k*-space data. For simplicity, we can assume the *k*-space window is a 2-D "top-hat" function with constant radius. The resulting blur would be proportional to a J1 Bessel function divided by *r*.

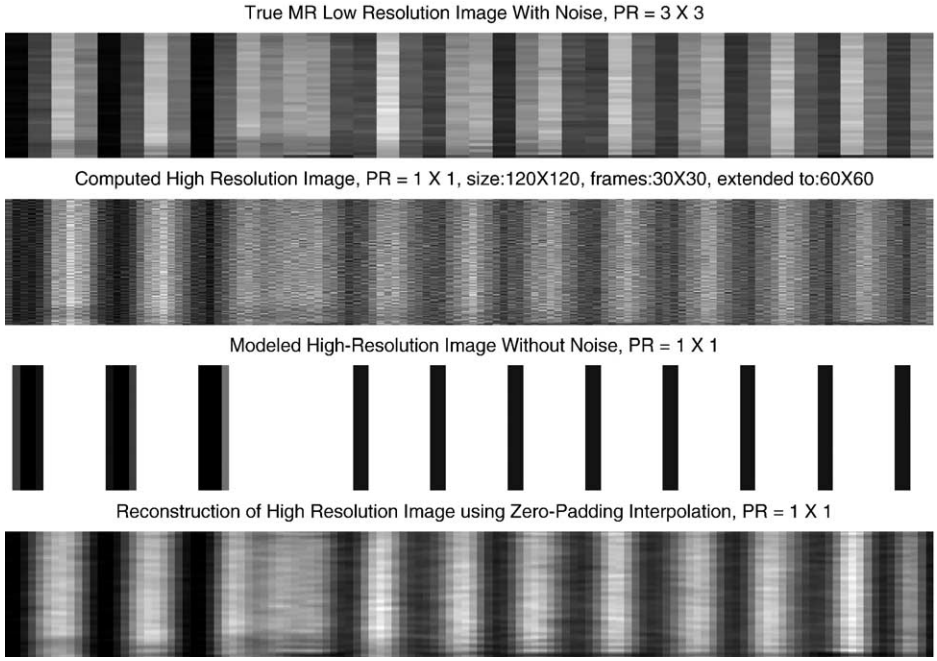

Fig. 24. Cluster 3: spacers width: 2.54 mm; distance between spacers: 0.5 mm.

Using nonbox blur function or extending the pixel's area of influence does not conform with the theoretical analysis (see Section 4) that shows that full reconstruction always exists since some submatrices may become singular.

The rectangular blur assumption is correct for the slice direction when selective excitation is used (i.e., 2-D scans). This hangs together with the results in Ref. [1] where super resolution worked (in the z direction) and where the assumption of a rectangular PSF is more nearly correct.

A 2-D-imaging case that more nearly satisfies the assumption of rectangular blur (at least in one direction) is line-scan imaging (for example, see Ref. [21]). This method is used for diffusion imaging because of its relative insensitivity to patient motion on some scanners or on some anatomy (such as spines) where single-shot EPI does not work well.

Line-scan imaging might work better (in one direction) than standard imaging with our resolution enhancement technique since there is no Fourier encoding in the y

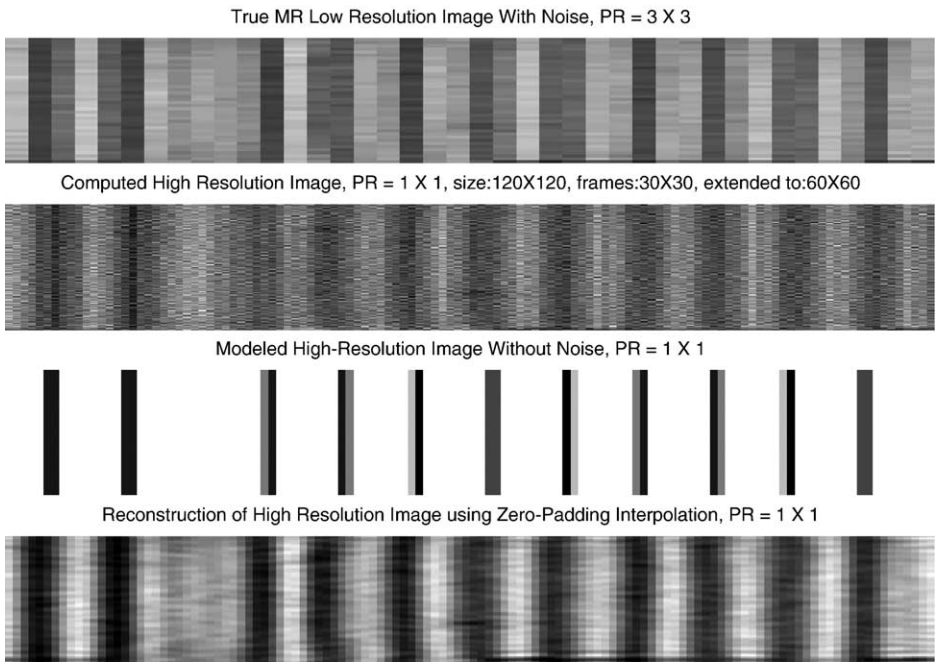

Fig. 25. Cluster 4: spacers width: 2.54 mm; distance between spacers: 0.375 mm.

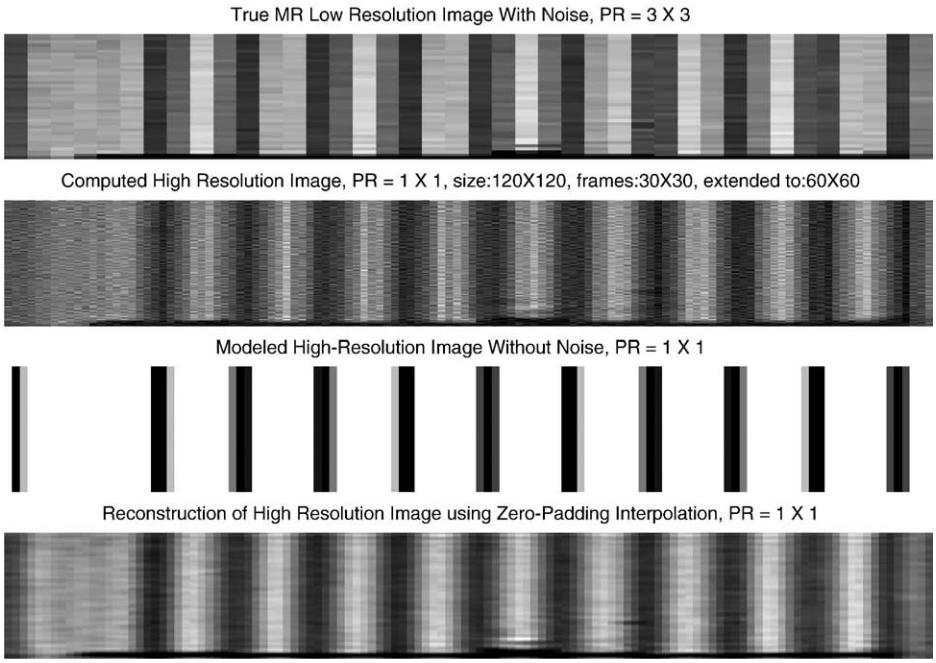

Fig. 26. Cluster 5: spacers width: 2.54 mm; distance between spacers: 0.63 mm.

direction. The PSF for the y direction is more nearly a rectangle (depending on the RF pulses), and so our method is expected to produce good results.

Subpixel shifts via frequency shifts. As discussed in Section 6.2, the use of subpixel shifts in the context of MRI can theoretically improve the reconstruction approximation. These methods usually have serious implementation implications, which may explain the difference in the results of the algorithm between the experimental and the modeled images. Better resolution may be achieved in the future by the use of physical shifts of the subject area (instead of changing the receiver oscillator frequency). We predict the results would be more similar to the ones of the modeled images. To perform the required physical shift, special hardware must be built.

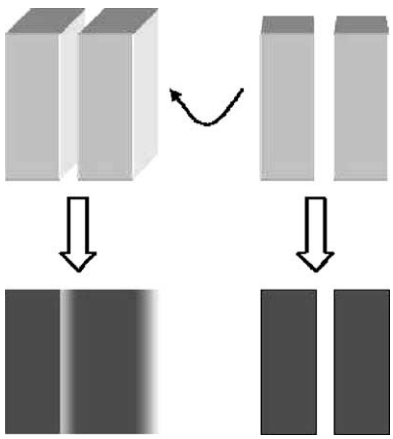

Fig. 27. Orientation example. Top row: 3-D phantom image — front view (on the left) and side view (on the right). Bottom row: 2-D image intensities as seen by the MRI.

Phantom orientation. In order to receive an accurate front-view image of the phantom (see Ref. [8]), the phantom should have been placed in a specific orientation inside the MRI device. Since the phantom was placed manually, such accuracy is very difficult to achieve. The effect of imperfect orientation with x, y and z directions on the final image is that sharp transitions from low-intensity regions (sheets) and high-intensity regions (regions containing water) become a continuous transition. Visual demonstration of the effect is presented in Fig. 27.

Homogeneity of the phantom. Unfortunately, the constructed phantom did not maintain the thin plastic sheets in a completely flat position. On the other hand, the model is based on the assumption that the sheets are completely flat. The model also assumes that the distribution of water in the spacer cavities is homogenous, whereas in reality there may be some irregularities in the water distribution in the cavities.

Imperfect magnetic field. Inhomogeneity of the static magnetic field (z direction) and nonlinearity of magnetic field gradients are two reasons that could also affect the resulting image. Those properties usually vary from one device to another and are included in the designer specifications. Such problems can cause the low-resolution MR images to not reflect the exact phantom, but the probability of this factor is not high.

8. Conclusions and open problems

We have given a mathematical model for the SRR problem and associated algorithms with highly useful properties. We note that our construction gives an improve-

ment factor of 3, but the generality of the method could give any improvement factor, e.g., had we used low-resolution pixels with dimensions 17×17, 19×19, 20×20, the improvement factor would have been 17. The results seem unsatisfactory. This may be due to the use of a rectangular PSF in the data processing and/or the use of frequency change for subpixel displacement of the images.

There are several open problems that arise in this setting:

(1) Most interesting is that we seem to have a new type of error control mechanism. While we can localize the reconstruction operation, it may actually make sense to use a variety of localized and nonlocalized reconstructions. The use of least-squares techniques for a subset of reconstructions is always possible. An alternative formulation of these arguments is that if we actually write down all the linear equalities, then many submatrices are nonsingular and lead to (partial) solutions of the variables.

(2) Optimization problems: a good optimization problem is "what is the smallest number of scans we can do to reconstruct the high resolution image?" Because of the redundancy in the equations (as seen from the fact that many alternative reconstructions are possible), it may be possible to reduce the number of scans significantly.

Scan selection problem. We now present a formulation of the above problem for the one-dimensional case, along with a motivation that such optimization is feasible.

Definitions:

Variables $\{x_i\}$; $1 \le i \le n$ represent high-resolution pixels

Scan $S_{i,j}=\{S^l_{i,j}\}$; $1 \le l \lfloor n/i \rfloor$ represent a collection of low-resolution pixels with dimension $1\times i$ and initial offset j; $0 \le j \le i-1$. Where $S^l_{i,j} = \sum_{p=i\cdot(l-1)+j+1}^{i\cdot l+j} (x_p)$.

Cover Two scans, $S_{i,j}$, $S_{k,t}$, where $i \neq k$ relatively prime, "cover" a set V of variables if we can reconstruct V using the 1-D algorithm (see Fig. 28).

Cover degree C_i defined as the number of scan pairs ($S_{i,j}$, $S_{k,t}$) such that each pair "covers" x_i. $C_i=0$ means that we cannot reconstruct x_i based on Lemma 4.1.

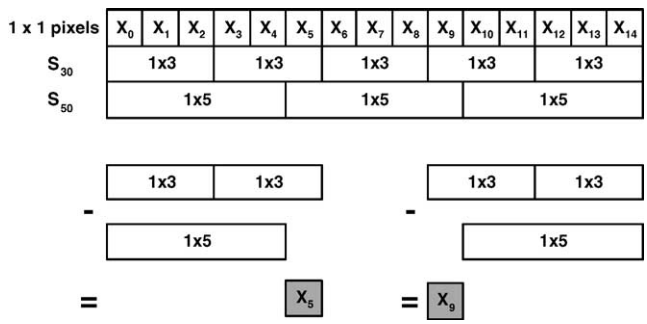

Fig. 28. Example of how two scans, S_{30} and S_{50}, cover the variables X_4 and X_9.

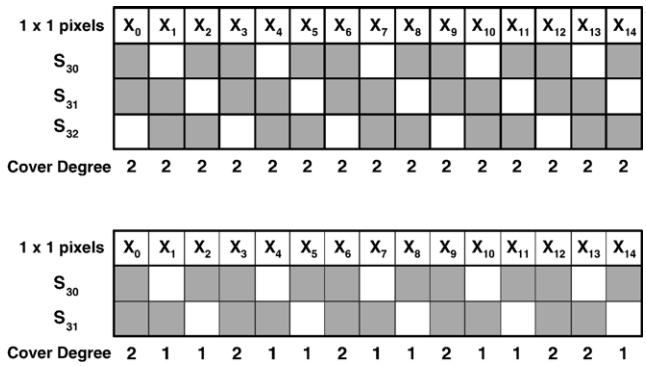

Fig. 29. Upper part: Each gray square in line i, $1 \le i \le 3$, represents a variable being covered by S_{3i} and all S_{5j}, $1 \le j \le 5$. Every variable is covered by exactly two pairs of scans. Lower part: The same without using scan S_{32}; all variables are still covered.

Goal: find G scans (offline or online) such that

(a) $|G|$ is minimal.

(b) $\forall i \cdot C_i \ge 1$ (this ensures we have enough equations to solve each variable using the 1-D algorithm).

Greedy algorithm:

1: **if** $\forall i \cdot C_i \ge 1$ **then**
2: Stop
3: **else**
4: Select the scan $S_{i,j}$ (from the unselected scans) that covers the maximal number of uncovered variables.
5: **end if**

Refinement: if all unselected scans $S_{i,j}$ at one stage contribute the same, select the one with the lowest i (it has a chance to cover more variables in the future).

Fig. 29 shows that in the one-dimensional case with resolutions of 1×5 and 1×3, one 1×3 scan can be removed (does not matter which) and still all variables are "covered"

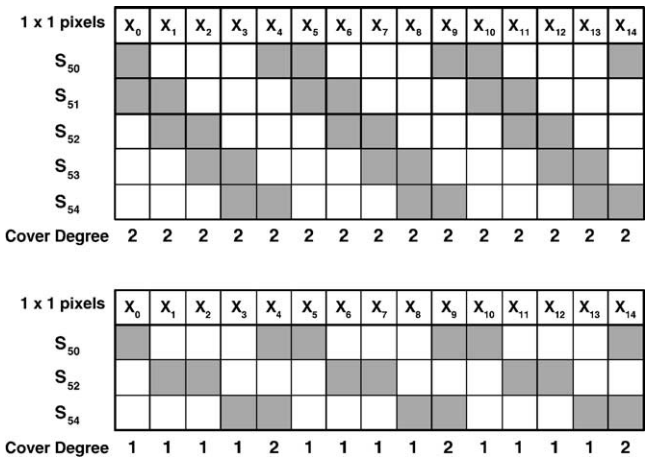

Fig. 30. Upper part: Each gray square in line i, $1 \le i \le 5$, represents a variable being covered by S_{5i} and all S_{3j}, $1 \le j \le 3$. Every variable is covered by exactly two pairs of scans. Lower part: The same without using scans S_{51} and S_{53}; all variables are still covered.

("cover" degree ≥ 1). The same holds for removing two 1×5 scans (as shown in Fig. 30).

Open questions:

(a) What/how many variables we wish to "cover"? or, how do we measure and decide that our result is good?
(b) What is the contribution (if any) to the validity/error of the result if a variable is "covered" more than once?
(c) What will be the value of an uncovered variable (interpolation/mean of neighbors/noise)?
(d) In the above target function, each scan is "charged" the same. We can think of a weighted version of the problem where each scan will have a positive weight e.
(e) The presented formulation of the problem is insufficient because it lacks the use of previously computed variables for computing other variables. This can be shown by looking at the previous 1×5 and 1×3 examples. The presented greedy algorithm yields a solution that leaves two 1×5 scans out; the optimal solution is much better leaving out four 1×5 scans.

(3) Another optimization problem is "Given a set of scans, what can we reconstruct?" — This is simply linear algebra, but if we add the constraints that we wish to minimize the (maximal, average) error propagation distance (to be defined, but should be clear), then we have a new type of decision/optimization problem. There is also a design problem: plan a set of scans so that reconstruction has "good" error localization. Maybe this can be phrased as a linear programming problem.
(4) An interesting direction, may be more interesting in practice than in theory, is to make use of *random* perturbations in the origin location so as to improve resolution. The advantages are many — primarily in that we do not have to set the origin exactly, all we do is set a random origin. The method should use a registration algorithm to recognize which offsets were actually used to generate the image.

9. Contact information

Further information on our work including full datasets of MR images can be found on: http://www.cs.tau.ac.il/~fiat/MRI/MRI.htm. Also available is a Matlab software that fully implements the presented algorithm.

Acknowledgments

We wish to thank Hezy Yeshurun, Nahum Kiryati, Leonid P. Yaroslavsky and Kostya Malsev for their very useful remarks and discussions on problems of SRR in MRI. Arie Yeredor for his remarks on generalization of the Nyquist sampling theorem for infinite band-limited signals. Kevin King for useful discussions in regard to GE scanner MRI data acquisition and processing and for useful comments in regard to the paper. Anna Altshuler and Janusz Hankiewicz for technical assistance. Zvi Regev for useful discussions in regard to the scanner's electronics. This research was supported in part by the National Institute of Health Grant NS37804 to Daniel Fiat.

References

[1] Greenspan H, Oz G, Kiryati N, Peled S. MRI inter-slice reconstruction using super resolution. Magn Reson Imaging 2002;20:437–46.
[2] Huang TS, Tsai RY. Multi-frame image restoration and registration. In: Huang TS, editor. Advances in computer vision and image processing, vol. 1. Greenwich (Conn): JAI; 1984. p. 317–39.
[3] Kim SP, Bose NK, Valenzuela HM. Recursive reconstruction of high-resolution image from noisy undersampled frames. IEEE Trans Acoust Speech Signal Process 1990;38:1013–27.
[4] Ur H, Gross D. Improved resolution from subpixel shifted pictures. Graph Models Image Proc 1992;54(2):181–6.
[5] Irani M, Peleg S. Motion analysis for image enhancement: resolution, occlusion, and transparency. J Vis Commun Image Represent 1993;4(4):324–35.
[6] Irani M, Peleg S. Super resolution from image sequences. 10th Int Conf Pattern Recognit 1990;2:115–20.
[7] Patti AJ, Sezan MI, Teklap AM. High-resolution image reconstruction from a low-resolution image sequence in the presence of time-varying motion Blur. Proc. ICIP, Austin, Texas; 1994. p. 343–7.
[8] Elad M, Feuer A. Super-resolution reconstruction of an image. IEEE 1996;391–4.
[9] Elad M, Feuer A. Restoration of a single superresolution image from several blurred, noisy and undersampled measured images. IEEE Trans Image Process 1997;6:1646–58.
[10] Elad M, Hel-Or Y. A fast super-resolution reconstruction algorithm for pure translational motion and common space invariant blur. IEEE Trans Image Process 2001;10(8):1187–93.
[11] Cheeseman P, Kanefsky B, Kruft R, Hanson R, Stutz J. Super-resolved surface reconstruction from multiple images. Technical Report FIA-94-12, NASA Ames Research Center, Artificial Intelligence Branch, October 1994.
[12] Shekarforoush H, Chellappa R. Data-driven multichannel super-resolution with application to video sequences. J Opt Soc Am 1990;16(3):481–92.
[13] Peled S, Yeshurun H. Superresolution in MRI: application to human white matter fiber tract visualization by diffusion tensor imaging. Magn Reson Med 2001;45:29–35.
[14] Zomet A, Peleg S. Super-resolution from multiple images having arbitrary mutual motion. In: Chaudhuri S, editor. Super-resolution imaging. Norwell (MA): Kluwer Academic; 2001. p. 195–209.
[15] Wehrli FW. Principles of magnetic resonance. 2nd ed. Magnetic resonance imaging, vol. 1. Mosley Year Book [Chapter 1].
[16] Papoulis A. Generalized sampling expantion. IEEE Trans Circuits Syst 1977;CAS-24:652–4.
[17] Yen LJ. On nonuniform sampling of bandwidth limited signals. IRE Trans Circuit Theory 1956;3:251–7.
[18] Fiat D. Method of enhancing MRI signal. US patent 6,294,914 B1, filed September 3, 1997, issued September 25, 2001.
[19] Golub GH, van Loan CF. Matrix computations, chapter 5: orthogonalization and least squares. 3rd ed. Baltimore: The John Hopkins University Press; 1996. p. 206–74.
[20] Liang Z-P, Lauterbur PC. Principles of magnetic resonance imaging: a signal processing perspective. New York: IEEE Press; 2000.
[21] Gudbjartsson H, Maier SE, Mulkern RV, et al. Line scan diffusion imaging. Magn Reson Med 1996;36(4):509–19.