A Pixel is an Artifact: On the Necessity of Zero-Filling in Fourier Imaging

XIAOLU ZHU, BOGUSLAW TOMANEK, JONATHAN SHARP

Medical Devices, National Research Council of Canada, Calgary, AB, Canada

ABSTRACT: MR imaging data is sometimes presented in a 'patchwork quilt' format with individual pixels visible as squares of uniform intensity. This phenomenon often arises by default from an image space convolution (performed implicitly by the graphics system) used to convert the sparse point sampling of the spatial domain offered by the discrete Fourier transform (DFT) into a sufficiently dense sampling to allow assignment of an intensity value to each addressable point on the display device. Typical examples are fMRI maps, spectroscopic images and zoomed-in views. These square patches are image structure not present in the object, i.e., artifacts. This form of image display is studied by both an image analysis method and by Fourier analysis. Image formation by display of the 2D DFT of an acquired k-space matrix as a 2D pixel array is a poor reconstruction because it does not ensure a faithful representation of the spatial frequency content actually present in the data. By analysis of the visual appearance of 2D pixel arrays we show that there are two principal effects: (a) attenuation of higher spatial frequencies (i.e., low-pass filtering); (b) introduction of artifactual high frequency image structure. These effects can lead to very poor performance with an artifact/signal ratio of over 200% in the corners of 2D k-space. Generated k-space maps demonstrate that both detrimental effects increase radially in k-space. The simple remedy is to zero-fill (resulting in image interpolation) until individual pixels become invisible in the displayed image. Alternatively, data modeling may be used. Concepts Magn Reson Part 42A: 32–44, 2013.

KEY WORDS: MRI; Fourier transform; zero-filling; image interpolation; *k*-space

INTRODUCTION

The Pixel Problem

Browsing through any recent MRI journal one quickly encounters images consisting of arrays of small squares of uniform intensity ("pixels"), resembling a "patchwork". To illustrate that this is a topical concern, in a randomly selected recent issue of Magnetic Resonance in Medicine (printed edition) small squares are visible in 14 figures (Page numbers

Received 23 April 2012; revised 20 December 2012; accepted 12 February 2013

Correspondence to: Jonathan Sharp. E-mail: sharpmri@gmail.com

Presented in part at 19th ISMRM Annual Meeting, Montreal Canada, 2011.
Concepts in Magnetic Resonance Part A, Vol. 42A(2) 32–44 (2013)
Published online in Wiley Online Library (wileyonlinelibrary.com).
DOI 10.1002/cmr.a.21256

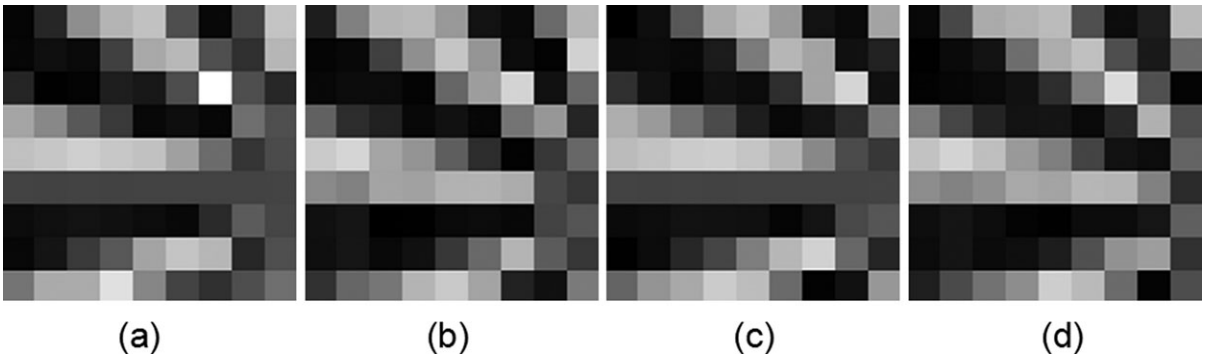

Figure 1 A demonstration that Fourier imaging performed with DFT reconstruction without zero-filling is not shift-invariant. These different pixelated images were all derived from the same raw data, using a digital Siemens star phantom. Different linear phase rolls in *k*-space were used to produce sub-pixel shifts resulting in images which vary significantly. (a) No shift; (b) shift down half a pixel (c) shift right half a pixel (d) shift down and right half a pixel.

are 87, 121, 122, 133, 134, 135, 147, 177, 186, 221, 222, 240) (*1*). Common examples are fMRI activation maps, parametric maps, spectroscopic metabolite maps and zoomed-in images. Where do these small squares come from? These squares exist neither in the object, nor in the raw Fourier *k*-space data. They are not apparent in a spreadsheet tabulation of the 2D data matrix resulting from 2D Discrete Fourier transform (2D-DFT) of a *k*-space matrix. The small squares are in fact features generated for image display because each point on the monitor screen needs to be assigned a value. Because these small squares appear as visible structure within the displayed image, but do not reflect object structures, there is really no other choice than to classify them as artifacts. We refer to the visibility of these small squares as "pixelation artifact," and the analysis of these artifacts is the topic of this article. We will establish the relationship between these patchwork displays and the spatial frequency content of the display.

These small squares present two types of problem. First, in a DFT-reconstructed MR image the intensity assigned to a small square does not equate to the mean object intensity over the square region. This might be contrasted with the quite different situation found with CCD optical detectors, in which the detector does actually integrate signal over small square regions.

Second, for any given *k*-space dataset there are, in any case, many possible reconstructions of these square-patchworks, because the spatial positioning of the reconstructed pixel array is arbitrary—it depends upon the linear phase roll given to the *k*-space data along the k_x and k_y axes before 2D-DFT. There is nothing in the data to justify the particular choice of any one 2D patchwork of small squares over any other. Figure 1 shows four different pixel patchwork grids, reconstructed from the same *k*-space data. The phantom shown is a portion of a Siemens Star phantom which consists of 24 alternating black and white 15° radial wedges in a sunburst pattern and is useful for non-directional resolution evaluation. Figure 1(a) demonstrates that for a nonzero-filled image there is no unique PSF as the imaging system response is not shift-invariant.

The images shown result from direct DFT of simulated *k*-space data. Three different linear phases were added to the *k*-space data, each equivalent to less than a pixel shift in the image domain. This results in images of significantly different appearance and illustrates that arbitrary selection of any single one out of the multitude is problematic and biased. Equivalently, consider that the registration of the patchwork pixel grid relative to the object depends upon the exact positioning of the object relative to the gradient coils, which is chance, and clearly not indicative in any way of the object structure, and therefore should not be reflected in the displayed image. The solution to these issues is to avoid the use of patchwork displays at all. We can conclude that pixels should be small enough to be invisible, i.e., some form of interpolated display should be used.

Another point to note is that there is important information contained in each data point in the spatial domain. Consider that (given adequate SNR) a 256 x 256 dataset offers significantly improved image quality over a 128 x 128 dataset. The difference between the two being additional high-resolution information, i.e., details on the pixel-scale. It is apparent that such

pixel–pixel variations are valuable and that consideration of visualization procedures is a worthwhile pursuit.

Zero-Filling

It is useful to include some introductory comments on zero-filling. It is sometimes stated that adding zeroes to a dataset (zero-filling) does not add information, with the implication that the procedure is therefore somehow unnecessary, superfluous, or merely cosmetic. However this is not the case. Information content and image display (i.e., a faithful spatial representation leading to reliable visualization) are different concepts. One might as well argue that since the DFT does not introduce new information, it is also superfluous! Or one might present image data in a 256 $\times$ 256 spreadsheet instead of as an intensity modulated display—after all the information content is the same! Clearly visualization matters. Zero-filling exerts its influence through the assumptions made about missing data—data not acquired, yet nevertheless impacting image display.

Interpolation

Considering the fundamental importance of display in any imaging modality, there is relatively little literature to be found on the interpolated display of Fourier-based MRI images. There is, to be sure, an extensive general literature on image interpolation, but this is mostly concerned with efficiency and accuracy of various image domain interpolation methods (*2–5*). MRI is a type of Fourier imaging, and so the image domain interpolation literature is not directly relevant.

In 1D NMR spectroscopy some form of interpolation is always used, even if it is merely joining adjacent data points on a graph with straight line segments. A noninterpolated display (the equivalent of 1D pixels) would have the appearance of a bar graph, which would certainly look odd. There is an NMR literature on zero-filling (*6*, *7*). Bartholdi and Ernst analyze the particular case of zero-fill by a factor of 2 (*6*). Lindon and Ferrige (*7*) show that there are situations in which a zero-fill factor as high as 32 can be beneficial in resolving overlapping peaks. Our findings are most in accord with this latter result.

Within the MRI literature there are a number of reports on the benefits of zero-filling and interpolation in a variety of circumstances (*8–17*). The most directly relevant work is by Bernstein et al. in their discussions of the corners of *k*-space (*14*). We focus exclusively on the case of Fourier imaging and MRI, but the conclusions may be relevant to other Fourier imaging modalities.

Spatial Frequency Analysis

We will present an approach analyzing the actual visual scene presented by a 2D array of finite-sized squares of uniform intensity ("image pixels"). The rationale is that the visual representation on the display device determines the image interpretation by an observer. We will also emphasis that there is an important distinction between the sampled spatial data output from the 2D-DFT algorithm, on the one hand, and the visual scene presented to the viewer by a 2D array of image pixels on the other. The spatial frequency analysis presented involves the derivation of *k*-space maps showing levels of pixelation artifact and signal loss. These maps can be used to maximize signal/artifact ratio by informing zero-fill strategies and the selection of the shape of the sampled *k*-space region used in image acquisition and processing.

The maps show that display of the 2D-DFT of an acquired *k*-space matrix as a 2D pixel array does not result in a faithful representation of the spatial frequency content actually present in the data. We will show that there are two effects: (1) attenuation of higher spatial frequency image content (i.e., a low-pass filtering effect); (2) introduction of artifactual high frequency image structure ("pixelation artifact"). In some situations the artifact level can exceed that of the signal. This approach is confirmed analytically, and verified with a series of test images. The analysis also shows that the remedy for pixelation artifacts is very simple—just perform sufficient zero-filling.

METHODS

K-Space Maps of Signal Loss and Pixelization Artifact

Figures 2(a,b) show two examples of the decomposition of a 2D patchwork array of finite-sized square image pixels into signal and artifact images. The procedure used to perform this decomposition is shown as a series of steps (A–F) in Fig. 2(c). The purpose of this analysis procedure is to derive measures of the signal and artifact content of a pixelated image display, as a function of spatial frequency. *K*-space maps of signal and artifact content are generated point-by-point. In order to only consider real objects, a pair of test points: (k_x, k_y) and ($-k_x$, $-k_y$) with complex conjugate values were used. This whole process

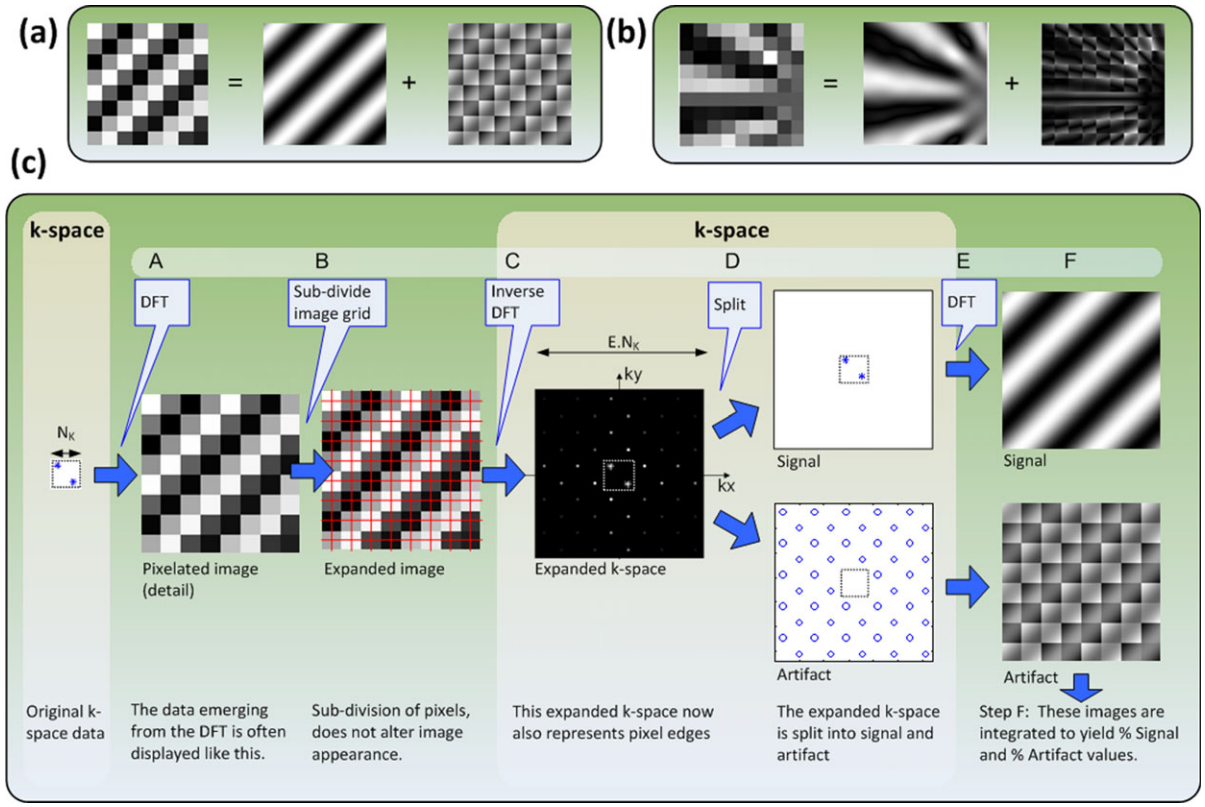

Figure 2 Figure (a) shows the separation of a pixelated image of a single spatial frequency into signal and pixelation artifact components. Figure (b) shows a second example, using part of a Siemens star phantom, to illustrate that the procedure may be applied to any pixelated image. Note that the "signal" image is not the ideal reconstruction as it has undergone a low-pass filtering effect. Figure (c) shows in detail how this analysis is achieved and its use to evaluate pixelation artifacts. First, an $N_K \times N_K$ *k*-space array is populated with two real-valued points, symmetric with respect to the origin, representing a single spatial frequency for a real object. This array is transformed to the spatial domain by 2D DFT to produce an $N_K \times N_K$ image (Step A); the data is then processed so that it is represented on an expanded, higher resolution pixel grid (Step B). So for an expansion factor of 5 each original pixel is split into 25 smaller pixels. This expanded grid is then transformed back to *k*-space by inverse 2D DFT (Step C), where it is split into two parts representing signal and pixelation artifact respectively (Step D). Finally these two parts are transformed back to the image domain (Step E) and integrated (Step F). The integrals of the signal and artifact images represent the values of signal and pixelation artifact.

(A–F) is performed for each pair of *k*-space coordinates (i.e., for each spatial frequency), with the resulting signal and artifact levels used to build-up *k*-space maps of signal and artifact respectively.

Step A - DFT

The starting point is an empty ($N_K \times N_K$) 2D *k*-space array, which represents the area that would be filled with acquired data in an imaging experiment. The array is populated with a single pair of non-zero test points, as just described, marked with star symbols [Fig. 2(c)]. Step A is the discrete Fourier transformation of this ($N_K \times N_K$) *k*-space array into the spatial domain to produce an ($N_K \times N_K$) image domain array.

Step B – Pixel Duplication

Step B is an expansion of the ($N_K \times N_K$) image domain data to a larger matrix. Each pixel in the ($N_K \times N_K$) image domain array is replaced with E^2 daughter pixels, where E is the linear expansion factor (e.g., E = 2, 1 pixel → 4 pixels). This produces a new matrix

of the size ($EN_K \times EN_K$). The screen area covered by the E^2 new pixels is considered equal to the area of the originating pixel so that the visual appearance of the image is not altered in any way. The image domain representation has changed, but the visual appearance has not.

We assume that the screen display pixels are small enough that they cannot be visually resolved, and so the visual appearance is determined solely by the data, not the display device. This is a good assumption, given the advanced state of today's display technologies. Note that that this pixel subdivision procedure is just a step in the analysis procedure employed in this article. (It is not proposed as an image processing procedure.)

Step C – Formation of the Extended *k*-Space

Step C is the inverse 2D-DFT of the expanded image to generate *k*-space domain information. Because each pixel has been subdivided, and because smaller structures require higher spatial frequencies for their representation, it is found that this new ($EN_K \times EN_K$) *k*-space is populated out to higher *k*-space coordinates. Thus the content of this new *k*-space is different from the original ($N_K \times N_K$) *k*-space, even though the image visual appearances are identical. The patchwork pixel array is not uniquely represented by the originating spatial data array. The conception of a 1-to-1 relationship between an acquired N^2 *k*-space dataset and an N^2 image is erroneous. The star markers denote the locations of the original signal data points. The small central dotted square represents the size of the original 128 by 128 *k*-space.

Step D – *K*-Space Decomposition

Step D is the decomposition of the ($EN_K \times EN_K$) *k*-space array from step C into two ($EN_K \times EN_K$) arrays, one representing the signal, and the other the artifact. The signal array is created by setting all values outside the central ($N_K \times N_K$) square region to zero. Conversely, the pixelation artifact array is the complementary array with the central ($N_K \times N_K$) region set to zeros. The ($N_K \times N_K$) square is shown as a small central square. The justification for this split is that the central ($N_K \times N_K$) square contains the original *k*-space signal (star symbols). All data outside this region is related to pixelation artifact (circle symbols), i.e., contains information on pixel shape—which was not present in the acquired *k*-space.

Step E – DFT to Form HR Images

Step E is the DFT of these two ($EN_K \times EN_K$) *k*-space matrices back into the image domain to form respectively the high-resolution (HR) signal and HR pixelation artifact images (magnitude). "High resolution" simply means that there are now more pixels/mm (a higher pixel density) than in Step A. After step *E* are shown expanded signal and pixelation artifact images. The original (128 ×128) *k*-space was expanded by a factor of 20 to (2,560 × 2,560). The cosine wave structure of the image is set by the choice of the original two *k*-space data points which were located at (32, 32), (-32, -32) for the example shown in Fig. 2(c). Each square (actually 20x20 pixels) was a single pixel before expansion. The signal image (upper) is an image with a smooth single frequency cosine wave as expected for a single spatial frequency. However this is still not an ideal image, as the low-pass filtering effect is still present. The pixelation artifact image (lower) reveals the shape and edges of the original pixels and also high frequency harmonics of the signal.

Step F - Integration

Step F is simply the integration (summation) of all pixel magnitudes in each output image, resulting in a single number for each image. These two numbers are used to populate the *k*-space maps for signal level and artifact level respectively. Both signal and artifact maps are normalized by the signal level at the center of the *k*-space map.

Test Data Generation

Reconstruction strategies were evaluated using accurate digitally generated *k*-space test data (Figs. 1 and 2). The phantom was first defined spatially as a set of spin coordinates on a fine grid (15 × 15 times finer than the pixel spacing). The *k*-space content corresponding to each spins was then created using a direct synthesis procedure, as follows. For a spin in the center of the image, a constant value was assigned for *k*-space. For spins spatially shifted away from the center of the image, a linear phase shift in the *k*-space was added. So if $h(x, y)$, in the image domain, and $H(u, v)$, in *k*-space, are an FT pair, then the FT pair for off-center spins is:

$$h(x,y)e^{j2\pi(xu_0+yv_0)} \overset{\text{FFT}}{\Rightarrow} H[(u-u_0),(v-v_0)] \quad [1]$$

The *k*-space data for the entire phantom consists of a summation of the contributions from

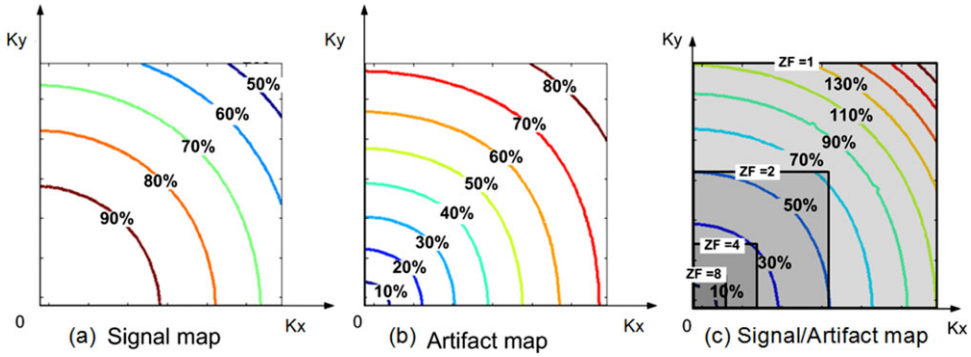

Figure 3 The process illustrated in Fig. 2 was carried-out individually for each spatial frequency representable in the original *k*-space matrix. Thus for each pair of *k*-space points a value for % signal and a value for % artifact is obtained. These values were plotted as the contour maps shown here. Only one quadrant of the *k*-space plane is shown. The original 128 × 128 *k*-space (containing 2 points, symmetric with respect to the origin) was expanded by a factor of 20. (a) Signal Map: shows the percentage of the signal compared with the center of the *k*-space. A radial loss of signal visibility is observed, ranging from 100% signal (no loss) at the center of *k*-space to 64.7% signal at $(k_X, k_Y) = (N_K/2, 0)$, $N_K = 128$. The signal visibility level falls to its lowest value of 41.9% at the corner of *k*-space, $(N_K/2, N_K/2)$. (b) The *k*-space artifact map represents the integrated pixelation artifact. Both this map and the signal map are normalized relative to the origin of the signal map. The artifact level is zero at the *k*-space origin, signifying no visible pixelation for a uniform image. The artifact level increases steadily radially in *k*-space, with the worst performance is at the corner of *k*-space where artifact level is 87.3%. (c) Signal-to-Artifact Ratio (S/A) Map: is formed by the ratio of the two previous maps. At $(N_K/4, N_K/4)$ S/A = 69%; at $(N_K/2, 0)$ S/A = 110%; at $(N_K/2, N_K/2)$ S/A = 210%. Overlaid on the S/A map is a *k*-space explanation of why zero-filling improves signal to pixelation artifact ratio. Sampled data is represented by the shaded area. Four different zero-filling levels are illustrated: ZF = 1, 2, 4, 8. As the level of zero-fill is increased the acquired data resides closer and closer to the origin of the map, where pixelation artifact levels are lower.

each spin. In this way the test *k*-space closely represents the *k*-space data that would be collected in a real experiment for a continuous object. This is particularly important for this study, to avoid any potential confusion between (nonphysical) sample pixelation effects and reconstruction pixelation effects.

K-Space Filter Design

To reduce ringing that occurs near sharp image edges, due to direct DFT of the *k*-space data, *k*-space filtering was used in Fig. 4. The form of this circular *k*-space filter was a simple transition band. For a 128 x 128 *k*-space, two circles are defined: an outer circle with a radius of 64, and an inner circle with a radius of 56. The filter has values outside of the outer circle set to 0; values between the two circles as 0.5; and values inside the inner circle as 1.

Zero-Fill Factor Definition

In all cases, the *k*-space matrix is assumed to be square. We define the 1D zero-fill factor (ZF) as:

$$\mathrm{ZF} = \frac{\#\text{ rows in final } k\text{-space matrix}}{\#\text{ rows in acquired } k\text{-space matrix}} \qquad [2]$$

For example, for an acquired *k*-space size of 256 x 256, zero-filled into a final *k*-space size of 512 x 512, then ZF = 2.

RESULTS

K-Space Maps

Figure 3 shows results for signal, artifact, and artifact/signal maps. The maps were generated with an expansion factor of 20, as an approximation to an infinite expansion. This is a reasonable assumption because based on simulation experiments, the artifact

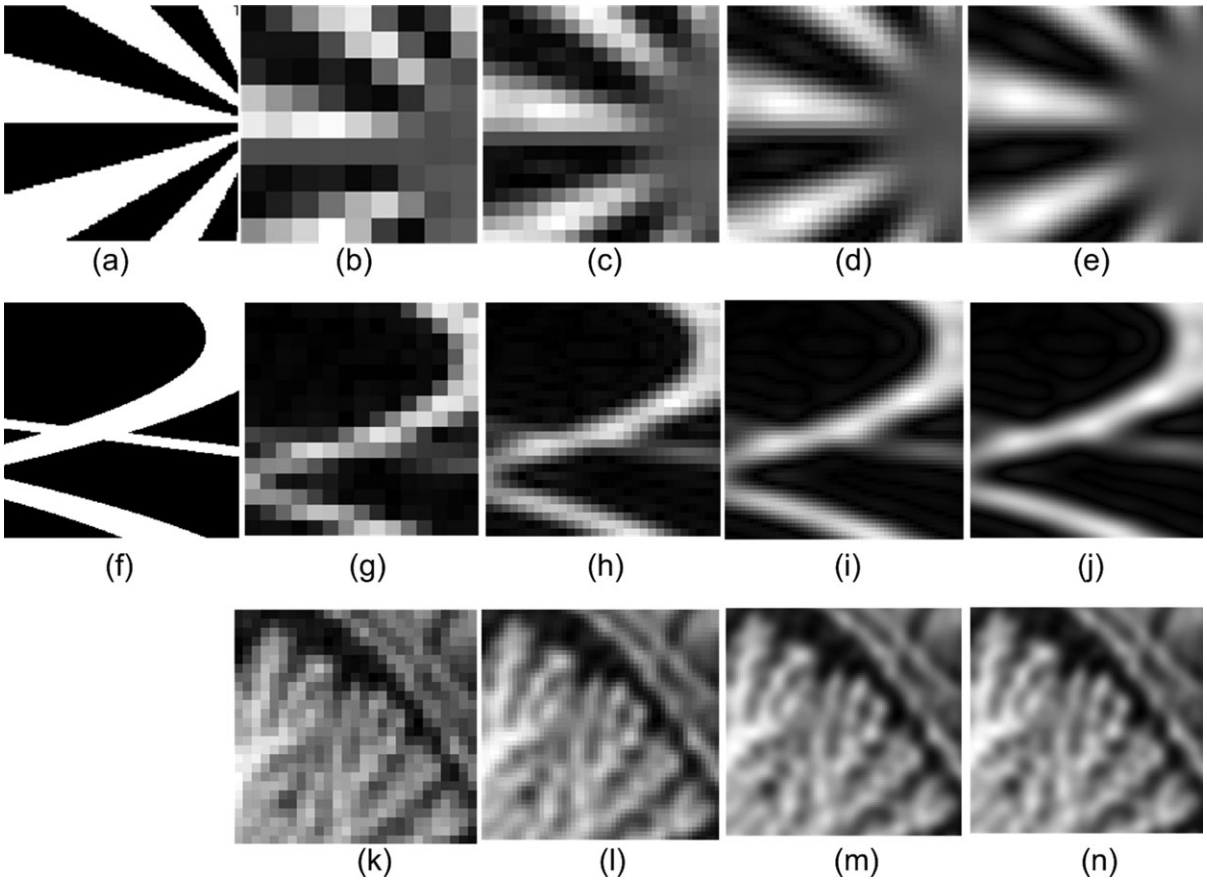

Figure 4 This figure shows that images with higher zero-fill factors have less pixelation artifact and show better details in the image. A part of the "Siemens star" resolution phantom is shown in (a–e). The phantom shown in (f–j) is composed of a superposition of three sine-wave bands and demonstrates the representation of curved contours in the image. Panel (k) is a part of a brain image. Panels (a) and (f) are grids with a set of spins shown as binary values, (spin exists = 1, spin does not exist = 0). Details of how to generate *k*-space raw data based on the spin grid are in Section 'Test Data Generation'. Since (k) is derived from data collected from an MRI imaging system, there is no grid to be shown. In (b), (g), and (k), each image is reconstructed with no zero-fill, and the pixelation artifact is very high. In (c), (h), and (l), ZF = 2; in (d), (i), and (m), ZF = 4; in (e), (j), and (n), ZF = 8. As zero-fill increases, the true image structure becomes progressively more visible. Ringing artifacts can be observed where sharp edges exist.

to signal level converges closely to a constant value when the expansion factor is >15.

Signal Map

The *k*-space signal map [Fig. 3(a)] represents the integrated image intensity of a spatial frequency relative to the corresponding pair of points in *k*-space. The map shows a radial loss of signal visibility, ranging from 100% signal (no loss) at the center of *k*-space to 64.7% signal at $(k_X, k_Y) = (N_K/2, 0)$. The signal visibility level falls to its lowest value of 41.9% at the corner of *k*-space, $(N_K/2, N_K/2)$. By employing ZF = 2 in each spatial dimension the lowest signal visibility improves to 80% at coordinates $(N_K/4, N_K/4)$. This map reveals a low-pass filtering effect inherent in the non zero-filled DFT image reconstruction process.

Artifact Map

The *k*-space artifact map [Fig. 3(b)] represents the integrated pixelation artifact. The map is normalized identically as the signal map, i.e., relative to the ori-

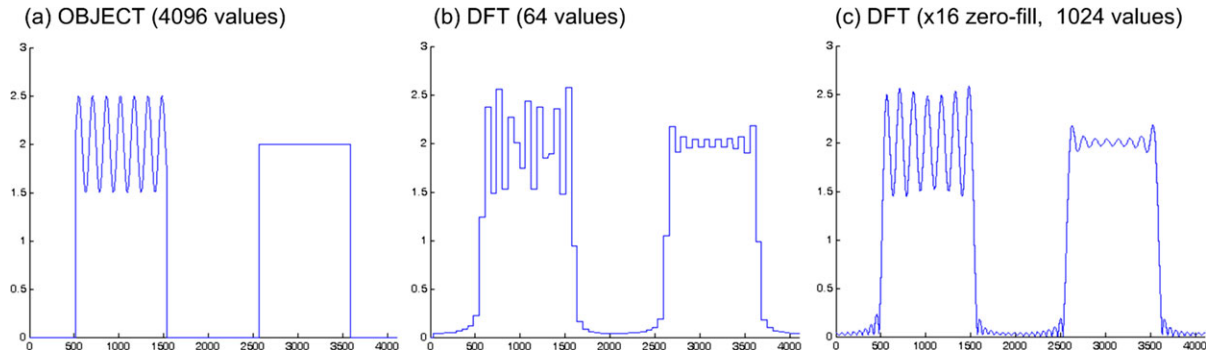

Figure 5 1D simulations showing the effect of zero-filling on truncated *k*-space data. (a) Object profile (4,096 points), used to generate *k*-space data. The profile contains two sharp-edged blocks, one overlaid with a sine function. (b) Profile showing the effects of truncation, using DFT of 64 *k*-space points, displayed in "staircase" mode to mimic pixelated image display. (c) As (b), but zero-fill by factor of 16 prior to DFT, showing much better visualization of object structure (sine wave). The appearance of the edge ringing (due to *k*-space windowing) changes and fine ringing on the baseline becomes resolved.

gin of the signal intensity map. The artifact level is zero at the *k*-space origin, signifying that there is no visible pixelation for a uniform image. The artifact level increases steadily radially in *k*-space. The worst performance is again at the corner of *k*-space where artifact level is 87.3%. This indicates that the intensity of these artifactual structures increase as the signal being represented approaches the pixel size.

Artifact-to-Signal Ratio (A/S) Map

This Artifact/Signal (A/S) [Fig.3(c)] map is formed by the ratio of the two previous maps, and expressed as a percentage. The A/S ratio is zero at the origin. At the corner of the *k*-space, (N_K/2, N_K/2), the A/S ratio is 210%, i.e., artifact is more than twice the signal amplitude. At the end of the $+k_X$ axis, coordinates (N_K/2, 0), the A/S ratio is 110%. If a zero-fill factor of 2 is used, then the worst A/S ratio improves to 69% at (N_K/4, N_K/4). This is still poor, implying that further zero-filling is required.

Levels of Zero-Fill

The *k*-space maps predict that Artifact/Signal ratio is minimized by using a circular *k*-space and a high level of zero-filling. The use of high levels of zero-fill (ZF) are illustrated in Fig. 4, which shows image results for two digitally defined images (using circular *k*-space filtering) and one in vivo dataset. The phantom in Figs. 4(a–e) is the center part of a "Siemens star" phantom. Figure 4(a) is a grid with a set of spins shown as binary values. Spin exist = 1, spin does not exist = 0. The structure of the "star" is not clearly seen for ZF=1 [Fig. 4(b)] because the pixel size is comparable to the size of the structure. The 15° angled edge is very jagged due to pixel edges. The jaggedness is progressively reduced for ZF = 2, 4, and 8 [Figs. 4(b–e)]. The higher zero-fill, the less effect of the pixelation artifact.

In Figs. 4(f–j), a phantom composed of a summation of three sine-wave bands is defined. This series demonstrates the image representation of curved structures for increasing levels of zero-fill. The three sine-wave bands are progressively clearer for increased zero-fill. Ringing is visible near sharp edges of the phantom. This will be addressed and explained in details in the following section.

In Figs. 4(k–n), a part of an in vivo brain image is reconstructed (Siemens 3T). There are branch-like structures in the image which are visible in Fig. 4(n). The same structure is only shown as blocks of squares in Fig. 4(k) (ZF = 1).

Truncation and Ringing Artifacts

In MRI, the FT of the *k*-space sampling function defines a window function $W(x)$ with which the object $O(x)$ is convolved to yield the image: $I(x) = W(x)*O(x)$. The window function $W(x)$ can result in ringing in $I(x)$ if high spatial frequencies from sharp edges present in $O(x)$ are truncated. $W(x)$ also fundamentally limits the ultimate image resolution.

In Fig. 5 simulated *k*-space data was obtained by FT of a digitally defined 4096 object profile, followed by *k*-space truncation to 64 points. Two recon-

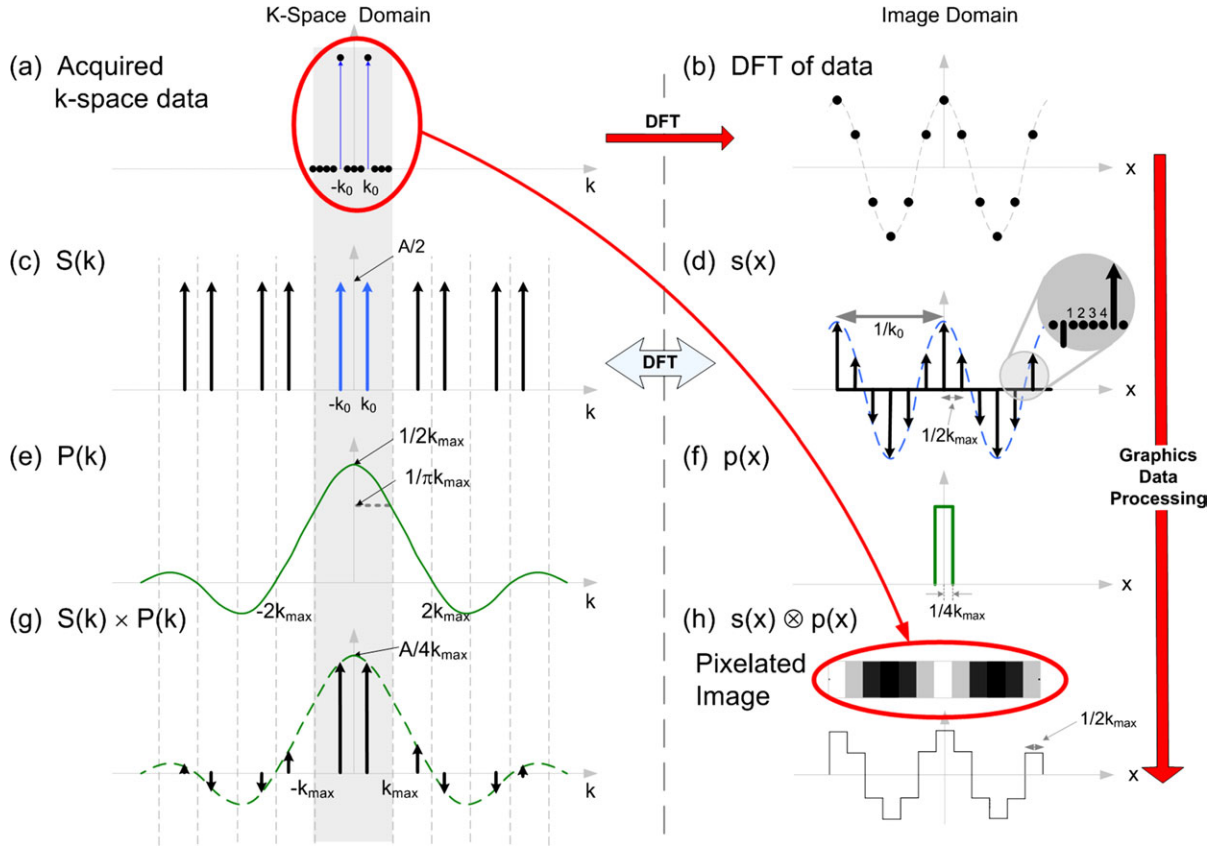

Figure 6 A graphical explanation of the transformation of acquired *k*-space data (a) into a pixelated image figure (h). Plots on the left are in the *k*-space domain, those on the right, the spatial image domain. In figure (a) the raw *k*-space data is shown as discrete sampled points over a limited *k*-space range (shaded). The first step is DFT resulting in a set of points in the image domain, figure (b). At this stage in the process there is a set of points, which is usually too sparse to be suitable for direct graphical display. Figures (d) and (f) represent graphics data processing steps necessary to obtain the image function shown in (h). Figure (d) shows a function $s(x)$, with explicit zeros inserted between the data points from (b), Figure (f) shows a pixel function $p(x)$, which when convolved with $s(x)$ yields the image function shown in figure (h). The *k*-space representation of $s(x)$ contains periodic repeats of the acquired data. Multiplication of this by the sinc function (FT of $p(x)$) yields the *k*-space representation of the pixelated image, which contains an attenuated version of the original data in addition to higher frequency artifactual content.

structions are shown, without [Fig. 5(b)] and with zero-filling [Fig. 5(c)]. Ringing is visible in both cases, demonstrating that zero-filling neither creates the ringing phenomenon, nor eliminates it. Zero-filling does however affect the appearance. Figure 5(c) shows that zero-filling simultaneously results in clearer visualization of both object structure and fine details of the edge-ringing. While the former is a clear benefit, the latter is more equivocal. Some windowing effects may only become visible as ringing after zero-filling, as can be seen by comparing the baselines in the profile in Fig. 5(b) with that in Fig. 5(c).

A detailed discussion of strategies for ringing artifact reduction is beyond our scope here, but we mention briefly some possible approaches: (1) the collection of more data; (2) filtering of the acquired data; or (3) modeling to fill-in the missing data. In this work we used a simple filter to reduce ringing. Techniques for reducing the effect of ringing have been studied extensively in the literature (*2*). A disadvantage of *k*-space data weighting filter functions (e.g., Hanning, Bartlett, Parzen, and Dolph-Chebyshev) is that they may reduce resolution. Data modeling methods require some additional information or assumptions.

ANALYSIS

Fourier Transform Analysis

The *k*-space map generation process (Steps A–F) were verified using Fourier transform theory (*18*). Figure 6 illustrates graphically the process for the 1D case. Original *k*-space data (representing a single spatial frequency, k_0) is shown as 12 data points in Fig. 6(a), with the extent of the sampled region shaded ($-k_{max}$ ... + k_{max}). The DFT of this data is shown in Fig. 6(b) as 12 points in the spatial domain. In the case, because the cosine is band-limited (i.e., no part of its *k*-space representation was truncated by the *k*-space window) these data points coincide with the analytical cosine function (dotted line). However, in MRI it is common that the acquired *k*-space data is a truncated version of the full *k*-space signal describing the object. In these cases, because the data points following DFT are samples of *I(x)* [i.e., the convolution of the object with a window function *W(x)*], *S(x)* may exhibit ringing at sharp edges. As was shown in Fig. 5, zero-filling results in a denser sampling of *I(x)*, but does not modify *I(x)*, so ultimately image resolution is determined by the *k*-space window, not by the level of zero-filling. Once *I(x)* becomes well-visualized, additional zero-filling has no further useful effect.

In Fig. 6(d) a function *s(x)* is shown which consists of zeros inserted between the 12 points of Fig. 6(b). Figure 6(f) shows a pixel function *p(x)* which when convolved with *s(x)* yields the pixelated image [Fig. 6(h)]. This process represents the data processing performed by the graphics display system when DFT data is displayed as an image without zero-filling. By examining the *k*-space domain representation of this process [in Figs. 6(c,e,g)] both the low-pass signal filtering effect (within the shaded region) and the origin of the *k*-space content corresponding to pixelation artifacts (Fig. 2) can be appreciated. In detail: the DFT of *s(x)* is *S(k)* which consists of the original *k*-space data plus 4 periodic repeats. The DFT of *p(x)* is *P(k)* which is a sinc function. The multiplication of *S(k)* and *P(k)* yields the *k*-space representation of the pixelated image, which is valid irrespective of the actual mathematical process used in any particular display system.

Low-Pass Filtering Effect (1D)

Figures 6(g,h) show how the amplitude of the 1D image signal is modulated by a sinc function:

$$\text{signal}(x, k_0) = \frac{A}{2k_{max}} \text{sinc}\left(\frac{\pi k_0}{2k_{max}}\right) \cos(2\pi k_0 x) \quad [3]$$

This "signal" function is a filtered version of *I(x)*, i.e. it exhibits both the filtering effect and *k*-space windowing effects from *W(x)*. The 2D version of this function is labeled "signal" in Fig. 2.

For 1D, the maximum attenuation due to the sinc function arises for the condition $k_0 = k_{max}$, i.e. $\text{sinc}(\pi/2) = 2/\pi = 0.6366$. To compare this theoretical result and the simulation experimental result, the Pearson correlation coefficient was calculated. A very close fit was found between the simulated signal map and the 2D sinc distribution (Pearson Correlation Coefficient = 0.9999). This result is also in exact agreement with Lindon and Ferrige (*7*) who in their discussion of zero-filling in NMR spectroscopy report that without zero-filling the worse case (a spectral peak lying midway between data points) results in a peak height reduction by a factor of 0.637. They report peak height recovery to 0.998 with a zero-filling factor or 16.

1D Pixelation Artifact—Analysis

The pixelation artifact represents the difference between the original pixelated image and the low-pass filtered version of *I(x)*, the "signal" function. To evaluate the pixelation artifact, we can sum the image domain contributions of all the pixelation artifact terms in Fig. 6(g) to give a 1D version of the total pixelation artifact:

$$\text{artifact}(x, k_0) \sum_{n=1}^{\infty} [A_n \cos(2\pi(2nk_{max} - k_0)x) + A'_n \cos(2\pi(2nk_{max} + k_0)x)] \quad [4]$$

where:

$$A_n = \frac{A}{4k_{max}} \text{sinc}\left(\frac{\pi(2nk_{max} - k_0)}{2k_{max}}\right) \quad [5]$$

$$A'_n = \frac{A}{4k_{max}} \text{sinc}\left(\frac{\pi(2nk_{max} + k_0)}{2k_{max}}\right) \quad [6]$$

and where *n* is the "order" of the artifact. The coordinates $2nk_{max}$ are shown by vertical dotted lines in Figs. 6(c,e,g). Thus *n* = 1 represents the 4 largest artifact peaks, consisting of the pair centered around $+2k_{max}$ and the pair centered around $-2k_{max}$.

If this whole process, as illustrated in Fig. 6, were repeated with a truncated MRI dataset, then the acquired spatial frequencies would undergo the mod-

Table 1 Tabulation of Maximum and Average Pixelation Artifact/Signal Ratio for Circular and Square *k*-Space Masks, for a Range of Zero-Fill Factors

		Shape of *k*-Space Region Populated With Data									
		Square Mask					Circular Mask				
ZF (1D)		1	2	4	8	16	1	2	4	8	16
Fraction of *k*-space matrix populated with zeroes		0	0.75	0.938	0.984	0.996	0.215	0.804	0.951	0.988	0.997
Artifact to signal ratio	MAX	208%	69%	28%	13%	8%	111%	43%	20%	10%	5%
	AVG	79.9%	32.1%	15.3%	7.8%	4.2%	51.1%	21.9%	10.7%	5.5%	2.9%
Maximum signal loss		58.1%	18.9%	5.0%	1.2%	0.2%	35.3%	9.9%	2.5%	0.7%	0.2%
Maximum artifact level		87.3%	56.4%	26.9%	13.2%	6.5%	71.7%	39.0%	19.8%	10.0%	5.3%

A mask corresponds to the region of *k*-space populated with acquired data, and is determined by the acquired data and any *k*-space masking applied. In all cases the *k*-space matrix is assumed to be square. In the case of the circular mask, the circle is inscribed within the acquired *k*-space region. The circular mask shows lower pixelation artifact/signal ratios, especially for the lower zero-fill factors.

ifications as shown (signal attenuation and introduction of pixelation artifacts), while the loss of data due to truncation would be manifested by the deviation of the pixel values from the object signal values (due to convolution with the window function).

Zero-Fill Improves Artifact to Signal Ratio

The effect of zero-filling is to cause the acquired data to lie closer to the origin of the *k*-space maps, where the artifact-to-signal ratio is lower. This is illustrated in Fig. 3(c) for 2D *k*-space. The coverage of acquired (non-zero) *k*-space is shown as a shaded area. The ratio of non-zero *k*-space to the whole *k*-space is $1/ZF^2$. As given in Table 1, the trend is that the more zero-filling, the smaller both the maximum and average artifact to signal ratio become. Complete statistics for both square and circular masks are given in Table 1. The circular mask shows lower pixelation artifact/signal ratios, especially for the lower zero-fill factors.

DISCUSSION

The Data Not Collected

For N collected data points, the DFT yields just N points in the image domain, which is very often in practice too few for adequate image display. This arises because it is not sensible experimentally to acquire a large amount of additional (and noisy) low or zero signal data, merely to make up the numbers. How should the absence of this data be accounted for in the image reconstruction and display process?

Consider three alternatives for the status of the un-collected *k*-space data: (A) assume that the uncollected data is zero (i.e., zero-fill); (B) assume the form of *k*-space shown in Fig. 6(g), which is partially periodic and also has a low-pass filtering effect (i.e., no zero-fill); (C) fill in with modeled data. Additionally, the collected data can be treated in one of two ways prior to DFT: filtered or not. These categorizations lead to 3x2 = 6 strategies. Our thesis in this paper is that option B (DFT without zero-fill) is a surprisingly bad choice and should almost never be used. We study the zero-filling option, but do not discount other potentially even better options such as data modeling. The transition between the acquired data and a region of zeros can cause ringing artifacts at edges [Figs. 4(c–e) and (h–j)]. In these cases filtering and/or data modeling may give better results. However we argue that avoiding zero-filling for this reason is rarely justified as the pixelation artifacts present a greater impediment to visualization than does ringing (Figs. 4 and 5).

The Meaning of *k*-Space

Fundamentally, the meaning of a pair of *k*-space data points is a perfectly smooth sinusoidal image domain intensity variation, with a magnitude proportional to the *k*-space value. This is justified by considering the source of the *k*-space data: the MRI imaging process, which involves a continuous object and continuous magnetic gradient fields, and assumes that the receiver sensitivity is uniform over the FOV. A problem with a non-zero-filled DFT image (considering the 1D case) is that the highest spatial frequency data acquired can never be visualized as a smooth sine-wave, but at best as a square-wave (represented as a series of pixels with alternating intensities).

The signal map [Fig. 3(a)] is an evaluation of the observed proportionality between the two domains. The map shows that by displaying DFT data as uniform intensity pixels without use of zero-filling, the proportionality is surprisingly poor, with higher spa-

tial frequencies experiencing greater attenuation, to a maximum signal loss of 41.9% at the corners of 2D *k*-space. In other words, the image is low-pass filtered and resolution is lost. This built-in sinc filtering behavior of the DFT is a strong effect, but does not seem to be sufficiently appreciated in the MRI literature.

The very poor performance in the corners of the *k*-space maps is consistent with the conclusion in the work of Bernstein et al. (*14*) that the data in the corners of *k*-space introduce artifact to the image when no zero-filling or filtering is applied.

Pixel Function

Another perspective on this is to consider the pixel function. In Fig. 6(f) the pixel function is square, resulting in multiplication of the *P(k)* function with a sinc, with poor results. Ideally the pixel function should be a narrow sinc, resulting in multiplication of *P(k)* with a square function—which would perfectly pick out the signal with no attenuation and reject all pixelation artifact. This ideal sinc interpolation procedure is approximated by zero-filling by a large factor. So, in other words, the convolution step using a square pixel function is precisely backwards!

Zero-Filling

How much zero-filling should be performed? The commonly used zero-fill factor of two is insufficient as it results in a large artifact/signal ratio (65%) at the corner of a square *k*-space. Minimum values on ZF can be set by specifying a maximum acceptable artifact/signal ratio (Table 1). For example, for a maximum acceptable A/S ratio of 10%, ZF = 8 and a circular mask is required. For A/S less than 5%, ZF = 16 is required. Figure 4 shows that a lot of structure can be extracted from relatively few pixels, which implies that quite high zoom factors can be profitably used, and thus also correspondingly high ZF levels.

Shape of *k*-Space

What do these results say about the shape of *k*-space? The pixelation artifact/signal map is radial in *k*-space, so on this basis, a circular *k*-space would be preferable over a square one. However, as higher levels of zero-filling are applied, the shape of *k*-space has progressively less effect on pixelation artifacts (Table 1). The *k*-space shape is not insignificant, but the level of zero-fill is a more dominant effect.

Another effect however is that a noncircular *k*-space exhibits a directionally dependent spatial resolution. In the absence of any prior knowledge of the object structure, a circular *k*-space is the unbiased choice. However, if a higher resolution is desirable in a certain direction, then this direction can be profitably aligned with the corners of a square *k*-space. This becomes a real issue when an object has anisotropic high resolution structure, such as a bar resolution phantom.

CONCLUSIONS

A pixel visible in an image as a small square patch is by definition an artifact because such patches do not exist in the object. Since there is nothing in the acquired *k*-space Fourier data requiring such a display mode, it is a mistake to display the DFT of an acquired data matrix as a 2D patchwork array of small squares. The problems that arise with such pixelated displays are an image with both surprisingly large visual attenuation of higher spatial frequencies (58.1% in 2D) and also introduction of artifactual pixel structure. The solution is straightforward: zero-fill until the pixels are no longer visible.

A second conclusion is that zero-filling (or perhaps data modeling of missing data) is always necessary. To visualize the important high resolution object structure, which is present in the variation between adjacent image domain data points (obtained by DFT from *k*-space data), the image should be zoomed sufficiently for this "pixel-level" structure to be visible. However this will result in the patchwork effect unless zero-filling is performed. So to simultaneously view high resolution information—but without visualizing pixel artifacts—zero-filling is necessary. The results support the use of a circular *k*-space.

ACKNOWLEDGMENTS

The authors thank Uta Frankenstein for providing the in vivo image data.

REFERENCES

1. Magn Reson Med 2010; 63(1).
2. Thevenaz P, Blu T, Unser M. Interpolation revisited. IEEE Trans Med Imag 2000; 19:739–758.
3. Kirshner H, Porat M. On the role of exponential splines in image interpolation. IEEE Trans Image Process 2009; 18:2198–2208.

4. Lehmann TM, Gonner C, Spitzer K. Survey: interpolation methods in medical image processing. IEEE Trans Med Imaging 1999; 18(Copyright 2000, IEE):1049–1075.
5. Carrato S, Tenze L. High quality 2 image interpolator. IEEE Signal Proc Lett 2000; 7:132–134.
6. Barthold E, Ernst RR. Fourier spectroscopy and causality principle. J Magn Reson 1973; 11:9–19.
7. Lindon JC, Ferrige AG. Digitisation and data processing in Fourier transform NMR. Prog NMR Spectroscopy 1980; 14:27–66.
8. Smith MR, Nichols ST. Efficient algorithms for generating interpolated (zoomed) MR images. Magn Reson Med 1988; 7:156–171.
9. Mitchell DK, Nichols ST, Smith MR, Scott K. The use of band-selectable digital filtering in magnetic resonance image enhancement. Magn Reson Med 1989; 9:353–368.
10. Constable RT, Kay I, Smith MR, Henkelman RM. High quality zoomed MR images. J Comput Assist Tomogr 1989; 13:179–181.
11. Du YP, Parker DL, Davis WL, Cao G. Reduction of partial-volume artifacts with zero-filled interpolation in three-dimensional MR angiography. J Magn Reson Imaging 1994; 4:733–741.
12. Parker DL, Du YP, Davis WL. The voxel sensitivity function in Fourier transform imaging: applications to magnetic resonance angiography. Magn Reson Med 1995; 33:156–162.
13. Eddy WF, Fitzgerald M, Noll DC. Improved image registration by using Fourier interpolation. Magn Reson Med 1996; 36:923–931.
14. Bernstein MA, Fain SB, Riederer SJ. Effect of windowing and zero-filled reconstruction of MRI data on spatial resolution and acquisition strategy. J Magn Resonance Imaging 2001; 14:270–280.
15. Elgavish RA, Twieg DB. Improved depiction of small anatomic structures in MR images using Gaussian-weighted spirals and zero-filled interpolation. Magn Reson Imaging 2003; 21:103–112.
16. Lin C, Bernstein MA. 3D magnetization prepared elliptical centric fast gradient echo imaging. Magn Reson Med 2008; 59:434–439.
17. Ferreira P, Gatehouse P, Kellman P, Bucciarelli-Ducci C, Firmin D. Variability of myocardial perfusion dark rim Gibbs artifacts due to sub-pixel shifts. J Cardiovasc Magn Reson 2009; 11:17.
18. Brigham EO. The Fast Fourier Transform and its Applications.Oppenheim AV, ed. Upper Saddle River, NJ: Prentice Hall; 1988.

BIOGRAPHIES

Xiaolu (Iris) Zhu is currently working as a telecom engineer in SNC-Lavalin. She received the Bachelor of Engineering degree in Electrical Engineering in 2006 from Beijing University of Posts and Telecommunications, Beijing, P.R. China. During the period of 2007–2008, she obtained the Master of Science degree in Electrical Engineering from the University of Calgary, Alberta, Canada. Her research interests are digital signal and image processing, biomedical imaging.

Boguslaw Tomanek is Head of the MR Technology Group at the National Research Council of Canada, an Adjunct Professor in the Department of Clinical Neurosciences at the University of Calgary. He received an M.Sc. in Physics from the Jagiellonian University, Krakow, Poland (1988) and Ph.D. in Physics from the Institute of Nuclear Physics, Polish Academy of Sciences (1995). His research comprises development of MRI technologies and multi-modal molecular imaging of cancer. His particular focus is on low field and inexpensive MRI systems.

Jonathan Sharp is currently a senior research officer at the National Research Council of Canada. He has degrees from University of Cambridge, University (1984, B.Sc.) of Aberdeen (1985, M.Sc.) and the University of London (1989, Ph.D.). He further trained at the University of Nottingham. His research interests include the basic mechanisms of spatial encoding used in magnetic resonance imaging.