

MRI Inter-Slice Reconstruction using Super-Resolution

H. Greenspan[§] G. Oz[§] N. Kiryati[§] S. Peled[¶]

[§]Faculty of Engineering
Tel Aviv University
Tel Aviv 69978, Israel

[¶] General Electric Medical Systems
Tirat Hacarmel 39120, Israel

August 1, 2001

Corresponding author:

Dr. Hayit Greenspan
Department of Biomedical Engineering
Faculty of Engineering
Tel Aviv University
Tel Aviv 69978, Israel

Phone: +972-3-6407398
Fax: +972-3-6407939
email : hayit@eng.tau.ac.il

Abstract

MRI reconstruction using super-resolution is presented and shown to improve spatial resolution in cases when spatially-selective RF pulses are used for localization. In 2-D multislice MRI, the resolution in the slice direction is often lower than the in-plane resolution. For certain diagnostic imaging applications, isotropic resolution is necessary but true 3-D acquisition methods are not practical. In this case, if the imaging volume is acquired two or more times, with small spatial shifts between acquisitions, combination of the data sets using an iterative superresolution algorithm gives improved resolution and better edge definition in the slice-select direction. Resolution augmentation in MRI is important for visualization and early diagnosis. The method also improves the signal-to-noise efficiency of the data acquisition.

Keywords: single-shot echo-planar imaging (EPI), magnetic resonance angiography (MRA), fast spin echo (FSE), signal-to-noise ratio (SNR), superresolution, thin slices.

1 Introduction

In cases in which true 3-D image acquisition is not effective or possible in MRI, it is common practice to acquire a set of 2-D slices. Such is often the case, for example, in T_2 -weighted imaging, diffusion-weighted imaging, and MR angiography. These are all imaging techniques that are important for early medical diagnosis and visualization purposes, and usually require coverage of extensive 3-D volumes in the imaged subject. The problem, as illustrated in Figure 1, is that a set of 2-D slices does not give a good isotropic 3-D image. MRI slices are usually thick, due to hardware limitations coupled with pulse sequence timing considerations. This results in resolution that is high in-plane and is low in the slice-select (or “through-plane”) direction. Data acquisition in all three orientations, would not provide self consistent 3-D data sets viewable in any oblique orientation.

In this work we address the challenge of achieving high-resolution, isotropic 3-D MRI images from 2-D MRI slices. A technique based on super-resolution is proposed in which *several sets* of 2-D slices are merged, to provide the high-resolution 3-D image. The method we present consists of two stages. First, the acquisition of a small number of multislice data sets, each volume shifted by a sub-pixel amount in the slice-select direction with respect to the other volumes. Second, the use of super-resolution post-processing in the inter-slice (z) dimension. Experimental validation of the technique shows that the resolution in the z direction is significantly improved. We show also that the SNR efficiency (SNR per unit acquisition time) of image data sets reconstructed using super-resolution is better than

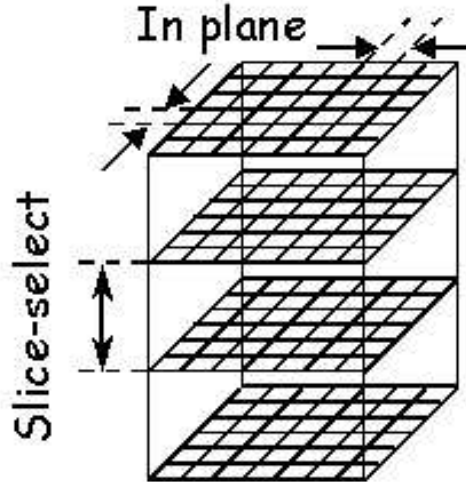


Figure 1: MRI slice acquisition. The resolution in the slice-select direction is usually much lower than in the in-plane directions.

that of images acquired directly with higher spatial resolution in the slice-select direction. Preliminary results appeared in a short abstract [4].

The structure of the paper is as follows: In section 2 a brief overview of MRI is given, with a specific focus on 2-D vs. 3-D MRI spatial encoding techniques. Section 3 discusses super-resolution algorithms. In section 4 the application of super-resolution algorithms to MRI data sets is investigated. Experimental results using echo-planar and fast spin-echo imaging methods are given in section 5. Computational aspects and quantitative performance evaluation are discussed in section 6 and conclusions are drawn in section 7.

2 Background

Magnetic resonance imaging utilizes the well timed application of magnetic field gradients for spatial localization of the particles contributing to the NMR signal. The Larmor equation (Eq. 1) states that the precession frequency of nuclear spins, ω , is linearly proportional to the magnetic field, B :

$$\omega = \gamma B \tag{1}$$

where γ is termed the gyromagnetic ratio of the nucleus in question. The two main methods of encoding the spatial location of the spins in MRI are:

1. Radio frequency (RF) pulses, applied simultaneously with the field gradients, excite a certain frequency spectrum corresponding (through the Larmor equation) to the position of the spins in space.
2. Temporal variation of the magnetic field gradients before or during signal reception causes spatially dependent frequency and phase modulations of the MR signal.

Conventional two-dimensional multi-slice MR imaging usually relies on the selection of 2-D planes (or slices) using method #1, and on in-plane encoding (method #2) using frequency encoding of the signal in one direction and phase encoding in the other in-plane direction. Usually each slice requires multiple RF excitations to be applied to it in order to be fully encoded. For reasons of pulse sequence timing, and signal-to-noise ratio (SNR), there are disadvantages to thin-slice excitation. RF pulses are finite in duration, thus a perfectly rectangular slice excitation is not feasible. The thinner the slice that should be excited by the RF pulse (i.e. the narrower the frequency band), the sharper the edges of the slice excitation profile are required to be. This translates to longer pulse durations, given the same magnetic field gradient strength. Thin slice excitation is particularly problematic when spectral-spatial pulses (used for fat suppression) are applied. Long slice-selective pulses increase echo times which increases the signal loss caused by dephasing of the magnetization.

Three-dimensional acquisition techniques provide thin sections and the ability to view interpolated data at high resolution from any angle. In true 3-D image acquisition, the third dimension is also phase encoded, and the RF excitation pulses excite a thick *slab* defining the whole imaging volume. Use of a short echo time is possible because of the short duration of the minimally-selective RF excitation pulses. T_1 -weighted imaging is compatible with 3-D acquisitions, and, in this case, a short repetition time improves the image contrast. True 3-D acquisition may not be advisable in other imaging cases, such as in T_2 -weighted imaging, MR angiographic (MRA) imaging and diffusion-weighted imaging.

True T_2 -weighting is difficult to obtain in reasonable imaging times by 3-D acquisition methods, although progress is being made in this direction using fast-spin-echo (FSE) methods [6]. The problem stems from the need for relatively long signal recovery between excitations to enable operation of the spin-echo mechanism that provides T_2 contrast. Since all the spins are excited by every pulse, the recovery time cannot be utilized and the sequence

takes a long time.

Multi-slab 3-D Fourier imaging in which a number of slabs are excited and each one is encoded by 3-D techniques attempts to address the problem of wasted time in 3-D T_2 imaging, but creates artifacts of its own. The artifacts include signal deterioration at slab edges, and increased point-spread function and truncation artifacts due to the small number of Fourier phase encodes in the slab-select dimension. A proposed solution to the problems of 3-D multi-slab encoding involves the use of non-Fourier methods such as Hadamard wavelets for encoding the third dimension [3].

Time-of-Flight (TOF) MR angiography (MRA) is another popular application that does not lend itself easily to 3-D acquisition. Spin saturation of the smaller vessels with slower flow results in signal drop off from these vessels.

For the important MRI application of diffusion imaging, no 3-D technique exists. Sequences that acquire raw data pertaining to the same slice or volume over many excitations cannot be modified to provide diffusion-weighted contrast because of motion artifacts resulting from phase inconsistencies in the data. Today, the most popular solution is to base diffusion imaging on 2-D single-shot techniques - mostly echo-planar imaging (EPI). Even the in-plane resolution here is sub-optimal, but usually better than the resolution in the third dimension. Isotropic resolution is particularly important in diffusion tensor imaging - the technique used for delineating white matter fiber tracts in the human brain *in vivo* [1].

In cases in which 3-D acquisition is not possible, it is common practice to use a set, or sets, of 2-D slices. In this work we propose a novel approach for achieving high-resolution isotropic 3-D images by merging sets of 2-D slices.

Newly developed parallel imaging techniques (see, for example, Ref. [9] and references therein) will allow faster acquisition, or higher in-plane resolution, at the expense of SNR. Increasingly high commercial gradient strengths and slew rates generated by local gradient coils will do the same. High resolution in plane will call for thin-slice acquisition which could bring the SNR down to values requiring signal averaging in order to provide clinically useful images. The ability to use super-resolution post-processing of thick slices, as shown in this work, might in some cases provide a much needed extra boost to the SNR.

3 Super-Resolution algorithms

Super-resolution algorithms are a family of techniques for creating a high resolution image from several lower resolution images of the same scene, taken from slightly different view-points. The first to address the problem have been Huang and Tsai [5]. They were followed by Ur and Gross [13], Irani and Peleg [7], Kim *et al* [8], Tekalp *et al* [12] and others. Super-resolution is a very active research area, motivated by emerging video technologies (e.g. [10]).

A general model of an imaging system is given in Figure 2. A high-resolution scene undergoes a geometrical transformation and optical blur. The low-resolution image is the outcome of sampling and additive noise. The principle of super-resolution algorithms is to acquire several low-resolution images $\{g_k\}$, with slightly different imaging conditions, and then estimate the high-resolution source, f , that best explains the low-resolution data. A variety of approaches can be found in the super-resolution literature.




Figure 2: Linear imaging system

3.1 The Irani-Peleg Algorithm

The algorithm we use is based on the iterative back-projection (IBP) method of Irani and Peleg [7]. A high resolution image is constructed iteratively to best explain the given data set of low resolution images. This is accomplished by minimizing the differences between the given low resolution images, and low resolution versions of the high resolution image. The latter are obtained by simulating the imaging process, taking geometric transformations and blur into account. Formally, if f is the high resolution image we wish to reconstruct, and $\{g_k\}_{k=1}^K$ are the given set of low resolution images, the imaging process can be modeled by:

$$g_k = (T_k(f) * h) \downarrow s + \eta_k \quad (2)$$

where T_k describes the geometric transformation between the k -th image and the reference frame, h corresponds to the imaging blur kernel, $\downarrow s$ is a down-sampling operator that decreases the resolution by the factor s , and η_k is an additive noise term. A flowchart describing the Irani-Peleg algorithm is shown in Figure 3.

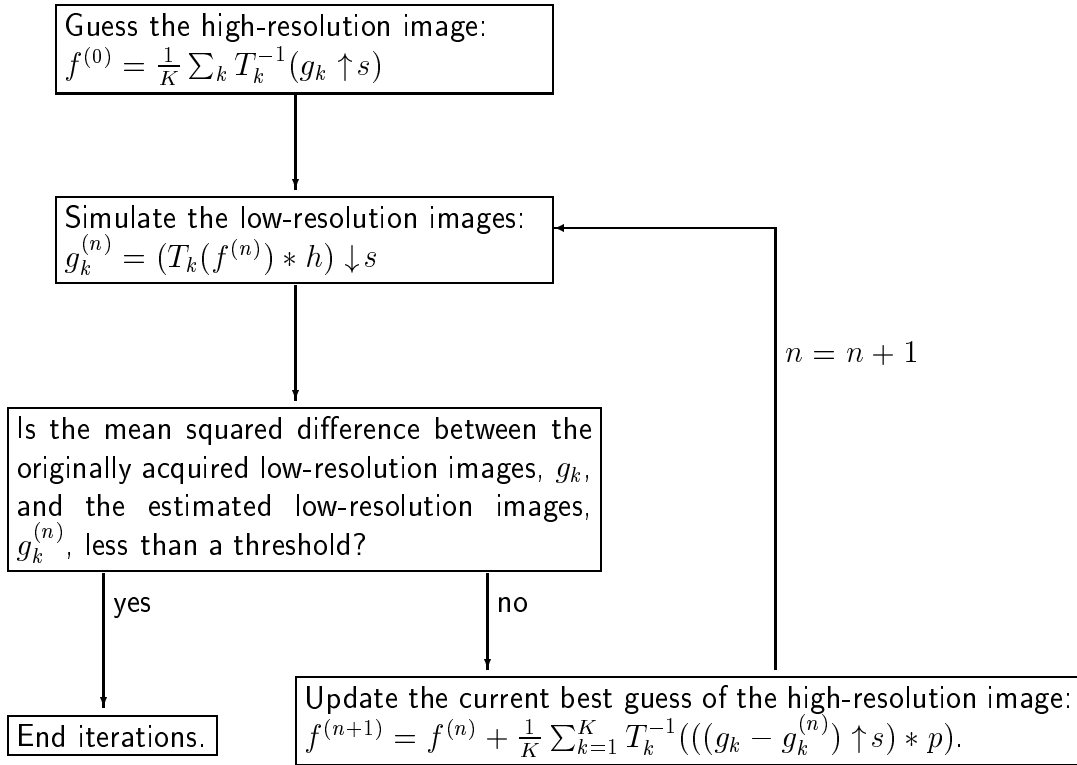


Figure 3: Block diagram of the Irani-Peleg superresolution algorithm.

In the Irani-Peleg algorithm, p is a “back-projection” kernel (see Figure 3). In order for the algorithm to converge, p should be chosen so that the following condition holds: $\|\delta - h * p\|_2 \leq 1$ [7]. Stability considerations may suggest that p be other than the inverse of h , at the expense of a slower convergence rate.

4 Super-resolution in MRI

Applying the super-resolution principle to MRI raises several fundamental as well as technical issues. It turns out that the differences between an MRI system and a common camera hinder the straightforward extension of the super-resolution algorithm to 3-D for use in MRI. These and other issues specific to the application of super-resolution in MRI are presented in this section.

4.1 Super-resolution dimensionality

Super-resolution, for images and image sequences, has customarily been treated as a 2-D problem. Note that super-resolution can also be applied to 1-D signals; the feasibility of applying a 1-D super-resolution procedure to 2-D images, improving resolution in one dimension at a time, depends primarily on the separability of the physical 2-D blur kernel. In our investigation of 3-D MRI data-sets we explored the possibility of extending super-resolution methodologies to 3-D. The 3-D super-resolution methodology is a novel approach for treating volumetric data. We introduce the general 3-D methodology next, and pursue an investigation of its applicability to MRI data-sets in the following sub-section.

The concept of obtaining a high resolution 3-D image volume from a set of low resolution 3-D image volumes is different from the extraction of a 2-D still image from a sequence of 2-D video frames, or even from the incremental generation of high resolution video frames from a low resolution input sequence, as suggested by Elad and Feuer [2]. The difference is the pooling, in true 3-D super-resolution, of several 3-D data sets to simultaneously obtain an entire 3-D higher resolution output set. An illustration of the acquisition of voxel-shifted data sets is shown in Figure 4. The geometric relationship between low and high resolution frames in 3-D sampling is shown. The large transparent boxes represent voxels in the low-resolution acquired data. The shaded box shows the position of a single high-resolution voxel within each of the low-resolution data sets. In the presented example, a voxel in the high resolution image volume is “covered” by 8 low resolution voxels.

The 3-D methodology is applicable to cases where:

1. Data acquisition covers a 3-D volume;

2. Successive data acquisitions add new information in each dimension.

The extension of the Irani-Peleg algorithm from 2-D to 3-D is conceptually straightforward, in principle requiring only 3-D models of the blur kernel h , the back-projection operator p , the individual geometric transformations T_k and their inverse. Note that 2-D images (or 1-D “needles”) can be extracted from the 3-D image.

4.2 In-plane vs. inter-slice investigation

An important issue for consideration when applying super-resolution to MRI data, is the dimensionality of the problem. An examination of Fourier-encoded MRI data sets reveals distinct characteristics of the in-plane vs. through-plane encoding. These differences impose constraints on the dimensions in which super-resolution can lead to useful results.

Fourier-encoded in-plane MRI data is inherently band-limited. This is due to the time limit of the acquisition mechanism and the fact that the information is gathered in the frequency domain (k-space). The spatial frequencies in the z (inter-slice) direction exhibit a less sharp cut-off. There is sufficient information in the z dimension such that sampling of the data in that direction results in aliasing, and thus provides the basis for using a super-resolution algorithm in enhancing the resolution.

In order to verify that the above statements do in fact reflect the reality of MR imaging, we acquired multislice 2-D image data sets shifted in all three spatial directions by half a voxel (as shown in Figure 4). A 3-D iterative super-resolution algorithm was applied. In Figure 5, we see an original low-resolution image (top left) with its original power spectrum (top right). This image is parallel to the $y - z$ plane, where y is one of the in-plane coordinates, and z is the slice-select direction. The output of the super-resolution process (double size in each dimension) is shown (bottom left), with its power spectrum (bottom right). The sharp frequency cut-off in the y direction is evident. A comparison between the low-resolution spectrum and high-resolution spectrum shows that the attempt to augment the in-plane image resolution is equivalent to zero-padding extension of the frequency domain. A spreading-out of the power-spectrum is evident in the z direction.

The investigation leads to the following conclusions: First, without prior information about the data, the prediction of higher spatial frequencies is not possible for Fourier-encoded




Figure 4: Acquisition of voxel-shifted data sets. The geometric relationship between low and high resolution frames in 3-D sampling of MRI data. The large transparent boxes represent voxels in the low-resolution acquired data. The shaded box shows the position of a single high-resolution voxel within each of the low-resolution data sets. The Field of View (FOV) defined by the acquisition parameters on the MR scanner is shifted in this case by half a voxel in each direction. Thus, a voxel in the high resolution image volume is “covered” by 8 low resolution voxels.

MR data. The best that can be done in the in-plane (x and y dimensions) is to interpolate, via zero-padding, the given data to the desired resolution. Previous results that show improved images after application of super-resolution to the in-plane dimensions [11] may be, in most part, replicated by zero-padding. An important second observation is that in the inter-slice direction, sub-voxel spatial shifts can in fact be utilized to increase the resolution. It is in this dimension that we apply the super-resolution algorithm¹.

Figure 6 shows the sub-voxel spatial shifts in the slice-select direction, as used in this work. The input low-resolution voxels are shown left, with the high-resolution voxels via sub-voxel shifts, shown right. Experiments are conducted with two, three and four shifts per voxel, enabling a substantial augmentation in the slice-select resolution. The number of shifts

¹In principle, with non-Fourier encoding, other dimensions should also be amenable to enhancement through super-resolution, although we have not investigated this route.




Figure 5: Spectrum analysis, (y, z) plane. The horizontal axis is the slice-select direction. Top row: Low-resolution input (left), spectrum of low-resolution input (right). Bottom row: High-resolution output (left), spectrum of high-resolution output (right).




Figure 6: Acquisition with 1-D voxel shifts in the slice-select direction. Low-resolution voxels (left). High-resolution voxels (right).

is determined so as to achieve (as much as possible) isotropic resolution in all dimensions.

4.3 Signal-to-noise ratio

When considering a method for resolution improvement in MRI, one should make sure that the signal-to-noise ratio (SNR) is not compromised. The signal-to-noise ratio (SNR) of an MR image is often measured by taking the mean of a high-intensity region of interest and dividing by the standard deviation of a region of noise outside the imaged object. For the purpose of our comparison between reconstruction methods, this definition is satisfactory.

RF excitation of a volume or slice leaves the affected nuclear spins in a state of partial saturation. Multi-slice 2-D imaging takes advantage of the time required for signal recovery after RF excitation in order to excite and encode many slices in an interleaved fashion. The time constant associated with this signal recovery is referred to as T_1 . Thus increasing volume coverage by adding more slices increases the acquisition time approximately in proportion to the number of slices, but may also improve the SNR due to a longer repetition time (TR) between individual excitations of a particular slice. However, when many slices need to be acquired, as is often the case in clinical settings, the incremental improvement in SNR of adding more slices (thus adding time between consecutive excitations of the same slice) is insignificant. In such cases, serial acquisition of 2 or 3 sets of a smaller number of thicker slices with correspondingly shorter repetition times, as described in this work, does not negatively affect the spin recovery. In such cases we find a clear SNR advantage to using our method of data acquisition with super-resolution post-processing.

4.4 Parameterization

Successful implementation of the super-resolution algorithm requires an estimate for the transformation and blur parameters that most closely correspond to the true imaging system parameters.

The point-spread function (PSF), or blur, h , in the slice-select direction can be inferred from the slice excitation profile. We measured typical slice profiles and found them to be well approximated by Gaussian functions, where the full width at half maximum (FWHM) was the originally selected slice width. In our work we experiment with two PSFs. The first is a rectangular pulse PSF (a crude estimation for the slice profile), hereon termed *Box-PSF*. The box width is taken as the selected slice width, in the desired high-resolution pixel units.

The second is a Gaussian PSF (*Gaussian-PSF*), with FWHM set to the selected slice width. In the implementation of the Irani-Peleg algorithm, the h filter is taken as one of the above two PSFs. In both cases, the p filter is taken as an impulse function. The choice for p satisfies the Irani-Peleg requirement for convergence.

A key requirement for successful implementation of super-resolution algorithms in general is precise image registration accurate to a small fraction of a pixel and capable of bringing all the input images to a common reference frame. The difficulty here depends on the number of degrees of freedom in the inter-image transformations that have to be accommodated. In the Irani-Peleg algorithm, the geometric transformations T_k need to be known to sub-pixel accuracy, preferably to 1/10th of a pixel. Then the iterative process converges rapidly and provides good results. The threshold used as a stopping condition in the Irani-Peleg algorithm was 2% (see Figure 3).

In our case, the accuracy and stability in the field of view of MRI machines allows us to acquire a set of images of a subject, each translated by a predefined sub-voxel vector with respect to the 3-D reference frame. Depending on the physical orientation of the slice-select direction, the shift will be implemented either by shifting the patient bed, or by moving the RF transmitter center frequency. In the super-resolution algorithm, this eliminates the need for registration, since all the geometric transformations T_k (and their inverse) are known in advance. We are provided with the keys necessary for successful super-resolution processing: all image volumes are brought to the same 3-D reference frame reliably, accurately and without any computational effort. It should be noted that we are assuming minimal movement of the subject. For the more general case of a moving imaged subject, the method proposed herein would also benefit from image registration.

5 Experimental Results

In order to evaluate the inter-slice super-resolution methodology in MRI, we carried out a number of experiments:

- *Phantom* experiments, in which acquired MRI data sets of objects and fruits are used, thus assuring no subject motion.

- *MR brain imaging* experiments, in which we use human brain data as input to the super-resolution algorithm.

All imaging was performed with an RF head coil on either a 1.5 Tesla General Electric Signa Echospeed MRI system or a 3 Tesla General Electric MRI system.

We present a sample result from each of the experimental paradigms. The presented results include visual comparisons between the input low-resolution image (zoomed to the desired size of the high-resolution result), the low-resolution image following Sinc interpolation (zero-padding), the result of interleaving the sets of low-resolution images and the result of using the super-resolution algorithm on the low-resolution image set.

Interleaving is a method to achieve a high-resolution image from a set of shifted low-resolution images by combining the pixels, one by one from alternating low-resolution image inputs, to generate a single large image. This method will give best results when the PSF of the sensor is the same size as the high-resolution pixel. In our case, the PSF is much bigger and its width is approximated by the size of the low-resolution pixel in the z direction.

5.1 MRI resolution phantom

The homemade phantom consists of long thin plastic partitions (“teeth”), lodged in a plastic block, placed 4mm apart, surrounded by Gd-DTPA-doped water. The imaging sequence consists of multislice fast spin-echo (FSE) with 16 slices, 3mm thick, approximately parallel to the plastic partitions. Three sets of multislice data were acquired, with 1mm shifts in the slice-select direction. The low-resolution input voxel size is $1 \times 1 \times 3$ mm. Following the super-resolution procedure, an output voxel will be a 1mm isotropic cube.

Figure 7 shows the results of super-resolution applied to the comb-phantom MRI data. The horizontal axis is the slice-select direction. In (a), the original low resolution data is shown, followed by zero-padding interpolation (b), interleaving (c) and super-resolution results in the inter-slice direction (d,e). The super-resolution results include a box-PSF (d) and a Gaussian-PSF (e).

Several observations may be made from the presented results. The visibility of the comb teeth has greatly improved by using super-resolution rather than zero-padding interpolation. Moreover, super-resolution brings out more information than interleaving. The super-




Figure 7: Super-resolution applied to comb-phantom MRI data. The horizontal axis is the slice-select direction. Shown are: original low resolution data (a); zero-padded data (b); interleaved slices (c); inter-slice direction improved by super-resolution (Box-PSF) (d); inter-slice direction improved by super-resolution (Gaussian-PSF) (e).

resolution brings out features that are inseparable in the source images. The implementation of the super-resolution algorithm with a Gaussian-PSF (e) gives better results than when using a Box-PSF (d). The comparison between the two PSFs is interesting: better estimation of the high resolution image is achieved by using a blurring filter, h , that more closely matches with the MRI imaging system and the MRI image characteristics in the slice-select direction.

5.2 Fruit

Similar results have been generated from a variety of fruits. Figure 8 shows images of a pomegranate. The imaging sequence was single-shot echo-planar (EPI), with 4 shifts in the slice direction. The slice thickness was 4mm, the in-plane resolution was 3 mm and the number of slices was 42. Again, the horizontal axis is the slice-select direction. Applying the super-resolution procedure in the slice-select direction, will provide us with a non-isotropic resolution, of $3 \times 3 \times 1$ mm voxels. Note that in this case the imaging system provides limited resolution in the in-plane dimensions as well as in the slice-select. In order to achieve isotropic voxels we use Sinc interpolation (zero-padding) in the in-plane dimensions and super-resolution in the slice-select dimension, for an output of 1 mm in all dimensions.

From top to bottom, left to right, are the original low resolution data (a), zero-padding interpolation (b), interleaving (c) and super-resolution result (d). We witness the resolution limitations in the z -direction, as exist to-date, along with the augmented visualization possible using a super-resolution approach. Many seeds in the fruit become visible only after the super-resolution processing.

5.3 Human brain

In the final example, we show super-resolution on human brain data. In this case we set the slice thickness wide enough in the low-resolution data acquisition so that we could also physically acquire high-resolution data with the same slice width as the super-resolution result. Our goal is to achieve a result that is as close to the high-resolution image as possible, using the low-resolution input and the super-resolution algorithm. The imaging sequence was fast-spin-echo (FSE) with 3 shifts in the slice direction. The slice thickness was 4.5 mm,




Figure 8: Pomegranate image: (a) original low resolution data; (b) zero-padding interpolation; (c) interleaved slices; (d) inter-slice direction improved by super-resolution.

the in-plane resolution was 1.5 mm and the number of slices was 22. The high-resolution data acquired is FSE with isotropic resolution of 1.5 mm.

Figure 9 shows the low-resolution image (a), zero-padding interpolation (b), super-resolution result with Box-PSF (c), super-resolution result with Gaussian-PSF (d), and the high-resolution image (e). There is a clear improvement in the resolution of the images as we progress from the low-resolution input to the Gaussian-PSF super-resolution result. We note that from careful visual inspection, there is still a difference in image quality between the super-resolution result (d) and the high-resolution acquisition (e). Nevertheless, the super-resolution result is close to the high-resolution goal and is seen to be consistent with it.




Figure 9: Human brain MRI. Shown are the low-resolution image (a), zero-padding interpolation (b), super-resolution result, Box-PSF (c), super-resolution result, Gaussian-PSF (d) and the high-resolution image (e).

6 Image Analysis - Resolution and SNR Efficiency

In addition to visual evaluation of the output images, quantitative measures of resolution and signal-to-noise ratio (SNR) are computed and used to evaluate the performance of the super-resolution algorithm.

6.1 Spatial Resolution

Quantitative comparison of resolution is not a trivial process. Currently we are basing our resolution estimate on the measurement of edge widths. A sample of edges are selected in one image, and then consistently compared across all the images. The width of each edge is measured by least-squares fitting it to a sigmoid function of the form:

$$y(x) = \frac{1}{1 + \exp(-a(x - c))} . \quad (3)$$

The parameter a is inversely proportional to the width, and c corresponds to the center location. Following the fitting step, a measure of “rise length” is computed, defined as the width (in high-resolution pixels) from 10% to 90% of the edge height. It is easy to show that:

$$\text{width}[\text{pixels}] = \frac{4.4}{a} . \quad (4)$$

For the quantitative analysis of super-resolution, we used an apple as the input source (due to the well defined outer edges and relative lack of structure close to those edges). Imaging was performed using the FSE sequence with an echo train length of 16, a slice width of 4.5 mm with no gap between them, and 3 shifts differing in spatial location by 1/3 of a pixel. The number of slices in each low resolution data set was 22 and the repetition time between RF excitations of the same slice was 8 sec. A high resolution data set of 66 contiguous slices covering the same volume was also acquired. We compared the width of 20 outer edges of the apple for 5 different data sets: (a) the input low resolution images following sinc interpolation (zero padding), (b) interleaved result, (c) super-resolution result, Box-PSF, (d) super-resolution result, Gaussian-PSF and (e) high-resolution source. A sample edge from the apple image is shown in Figure 10, along with the corresponding sigmoid function that was fitted to it, for each of the cases compared. Results are summarized in Table 1.




Figure 10: A sample edge from the apple image and the corresponding sigmoid function fitted to it, for each of three cases compared. The squares, circles and triangles show points on the sample edge, for zero-padding interpolated low-resolution image, super resolution image and high resolution image, respectively. The dashed, solid and dotted lines show the sigmoid corresponding to the low-resolution, super resolution and high resolution data, respectively.

	<i>Zero-padded</i>	<i>Interleaved</i>	<i>Super-resolution Box-PSF</i>	<i>Super-resolution Gaussian-PSF</i>	<i>High resolution</i>
<i>Acquisition time [min:sec]</i>	1:28	4:24	4:24	4:24	4:00
<i>SNR</i>	287	276	170	124	95
<i>Edge widths [pixels] (7 of 20)</i>	2.9	3.0	2.5	1.9	1.9
	3.9	3.7	3.2	2.6	2.3
	3.7	3.5	2.6	2.0	2.3
	3.5	3.8	3.3	2.8	2.7
	3.3	2.8	2.1	1.8	1.0
	4.0	4.2	3.1	2.4	3.1
	3.5	3.8	3.1	2.3	1.8
	:	:	:	:	:
<i>Mean [pixels]</i>	3.7	3.7	2.9	2.2	2.3
<i>Comparison to High Resolution</i>	+58%	+57%	+25%	-3%	0%

Table 1: Quantitative measures of SNR and resolution. SNR values are given as absolute ratio values. Resolution is measured as edge widths in the z direction. Results are given for an apple input. The edge widths in each row correspond to the same edge. In the two last rows the mean edge width in each data set is listed, as well as the percentage difference from the high-resolution data.

Several points may be learned from Table 1. Looking across the columns, we note that the resolution in edge width improves as we shift from the zero-padded input to the high-resolution source. A clear improvement is present in the edge-width of the super resolution result, in comparison to the zero-padded low-resolution source and the interleaved results. The super-resolution process with Box-PSF gives high-resolution results but not as high as a true high-resolution acquisition (in cases where such an acquisition is possible). The mean edge width of the Gaussian-PSF super-resolution result is (almost exactly) identical to the mean width of the edges in the high-resolution source. This statistic indicates a successful augmentation of the image resolution via the super-resolution procedure.

6.2 SNR Efficiency

The SNR values for the corresponding images decrease as expected with the increased resolution. As one collects more data in the MRI process (a long sequence), one obtains a result with higher SNR for the same resolution. The ratio between the SNR of the result and the time length of the data acquisition sequence is called *SNR Efficiency*. In an MRI process the goal is to obtain high resolution images with a high SNR efficiency. Table 1 shows the SNR obtained from each of the methods under comparison.

The sequence time of the high-resolution source was 4:00 minutes (=240 sec) and the sequence time for each of the low resolution data sets was 1:28 min which gives 4:24 minutes (=264 sec) for complete collection of three shifts. Division of the SNR by the image acquisition times gives a SNR efficiency of 0.47 sec^{-1} for the super-resolution result, and 0.40 sec^{-1} for the high-resolution acquisition. It is possible to state from these results that the super-resolution SNR efficiency is better than a high-resolution MRI collection (at the same effective resolution as measured by the edge widths).

7 Conclusion

This paper presents and demonstrates MRI inter-slice reconstruction using super-resolution. With inter-slice super-resolution we can effectively break limits on slice thickness posed by the physical properties of existing MR imaging hardware.

For large volume coverage using multislice 2-D imaging techniques, the acquisition of

multiple data sets does not add significant time to the data acquisition when compared to the acquisition of thin slices. In such cases, the SNR efficiency of this method appears to be better than 2-D thin slice acquisitions. From a purely theoretical point of view, the above finding can be explained by the following argument: acquisition of n sets of 2-D slices that are n times thicker results in a SNR per unit volume that is \sqrt{n} greater than that of a similarly timed total acquisition of one set of thin slices. In any case, super-resolution provides the flexibility to determine the extent of the trade-off between the high SNR of the thick slices, and the high anatomical accuracy of thin slices.

The overall spatial accuracy and stability in the field of view of MRI machines, in all three axes, is far better than their voxel size. This means that the imaged subject can be reliably and accurately shifted by *predetermined* sub-voxel steps between scans. Although for Fourier-encoded MRI data-sets we showed that super-resolution may only be usefully applied in the slice-select direction, a 3-D super-resolution methodology is suggested for general volumetric data-sets.

In conclusion, we are proposing a novel framework of using super-resolution algorithms in MRI resolution augmentation. The method is easily implementable because it is based on a simple reconstruction algorithm and on existing pulse sequences. In the MRI application domain super-resolution may be of significant potential importance.

Acknowledgments

This research was supported in part by the Ela Kodesz Institute for Medical Engineering and Physical Sciences, and by the Adams Super-Center for Brain Studies, Tel-Aviv University.

References

- [1] P. J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractography using DT-MRI data. *Mag. Reson. Med.*, 44:625–632, 2000.
- [2] M. Elad and A. Feuer. Super-resolution reconstruction of image sequences. *IEEE Trans. Patt. Anal. Machine Intell.*, 21:1646–1658, 1999.

- [3] G. Goelman. Fast 3D T₂-weighted MRI with Hadamard encoding in the slice select direction. *Magn. Reson. Imag.*, 18:939–945, 2000.
- [4] H. Greenspan, S. Peled, G. Oz, and N. Kiryati. MRI inter-slice reconstruction using super-resolution. In *Proceedings of MICCAI 2001, Fourth International Conference on Medical Image Computing and Computer-Assisted Intervention (short abstract)*, Lecture Notes in Computer Science. Springer, October 2001.
- [5] T. S. Huang and R. Y. Tsai. *Multi-frame image restoration and registration*, volume 1, pages 317–339. JAI Press Inc., 1984.
- [6] J. P. Mugler III, S. Bao, R. V. Mulkern, C. R. G. Guttmann, R. L. Robertson, F. A. Jolesz, and J. R. Brookeman. Optimized single-slab three-dimensional spin-echo MR imaging of the brain. *Radiology*, 216:891–899, 2000.
- [7] M. Irani and S. Peleg. Motion analysis for image enhancement: resolution, occlusion, and transparency. *J. Vis. Comm. Image Rep.*, 4(4):324–335, 1993.
- [8] S. P. Kim, N. K. Bose, and H. M. Valenzuela. Recursive reconstruction of high resolution image from noisy undersampled multiframes. *IEEE Transactions on Acoustics Speech and Signal Processing*, 38(6):1013–1027, 1990.
- [9] B. Madore and N. J. Pelc. SMASH and SENSE: Experimental and numerical comparisons. *Magn. Reson. Med.*, 45:1103–1111, 2001.
- [10] M.I. Sezan P.E. Eren and A.M. Tekalp. Robust object based high resolution image reconstruction from low resolution video. *IEEE Trans. on Image Processing*, 6:1446–1451, 1998.
- [11] S. Peled and Y. Yeshurun. Superresolution in MRI: Application to human white matter fiber tract visualization by diffusion tensor imaging. *Magn. Reson. Med.*, 2000.
- [12] A. M. Tekalp, M. K. Ozkan, and M. I. Sezan. High resolution image reconstruction for lower resolution image sequences and space-varying image restoration. *IEEE International conference on acoustics, speech, and signal processing (San Francisco, CA)*, March 23-26:III–169–172, 1992.

- [13] H. Ur and D. Gross. Improved resolution from subpixel shifted pictures. *CVGIP: Graphical Models and Image Processing*, 54:181–6, 1992.