

Tina Memo No. 1999-005

Short Version published in: JMRI, 10, 4, 582-588, 1999.

Renormalised Sinc Interpolation.

N.A.Thacker, A.Jackson, D.Moriarty and B.Vokurka.

Last updated
1 / 10 / 2000



Imaging Science and Biomedical Engineering Division,
Medical School, University of Manchester,
Stopford Building, Oxford Road,
Manchester, M13 9PT.

Abstract

It has long been established that the appropriate way of reslicing volume MR images is to use the method of sinc interpolation [2, 4]. We have recently needed to implement this method ourselves and have found, like other authors before us, that large convolution kernels are needed in order to produce accurate reslice data, suitable for subtraction. This requirement has led many groups to investigate the use of specialised hardware and software in order to perform data analysis within sensible timescales. However, we have found that the major component of the error introduced from interpolation with small kernels, is actually due to a first order normalisation problem introduced by truncation. In this paper we demonstrate the characteristics of this problem on real data and show how it can be eliminated, so that accurate reslice data can be obtained with small kernels. Unlike other recent suggestions for correcting such effects [3], the required changes in computation are simple and significantly reduce the processing requirement for a given interpolation accuracy.

Renormalised Sinc Interpolation.

There are several techniques that one can adopt to solve the problem of image interpolation. One is to assume a particular prior functional model for a local region of the image data, estimate the function parameters from a maximum likelihood metric and then recompute intermediate sites from the functional model. Using a least squares metric, and provided that there are fewer parameters than data points, it is always possible to determine the coefficients for this approach from a set of convolution kernels, (one per co-efficient) with a size given by the specified local region for the model. As these coefficients can be precomputed the computational requirement for a model of N parameters can be simply estimated from the cost of the convolutions and the functional evaluation required for the interpolated value. The execution time is dominated by the convolution process and one such process is required per parameter in the model. The problem with this approach is the difficulty in specifying a suitable functional model on a particular data set, this must be done so that an accurate fit to the data can be achieved with a limited number of parameters. Inaccurate modeling due to fit restrictions results in a smoothing process which may be undesirable. It is generally accepted that linear models are simply not sufficient due to an inability to model discontinuities in the data. This is particularly a problem for medical image subtraction algorithms. Quadratic or cubic models, which have reasonable modelling capabilities at least in 2-D, and relatively modest computational requirement are popular. Unfortunately, there is no generic model which will work well for all images.

In the area of image processing, image interpolation (for purposes of warping or reprojection) is often followed by a convolution process. In these cases it is perfectly acceptable to interpret the required shift of the image as an equivalent shift in a smooth kernel function. For rapid processing applications this is an attractive option as the interpolation process can be effectively eliminated altogether in computational terms. Shifted convolution is such a generically useful process that it has even been suggested as the basis for general purpose machine vision hardware [5]. If the convolution coefficients required at any subpixel location are stored in a discrete lookup table this entire process can be executed in the time required to perform a single multiply accumulate operation and can be quickly accelerated in hardware [5]. However, for medical image analysis the requirement is generally for high quality images which can be used to identify small changes in the image via subtraction or statistical analysis. Thus an additional convolution process is unnecessary and unwanted.

In the case of MRI, the images are initially acquired in the spatial frequency domain (or K -space). It is therefore natural to consider interpolation by reconstruction from the Fourier components. When this is done the result in the spatial domain is sinc interpolation. In some respects this is still a specific form of functional interpolation, where the assumed function form is a Fourier model. It would therefore be incorrect to believe that the images interpolated this way are actually what would have been obtained if the subject had been translated in the spatial domain to the required position. That would only be true if the original data were exactly defined by the truncated Fourier description in the original data set. However, from the point of view of MR image analysis, interpolation with a sinc function is at least consistent with the way the image has been acquired and no information is lost from what was available at the scanner, provided that the sinc kernel has a size equal to the image. For computational reasons it is inappropriate to apply such a large kernel, but as the kernel itself is quite spatially compact it is possible to truncate it at some level. An acceptable kernel size for medical subtraction work is $13 \times 13 \times 13$,

implying that it takes 6 pixels for the sinc function to reduce to an acceptable level. The computational implications of using a large scale kernels is offset somewhat by the fact that the coefficients required for sinc interpolation can be decomposed [2] so that the computation is linearly proportional to the kernel size rather than its volume. Thus despite its somewhat complicated form, sinc interpolation seems to offer as good a solution as the best techniques in terms of computational requirement on equivalent sized kernels.

Methods

Our starting point was the work of Hajnal [4]. We reimplemented the decomposable sinc interpolation algorithm with a cosine Hann (Hamming) window as described in their paper.

$$I(x, y, z) = \sum_X \sum_Y \sum_Z I(X, Y, Z) \cdot HS(x, X, R) \cdot HS(y, Y, R) \cdot HS(z, Z, R)$$

with

$$HS(a, A, R) = \frac{\sin(\pi(a - A))}{2\pi(a - A)} (1 + \cos(\pi(a - A)/(R + 1)))$$

The Hamming function eliminates problems with oscillatory effects at discontinuities and guarantees that the convolution coefficients fall off to zero at the edge of the sinc kernel (ie: at $a = R + 1$). We also confirmed that this process was the same as used in our version of Woods' reslice software [7], SPM [8] and the work of Bydder [1]. We therefore believe that this reslice technique is the one being used by the vast majority of the research community [6]. In Hajnal's work the fractional error on grey level value was computed after an image had been rotated with a forward followed by an inverse reprojection. We have essentially repeated these experiments, though we have used as our reference the sinc interpolated image from a large kernel. This is a valid test as effectively no information is lost when using a large kernel. However, using only one projection process allows us to associate residual errors directly to a single computation process. Our estimated relative fractional error differs by a factor of $\sqrt{2}$ from the previous work, exactly as expected if the errors introduced by the two stages of interpolation are uncorrelated [Figure 1a]. The data seems to improve by an approximate factor of two at each increase in kernel size. We also looked at the performance of a simple linear interpolation algorithm, where each voxel was estimated from the four nearest voxels in the original image. To our surprise, this linear interpolation worked much better than many of the large sinc kernels and an 11x11x11 kernel was needed in order to surpass its performance. Initially, we could think of no reason why a technique making use of so little data on such a small scale in the local region of the image could be so good in comparison to techniques which typically use a thousand values.

We worked with a range of typical NMR scans of the brain and skull and noted that the frequency distribution of fractional error estimates (the histogram) was not independent of scene content. In particular image sets which included a large grey level offset from the origin, before the dynamic range of the tissue data, had very poor residuals which exhibited systematic shifts in output value. In addition small shifts (fractions of a pixel movement across the whole image) could produce very strange and systematic interpolation errors, often producing beating effects across the image plane and sometimes mimicking the process of signal loss due to magnet inhomogeneity. On investigation we found that this position related error was due to the differing amount of truncation of the sinc function as a function of the subvoxel location used for interpolation. In addition we ran our most frequently used reslice software [7] and observed these same effects. To demonstrate this in a real application we have taken some images requiring registration on a Multiple Sclerosis study [Figure 2a]. This data has a dynamic range of 1000 greylevels with an estimated noise level 20(RMS). We have resliced the data according to the estimated alignment parameters using 4x4x4 and 14x14x14 kernels. The histogram of difference between the two reslice results is shown as an image below [Figure 2b]. A systematic offset in interpolation error is clearly evident. We also note that the use of even sized kernel introduces an asymmetry in the sinc truncation which is also undesirable. We would recommend that sinc interpolation be implemented with an odd sized kernel.

If sinc interpolation is used to reslice a completely uniform data block with constant grey level value then the resliced image will have an error on the output values proportional to the missing quantity of truncated kernel. This intuition suggests a natural way of correcting this effect. If the truncated sinc kernel were renormalised (scaled) to have a constant integral, this when applied to a uniform data set




Figure 1: Error and Time Plots

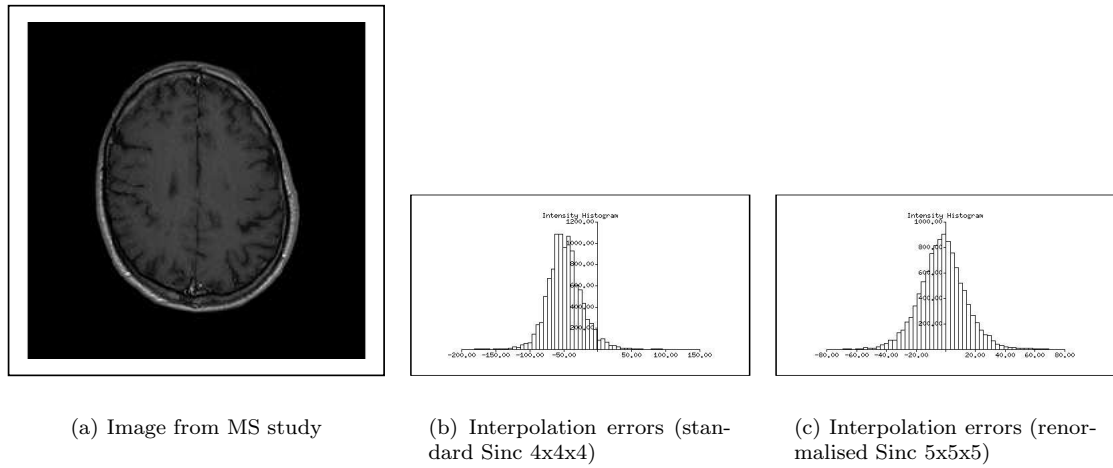


Figure 2: Interpolation Errors

must produce output data with identical values. This is statistically equivalent to modelling the mean bias using PCA analysis. It can be accommodated within the standard algorithms for sinc interpolation or the framework of the so-called 'fast-sinc' techniques [2]. It has to be said that in the image processing community such a renormalisation would probably be considered common practice. On an NMR image this renormalisation will completely remove the first order error on the estimated greylevel leaving only error terms due to the inadequacy in the assumed Fourier model. The renormalised kernel is simply computed by replacing HS in the original equation with;

$$HS'(a, A, R) = HS(a, A, R) / \sum_A HS(a, A, R) = HS(a, A, R) / |HS(a)|$$

The effect of this is most quickly computed as a renormalisation of the final greylevel result.

$$I'(x, y, z) = I(x, y, z) / (|HS(x)||HS(y)||HS(z)|)$$

Results

We repeated our performance study with the renormalised kernel for the previous set of data, which exhibited particularly poor asymmetry in its reslice error distribution. On the basis of the relative error measure, the interpolated images were now far more accurate and no longer display any of the systematic effects of the original technique. The general behaviour of the curve is again an approximate improvement of 2 at each increase in kernel size [Figure 1a]. However, the performance figures show an immediate improvement over linear interpolation even for a 5x5x5 kernel, with a performance for this kernel close to that obtained with the original 13x13x13 kernel. Results on data sets with no large offset produced a less striking improvement, typically only of the order of 20 % . However, with the new technique the absolute error on the interpolated data was far more consistent regardless of the data set. Indeed, with the first order error eliminated, remaining errors should be proportional to image contrast (ie: the dynamic range). As expected, the results seem to be comparable with those from the technique presented in [3], which is to appear in the latest version of the SPM package [8].

For the MS data we show here the histogram of residuals (from the 14x14x14) for the renormalised 5x5x5 kernel [Figure 2c]. The systematic shift in the reslice error, which would generally cause problems in later processing, has been eliminated resulting in data which is sufficiently accurate to be used in a subtraction study.

Discussion and Conclusions

During a re-implementation of the standard approach to NMR reslicing we have discovered that the major contribution to interpolation error on a large fraction of image data sets is due to position dependent sinc normalisation. We have established that this level of error is consistent with that reported in the literature in this area, both in terms of published figures and available software. From this study we conclude that the contribution to the error on the estimated pixel from standard sinc interpolation appears to be proportional to the output data value. Indeed it must be if the main cause is the difference in the interpolation kernel values from the ideal large kernel, as the error propagated via a convolution process will scale with the image. For renormalised sinc interpolation this component of error is no longer present. The level of improvement in percentage interpolation error is variable depending on the data set, but the resulting absolute values now seem to be less variable and better related to the dynamic range of the data. Given that most NMR sequences do not give a measurement accuracy of better than 1 % of the dynamic range we find that there is generally little point in interpolating data with a kernel much bigger than 5x5x5 using the renormalised sinc.

The time required for interpolation with the improved method is almost identical to equivalent sized kernels in the standard technique. Theoretically, for a normal sinc with precomputed kernel co-efficients, the execution time would scale with the volume of the kernel. For our implementation, without precomputation, this relationship becomes evident for large kernels [Figure 1b], (a 256x256 slice on a multi-user 167 MHz SUN Ultra sparc 1 and a Pentium II 300 PC). A comparison of using Wood's software gives a ratio of 4x4x4 to 14x14x14 reslice time of 30.6. The ability to use a 5x5x5 kernel in place of a 13x13x13 thus represents a computational saving factor of between 7 and 30.6 (equivalent to parallelising over of the order of twenty equivalent machines). Using the timing for the linear interpolator as the minimum time for setting up a loop over a slice of data and the per voxel convolution rate, we can estimate the time curve for the 'fast sinc' [Figure 1b]. The dominant contribution to this algorithm is expected to be the three loops over the data volume not the convolution stage. As a consequence this technique is not likely to be much faster than the standard techniques for small kernels. Indeed given the problems associated with large volume memory paging for the larger data sets, the theoretical improvement may be quickly lost. Also as the entire data set must be resliced to achieve the mean slice performance, the fast technique may not be justified in interactive applications where the user may only be interested in a few slices. In our software we also provide the facility to select a region for reslice, this typically gives another factor of 4 improvement.

References

1. G.M.Bydder, The Mackenzie Davidson Memorial Lecture: Detection of Small Changes in the Brain with Serial Magnetic Resonance Imaging, BJR, 68, 1271-1295, 1995.

2. P.A.Freeborough, R.P.Woods and N.C.Fox, Accurate Registration of Serial 3D MR Brain Images and Its Application to Visualising Change in Neurodegenerative Disorders. *Jou. Comp. Ass. Tom.*,20(60), pp 1012-1022, 1996.
3. S.Grootenk, C.Hutton, J.Ashburner, K.Friston and R.Turner, Characterisation and Correction of Movement-Related Interpolation Effects in MRI, *proc. ISMRM*, Manchester, 1997.
4. J.V.Hajnal, N.Saeed, E.J.Soar, A.Oatridge, I.R.Young and G.M.Bydder, A Registration and Interpolation Procedure for Subvoxel Matching of Serially Acquired MR Images, *Jou. Comp. Ass. Tom.*,19(2), pp 289-296,1995.
5. R.Lane, N.A.Thacker, N.L.Seed and P.A.Ivey, 'A Generalised Computer Vision Chip.' *Real Time Imageing*, 2, 203-213, 1996.
6. J.L.Otsuni, A. K.Santha, V.S.Mattay, D.R.Weinberger, R.L.Levin and J.A.Frank, Analysis of Interpolation Effects in the Reslicing of Functional MR Images, *Jou. Com. Ass. Tom*, 21(50), pp 803-810, 1997.
7. <http://bishopw.loni.ucla.edu/AIR3/index.html>
8. <http://www.fil.ion.ucl.ac.uk/spm/>