Parallel image restoration with spatially variant point-spread-function - Description and first clinical results.

Andrew Shearer[a], Gerard Gorman [a], Triona O'Doherty[a], Wil van der Putten[b], Peter Mc Carthy[c], Lukasz Jelen [a]

[a]Information Technology Centre, National University of Ireland, Galway, Ireland.
[b]Department of Medical Physics, University College Hospital, Galway, Ireland.
[c] Department of Radiology, National University of Ireland, Galway, Ireland.

ABSTRACT

In this paper we present a parallel code which performs an iterative image deconvolution using either a *spatially-invariant point spread function* (SI-PSF) or a *spatially-variant point spread function* (SV-PSF). The basic algorithm (Richardson-Lucy Maximum-Likelihood iterative procedure) is described as well as a description of the parallel implementation. Applications and results in the area of medical x-ray imaging are discussed.

Keywords: Spatially variant PSF, computed radiography, deconvolution, wavelet denoising

1. INTRODUCTION

Digital radiography presents a number of diagnostic and computational problems. One of them is the inherent under-sampling of the image by the relatively large pixel sizes - set by physical limitations and by scattering within the detector and patient. One solution, deconvolution through a measured PSF has been addressed in previous conferences[1].[2] This paper intends to bring these papers together and report the first results of a clinical trial where radiologist were presented with deconvolved and normal images at random on two different occasions.

The problem of image restoration (deconvolution), can be expressed in the following generic form;

$$\phi(r) = \int \psi(\xi) P(r - \xi) d\xi \tag{1}$$

where ψ is the function we wish to find, ϕ is the measured function and P is the kernel of the equation, or in particular it is the point-spread-function (hereafter PSF). P represents a local averaging of the function values associated with the remoteness of the observer and/or the impulse response of the measurement device. Fundamentally, imaging involves the detection and counting of photons, therefore the functions ϕ, ψ and P can be considered to be probability density functions. Thus, by definition these functions must be non-negative everywhere. This fact offers a powerful constraint in the determination of ψ which we will use later.

In the case of a digitalised image, Equation 1 can re-expressed as the quadrature;

$$\phi_i = \sum P_{ij} \psi_j \tag{2}$$

In this form, ϕ represents a pixel value. Assuming that ψ and ϕ are the same length; ψ can be determined through straight forward matrix inversion. However, this assumes that ϕ is free of errors/noise. This assumption is invalid in experimental sciences. As pointed out by Lucy[3] and others, a straight forward matrix inversion on data which contains noise leads to the amplification of the noise. However, iterative inverse techniques have been developed to deal with this problem. One approach is the Richardson-Lucy iterative procedure,[4].[5] This technique utilizes the non-negativity constraint on ψ_i, and uses the *Principle of Maximum Likelihood* as a goodness-of-fit criterion. We can consider the log-likelihood of some estimate ψ to be,

$$H = \sum \tilde{\phi}_i ln(\phi_i) \tag{3}$$

where ϕ_i is calculated from Equation 2. Now the problem involves the maximization of $H(\psi_i)$ subject to the constraint $\psi_i \geq 0$ for all i. This constrained optimization problem can be solved iteratively using the algorithm,

$$\psi_j^{r+1} = \psi_r^r \sum \frac{\tilde{\phi}_i}{\phi_i^r} P_{ij} \tag{4}$$

where $\psi_j^0 = constant$. Since each iteration increases the likelihood, and $H(\psi_i)$ is a convex function, the maximum H is approached monotonously.

In practice the image restoration algorithm described has a tendency to break up into near delta functions after high numbers of iterations. So in general one stops the iterative procedure after some number it iterations when a near optimal solution has been found. To overcome this limitation Starck et al[6] proposed a regularization technique based on a wavelet denoising algorithm applied to the residual between iterations. The residual is defined as,

$$R_i^r = \tilde{\phi}_i - \phi_i^r \tag{5}$$

Specifically the algorithm performs a hard threshold of the wavelet coefficients of the residual. The threshold level is set by a study of variation of Gaussian noise in this wavelet space. Although the algorithm proposed by Starck is effective, it is computationally demanding due to the redundant a'trous wavelet transformation applied. In addition, the method used to set the denoising threshold is based on the assumption of Gaussian noise. In order to improve upon this we use an algorithm based on a non-redundant transform - the Daubechies wavelet[7] which is more efficient in terms or storage. We denoise the residual between iterations using a soft thresholding (wavelet shrinkage) which is statically more attractive than hard thresholding.[8] We select the level dependent soft threshold as a fraction k of the wavelet coefficient standard deviation.

2. PARALLEL APPROACH

The approached outlined above is computionally very intensive - single processors are unlikely to be able to perfrom the deconvolution on large images within a reasonable period of time (minutes). However, deconvolution is an ideal algorithm for parallelization. If some continuous subset of an image is selected, it can be processed independently from the rest of the image.[9] For example, consider some image $I(i,j)$, where $0 <= i <= max_i$ and $0 <= j <= max_j$ and we wish to process some subsection of this image, $\iota(s,t)$, where $min_s <= s <= max_s$ and $min_t <= t <= max_t$. Then, this region along with a *halo* (or *padded*) region is selected. The width of the halo region must be $\geq$ the radius of the point-spread-function. this is to avoid errors at the edges. If the subsection is near the extremes of the image such that the halo extends over the bounds of I, then the external points are treated as having value zero. Figure 2 illustrates this point.

Thus an algorithm for processing the whole image would look like:

- divide the image into N sub-images
- include a halo region for each sub-image
- process each of the sub-images
- put the sub-images back together to form the total deconvolved image

If we want to execute this in parallel, it is natural to assume a task farm model (master-slave model). The master first reads in the image to be deconvolved. It then decides how the image is to be subdivided, and thus it form a list of *jobs*. It then distributes these jobs among idle *slaves* for processing, and awaits results (Figure 2). When a slave is finished with a subsection, it returns the result to the *master* and waits for another (Figure 3). In practice we find the master is usually idle, thus, in a parallel system who's processor supports multi-tasking, it is possible to run a slave on the same processor as the master without degradation of performance.

It is apparent from the above discussion that the more subsections I is divided into for processing, the more actual computation must be done, as there is more padded area to be processed. So although there is no reason why

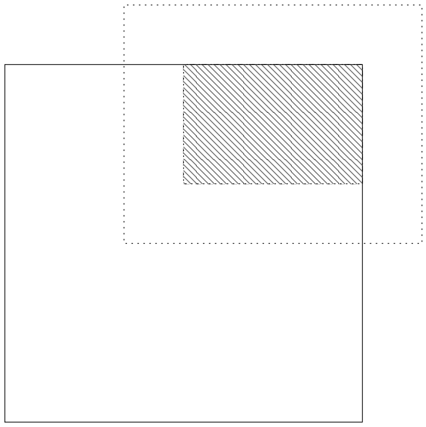

Figure 1. The solid line indicates the image I. The region in hatch is the region to be processed. The total frame which must be processed in order to minimize errors at the edges is indicated by the dotted line. The width of the halo region is selected to be equal to the radius of the PSF.

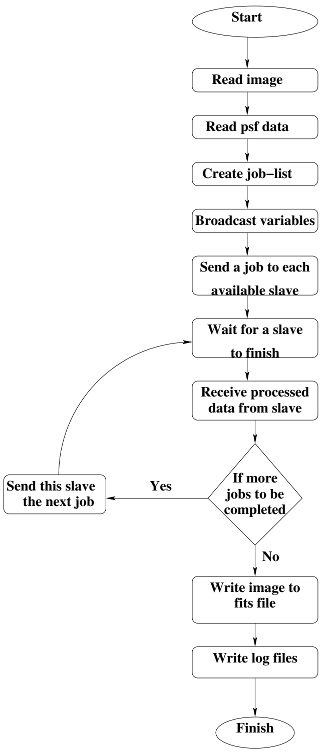

Figure 2. Flow diagram of master process.

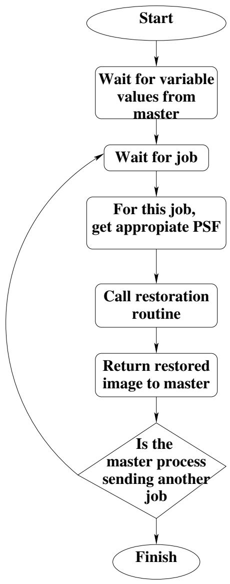

Figure 3. Flow diagram of slave process.

the image could not be subdivided into the same number of jobs as pixels, this would be in effect require the same about of computation as an image of extent;

$$N'_p = N_p * (N_{psf} + 1) \tag{6}$$

where N'_p is the effective number or pixels processed, N_p is the number or pixels in the original image and N_{psf} is the number of pixels in the PSF. If a globally invariant PSF is being used in the deconvolution, then it follows that a balance must be found when choosing the number or subsections the image is to be divided into, and the number or processors available. In general this would have to be done empirically for various systems.

In many cases however the assumption of an globally invariant PSF is gross. This variance is associated with a variance in the impulse response of the imaging apparatus. In this case one desires a spatially-variant PSF (SV-PSF). With this function, a PSF can now be generated for any location in the detector. We will assume that this is a smooth, slowly varying function (this is what one would expect from most well behaved instruments). So, for sufficiently small subsections of the main image, a good approximation to a SV-PSF is obtained by generating a PSF for the centre most pixel in the subsection, and then assuming that this PSF is locally invariant. The assumption of a locally invariant PSF has the additional benefit of allowing us to carry out the deconvolution in the Fourier domain. This has a significant numerical advantage.

From an implementation point of view, very little overhead is incurred when a slave (having a PSF model) has to generate a PSF for each job. It also had a negligible impact on communication overheads as the PSF-model only have to be broadcasted once at the beginning and thereafter only two additional integers must be sent for each job (the global coordinates of the centre most pixel in the sub-image).

From the above discussion one fact that should be now clear that the number of pixels processed per job is going to vary as,

$$N_J = \frac{N_p}{n} + N_{psf} \tag{7}$$

where N_J is the number of pixels per job and n is the number of jobs the image is subdivided into. However, while we are effectively linearly increasing the amount of computation with n, it should be kept in mind that for sufficiently accurate results an n sub-images (with their respective PSFs) may be required, thus the additional pixels which are added to the over all computation can not be viewed as being a side effect of parallelising the algorithm.

3. APPLICATIONS

The above treatment is quite general and the code can be applied to any image. One of the first issues with arise in these applications is the determination of the PSF. If we consider Equation 1 we see that if the function ψ is the Kronecker Delta function, $\delta(x)$, then, $\phi(r) = P(r - x)$. In the case of medical x-ray images it is not straight forward process to measure the PSF. We described our approach in where we were able to get a reasonable estimate of the PSF at different locations by placing a shield near the x-ray source. The shield had holes cut in it via laser which were finer than the maximum resolution of the instrument (the effective point source of x-rays providing a delta function). A sheet of perspex then provides a scatter coefficient which is not too different from human tissue. Thus the PSF can then be determined in the the usual way. When a sufficient number of PSFs have been measured from the field of view we can then construct a SV-PSF model.

3.1. Medical Imaging

Computed radiography has the potential to provide radiologists with a tool which can detect tumours earlier and with greater accuracy then film based systems. Although a digital radiography system can provide much greater contrast when compared with a conventional film system, the ability to detect small artifacts associated with breast cancer is limited by a reduced spatial resolution due to lack of screen sharpness and scatter induced fog. In this case screen sharpness refers to how accurately the imaging plate detects the x-rays, i.e. the impulse response of the detector. Screen sharpness is reduced due to scattering of x-rays in the imaging plate. We model the radiological image formation process as the convolution of a PSF (which we have measured) with the projected tissue density source function. When applied to a University of Leeds TORMAX breast phantom the results show as much as a two-fold improvement in resolution at the 50 percent MTF level.[1] Our results show that the regularized deconvolution algorithm significantly improves the signal-to-noise ratio in the restored image. The image sizes involved in this work are quite large (3730 x 3062). In the future oversampling of x2 and x4 will result in images whose size is in excess of 100 million pixels. Thus the parallel processing of the image can give an important speedup in processing speed which will be particularly important in the field.

We performed three trials. The first was using artifical objects which were blurred by convolution with a known point-spread function. These were then further modified by the addition of gaussian noise to produce our input images. These images were then deconvolved using the procedure outlined above. Figure 4 shows the test object, after blurring and noise addition, after deconvolution without regularisation and using finally denoised using the

Figure 4. Test object images. From top left hand corner - Truth image, Blurred truth image, Deconvolved image and Regularised deconvolved image

wavelet shrinkage procedure described. Figure 5 shows traces across these lines with an indication of the noise amplification under the 'normal' Lucy algorithm and using the regularised algorithm. Clearly the signal to noise ratio and resolution has been improved. Edge effects such as sinc function ringing have been cutback on the regularised curve.

In the second trial we obtained an improvement in resolution of about 2:1 after deconvolution. Figure 6 illustrates these results. These resuls were presented at the SPIE medical imaging conference in 1999.[1]

Finally we under-took a clinical trial. The first results have been described.[2] Our approach was to give the radiologists images which had virtually no noise amplification so that the appearance of the noise was not unusual. This was achieved at the expense of some enhancement. The images show the improvement obtained on a sample of these CR images.

Looking at the before/after image pairs from Figures 7 to 12 gives one an idea of how an x-ray may be improved to aid diagnosis.

4. CLINICAL TRIAL

The present trial involves assessing CR chest images which were obtained in an x-ray department in two consecutive days. These chest examinations were performed on adult outpatients (of both sexes) who required radiology for diverse pathologies. Thirty of these images were randomly chosen by an independant observer who didn't take part in the experiment.

For the purpose of this trial the raw format of the image was obtained i.e. no inherent processing was carried out on the image by the Agfa processor. The reason for doing this is that the PSF is measured for a raw image and in order for the algorithm to perform accurately the conditions in which the PSF is measured must remain constant with the conditions in which the image is taken.

Each of the thirty chest images were processed using the regularised deconvolution algorithm. It was found that 15 iterations of the algorithm produced the best supression of noise without affecting the anatomical detail. Hence,

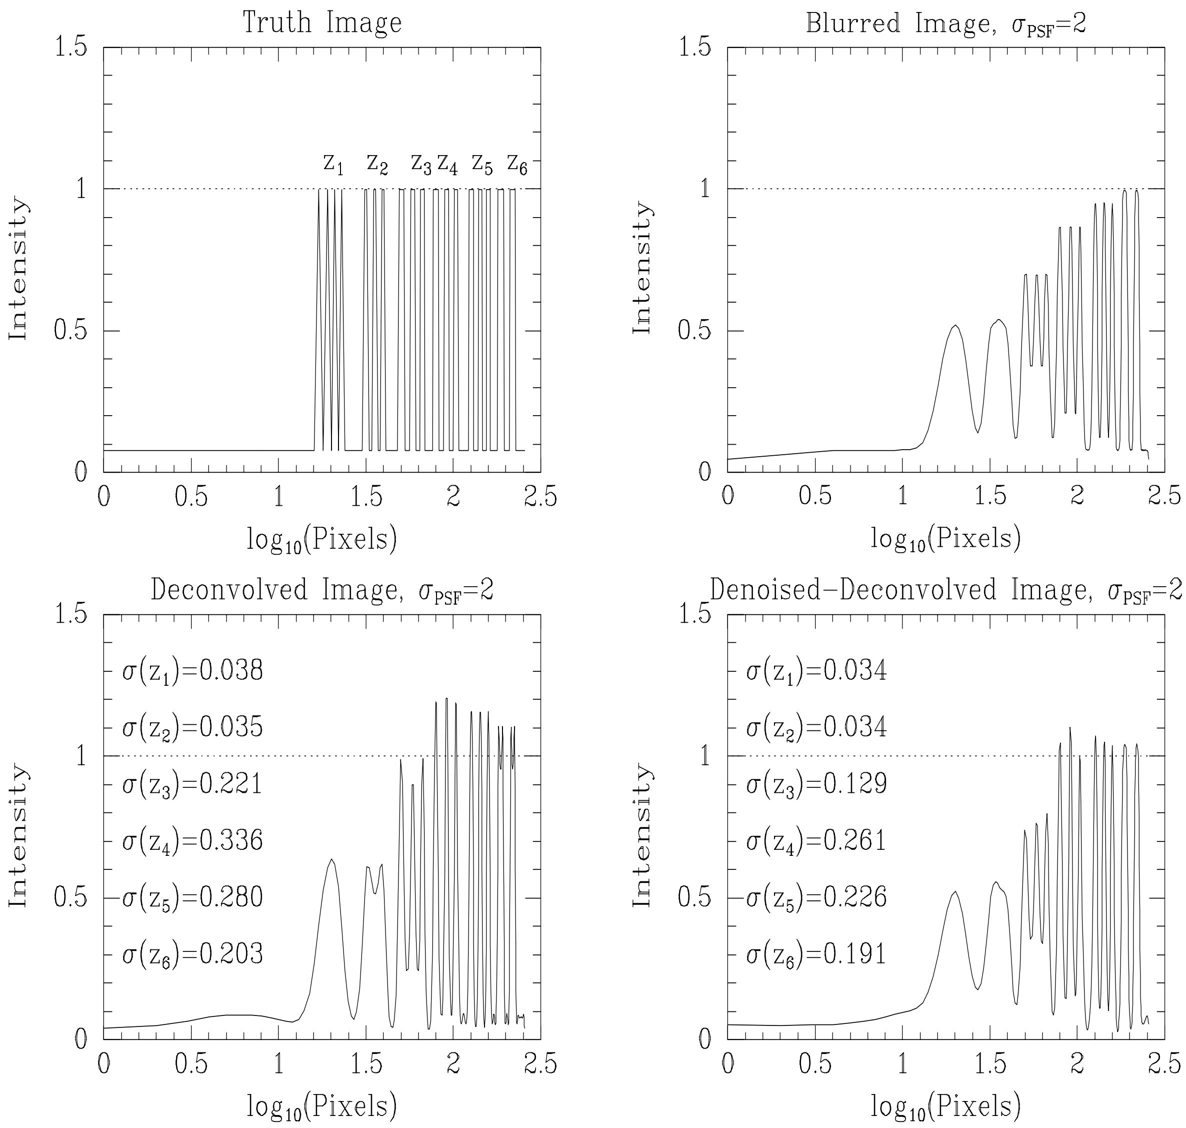

Figure 5. Profiles across resolution chart. The σ values indicate the better noise environment after regularisation.

each digitised x-ray was processed using 15 iterations. This task was carried out on a DEC Alpha workstation. It was very computationally intensive because image size was 3730 x 3062 pixels and 12 bits deep i.e. 4096 grey levels.

Six radiologists took part in the trial. All reading sessions were performed under low ambient light conditions. The radiologists were presented with digitized chest x-ray on an "interface" called saoimage. This allowed them to zoom in and out on any part of the image and also allowed them to adjust the contrast settings to give them the best possible visibility. The time spent reading each image was unlimited.

In order to assess the impact of the processing algorithm on clinical images, Guidelines on Quality Criteria for Clinical Radiographs[10] were used. These guidelines have been developed by the Commission of European Communities (CEC). They set out diagnostic requirements for a radiograph of a normal adult for common examinations. The radiologists were asked to score five criteria for each image on a rating scale of 1 to 5 (1=poor and 5=excellent). If the processed image was given a higher score than the original image then this was termed an improvement. If the situation was reversed, this term was given a negative value.

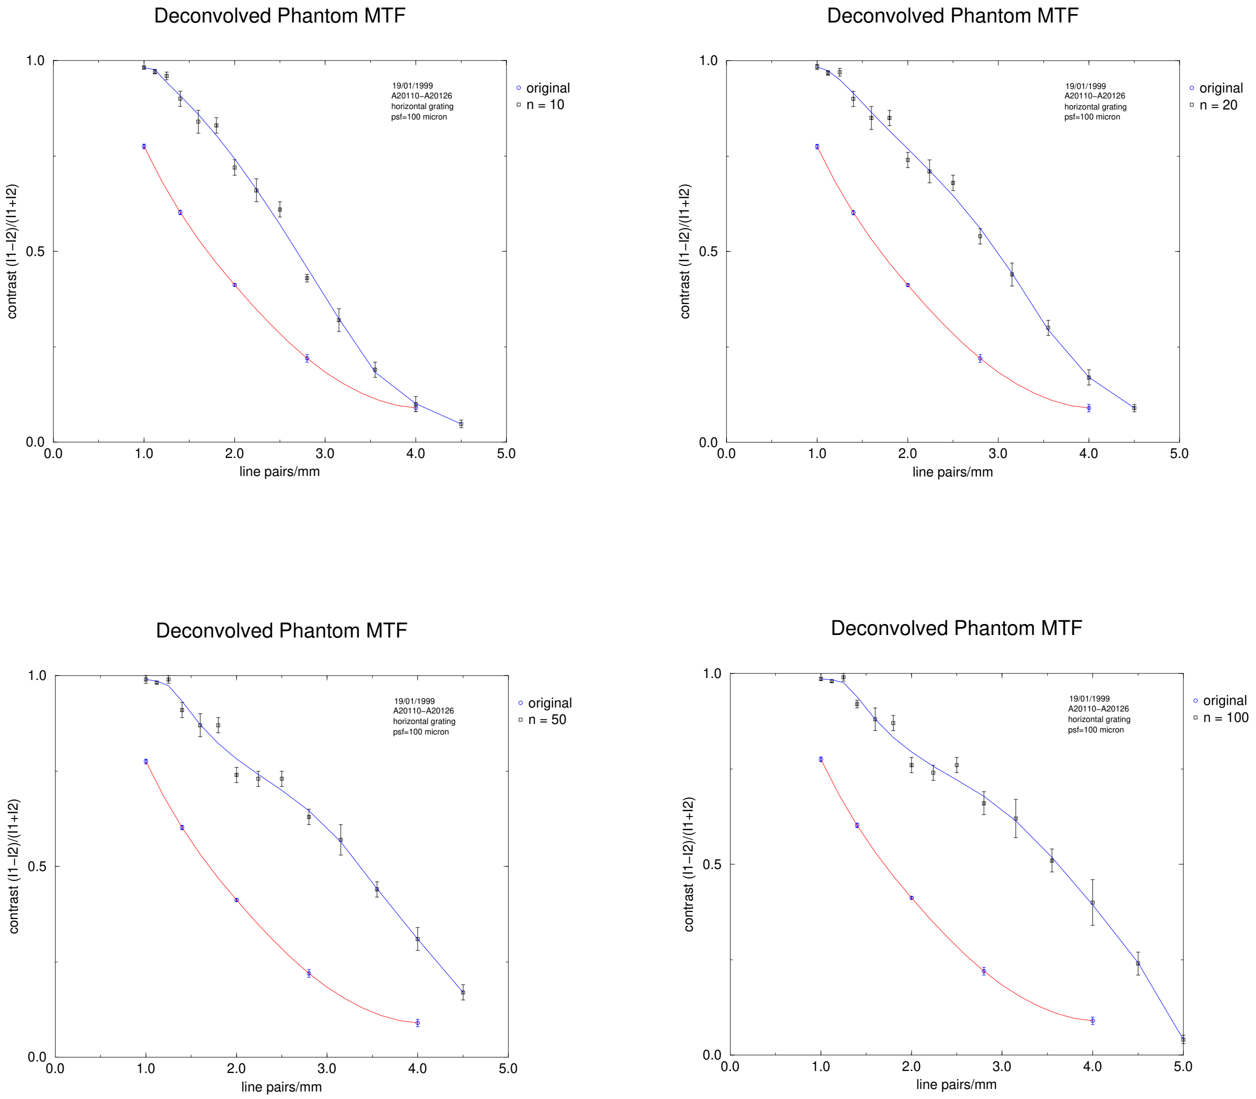

Figure 6. Modulation transfer function for 10,20,50,100 iterations with regularization

5. RESULTS

The results obtained are presented in Table 5. The results are reported as a score out of five. This method is used as because we are depending on human interpretation of image quality, this scale is non linear.

The results show that for all the image criteria scored an improvement is noted in the processed images. Visibility of the vascular lung is significantly improved when the x-ray is processed. In the borders of heart and aorta and the retrocardiac lung, spine and mediastinum there was a slight but insignificant preference for the processed image. Contrast was deemed to have improved substantially. Considerable improvement in image noise was noted on the processed x-rays. This shows the accuracy of the noise reduction in the algorithm. The standard error value is due to the fact that the hospital where this trial was carried out is not yet reporting from digitised images and so the radiologists vary in their idea of how anatomical visibility should be scored.

This image processing algorithm is particularly effective in improving the visibility in low contrast areas by significantly removing the noise. Hence features which appear in low contrast regions can be identified more readily when this algorithm is applied to the image. In general the algorithm works very well in balancing noise reduction with improved visibility of anatomical detail in the x-rays.

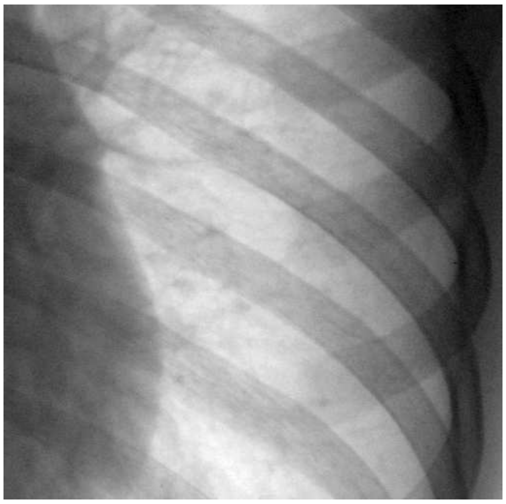

Figure 7. Original chest image

Figure 8. Deconvolved chest image

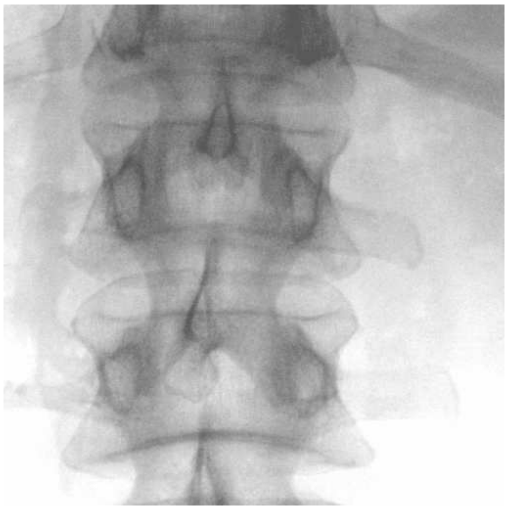

Figure 9. Original spine image

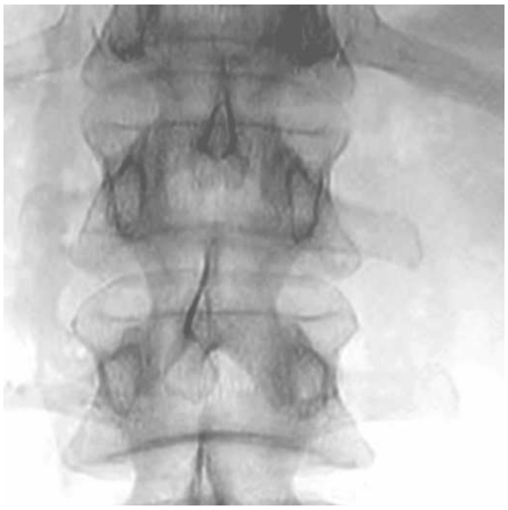

Figure 10. Deconvolved spine image

Figure 11. Original ball+socket image

Figure 12. Deconvolved ball+socket image

Image criteria	Orig Image (Mean score)	Mean Improvement	Std Error
Vascular lung	2.6	0.47	0.23
Borders of heart and aorta	3.6	0.19	0.26
Retrocardiac lung, spine and mediastinum	2.7	0.17	0.21
Contrast	3.2	0.33	0.21
Noise	3.4	0.60	0.20

Table 1. Results of the clinical trial

6. CONCLUSION AND FURTHER WORK

Using a SV-PSF can give much improved resolution when performing iterative deconvolution on images. By subdividing images into many sub-images each with their own PSF, iterative deconvolution can be carried out efficiently. Our code will be released under GNU Public License.[12] Is built on the MPI library. Our future trials will concentrate on different deconvolution recepies for different pathologies.

ACKNOWLEDGMENTS

We gratefully acknowledge the support of Enterprise Ireland, the Irish Research and Development Board for their continuing support of this work. The Western Health Board is also thanked for their kind support. We would like to especially thank the radiologists of University College Hospital, Galway who gave of their time and made this clinical trial possible.

REFERENCES

1. P. Abbott, A. Shearer, T. O'Doherty, W. van der Putten: Image deconvolution as an aid to mammographic artefact identification I: basic techniques Proc. SPIE **3661** (1999) 698-709
2. T. O'Doherty, A. Shearer, W. van der Putten, P. Abbott Image deconvolution as an aid to feature identification: a clinical trial Proc. SPIE (2000)
3. L. B. Lucy Astronomical inverse problems RvMA **31L** (1994)
4. W. H. Richardson: Bayesian based iterative method of image restoration J. Opt. Soc. Am. **62** (1972) 55
5. L. B. Lucy: An iterative technique for the rectification of observed distributions AJ **75** (1974) 745
6. J. L. Starck and F. Murtagh: Image restoration with noise suppresion using the wavelet transform Astronomy and Astrophysics **288** (1994) 343-348
7. I. Daubechies: Ten lectures on wavelets SIAM (1992)
8. D. L. Donoho: De-noising by soft-thresholding. IEEE Transactions on Information Theory **41** (1995) 613-
9. W. H. Press, S. A. Teukolsky, B. P. Flannery and W. T. Vetterling: Numerical Recipes: The art of Scientific Computing Cambridge University Press (1986) 413-414
10. Carmichael et al: European Guidelines and quality criteria for diagnostic radiographic images. Commission of European Communities (CEC, 1996) EUR 16260
11. D. C. Wells, E. W. Greisen and R. H. Harten A&AS **44** (1981) 363
12. http://www.gnu.org