Edge Localized Image Sharpening via Reassignment with Application to Computed Tomography

Timothy Dorney, Srikrishna Bhashyam, Andrew Doran, Hyeokho Choi,
Patrick Flandrin,† Richard Baraniuk

Department of Electrical and Computer Engineering, Rice University, Houston, TX 77251-1892, USA

† Ecole Normale Supérieure de Lyon, Laboratoire de Physique (UMR 5672 CNRS)
46 allée d'Italie 69364 Lyon Cedex 07, France

ABSTRACT

Traditional filtering methods operate on the entire signal or image. In some applications, however, errors are concentrated in specific regions or features. A prime example is images generated using computed tomography. Practical implementations limit the amount of high frequency content in the reconstructed image, and consequently, edges are blurred. We introduce a new post-reconstruction edge enhancement algorithm, based on the reassignment principle and wavelets, that localizes its sharpening exclusively to edge features. Our method enhances edges without disturbing the low frequency textural details.

Keywords: Image enhancement, edge sharpening, computed tomography, wavelets, reassignment

1. INTRODUCTION

The projection-slice theorem forms the foundation for tomographic reconstruction and is fundamental in a variety of fields (e.g., biomedical imaging, synthetic aperture radar, optical interferometry).[1-5] In biomedical computed tomography, cross section images of the human body are produced from data obtained by measuring the attenuation of x-rays along a large number of rays through the cross section (See Fig. 1).[2] The attenuation function at each projection angle is collected. The projection-slice theorem then states that the collection of 1-D projections, obtained from *all* possible angles, exactly reconstructs the 2-D Fourier transform of the original image. A 2-D inverse Fourier transform completes the reconstruction. Practical limitations, however, must be considered.

First, practical systems provide only a limited number of projection angles and a finite number of receivers. As a result, only a limited number of sampled 1-D projections is available. Second, rotational movement of the emitter/detector around the patient produces a polar sampling grid for the collection of 1-D projections. With a polar grid, the sample density is non-uniform, with a larger number of samples near the center and sparse sampling at the edges of the grid. Both of these practical considerations require some form of interpolation to produce the 2-D Fourier transform on a rectangular grid.[6] Third, digital implementation of the inverse Fourier transform requires the projections to be band limited. This bandwidth may be decreased due to high frequency noise considerations, but this also limits the desirable high frequencies created by the structures within the cross section.[7] Aliasing artifacts are then apparent in the inverse Fourier transformed image. Consequently, the vast majority of the errors induced by all three of these considerations occur at edges.[8] High pass information is essentially lost in practical implementations.

In this paper, we introduce a new, nonlinear image sharpening algorithm that is localized exclusively on edges. The algorithm operates on the post-reconstruction image and is based on the reassignment principle[9,10] and wavelets.[11,12] Edges are identified using the wavelet transform and are adaptively rebuilt using reassignment on the wavelet coefficients. We

E-mail: tdorney@ieee.org, skrishna@rice.edu, awd@zai.com, choi@ece.rice.edu, †flandrin@ens-lyon.fr, richb@rice.edu

Web: www.dsp.rice.edu, †www.ens-lyon.fr/~flandrin

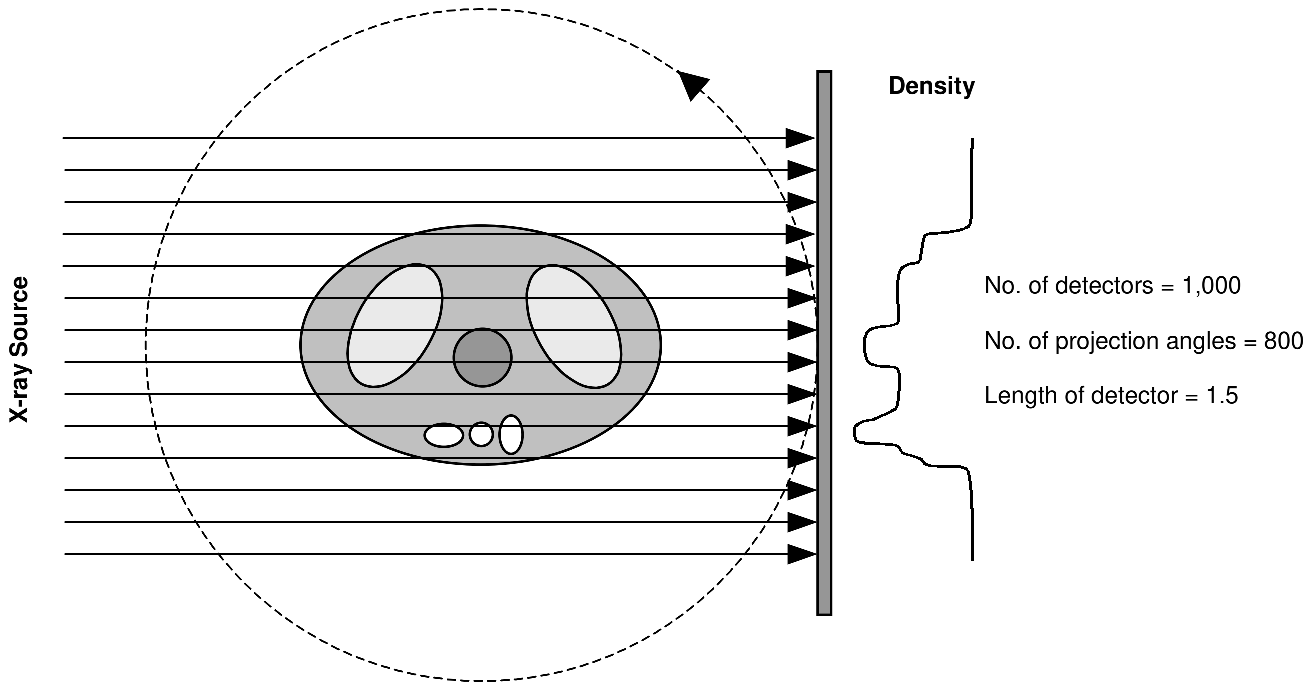

Fig. 1. Our projection-slice implementation collects 1,000 parallel rays onto a detector that has a normalized length of 1.5. The detector length is greater than the unit square in which the thoracic phantom resides to allow projections from any angle without clipping the phantom's features. A total of 800 different angles are used. We follow the general rule of having the number of projection angles equal or greater than $\pi/2$ times the number of desired pixels.[8] Our final reconstructed images have 512×512 pixels.

localize the filtering operation, because the vast majority of the error resides only in the edge features. Also, subtle variations (smooth regions) in gray scale may represent important features (e.g., a poorly differentiated tumor). Alteration of non-edge features in such cases is unacceptable.

This paper is organized as follows. Section 2 provides a brief introduction to the projection-slice theorem and some of its practical limitations. Section 3 discusses the concept of reassignment and the approach we used for this application. Section 4 details the Haar undecimated discrete wavelet transform and our reassignment algorithm. Section 5 contains the simulation results, and contrasts the algorithm's effect on edge versus non-edge features. Conclusions are in Section 6.

2. PROJECTION-SLICE THEOREM

The projection-slice theorem involves reconstructing a two-dimensional (2-D) function from its 1-D projections at all possible angles. The 1-D Fourier transform of the projection of a 2-D function, $g(x,y)$, at an angle ω is equal to the slice of the 2-D Fourier transform of the function at an angle ω. If we know the 1-D projections of a 2-D function at all possible angles, we know the 2-D Fourier transform of the function which is equivalent to knowing the function. We can, therefore, reconstruct the function exactly. One commonly used technique to reconstruct the original image from its projections is filtering and then back-projection.[1,2]

The image can be reconstructed from the projections using the Radon inversion formula:

$$g(x, y) = \int_0^{\pi} \int_{-\infty}^{\infty} P(f)\left|f\right| \exp\left(j\, 2\pi f \left(x\cos\omega + y\sin\omega\right)\right) df\, d\omega \qquad (1)$$

where $P(f)$ is the Fourier transform of the 1-D projection $p(t)$.[1-3,5] From the inversion formula (1), the first step in the reconstruction (inner integral) is to filter each of the projections using a filter with frequency response $|f|$. The filter has a

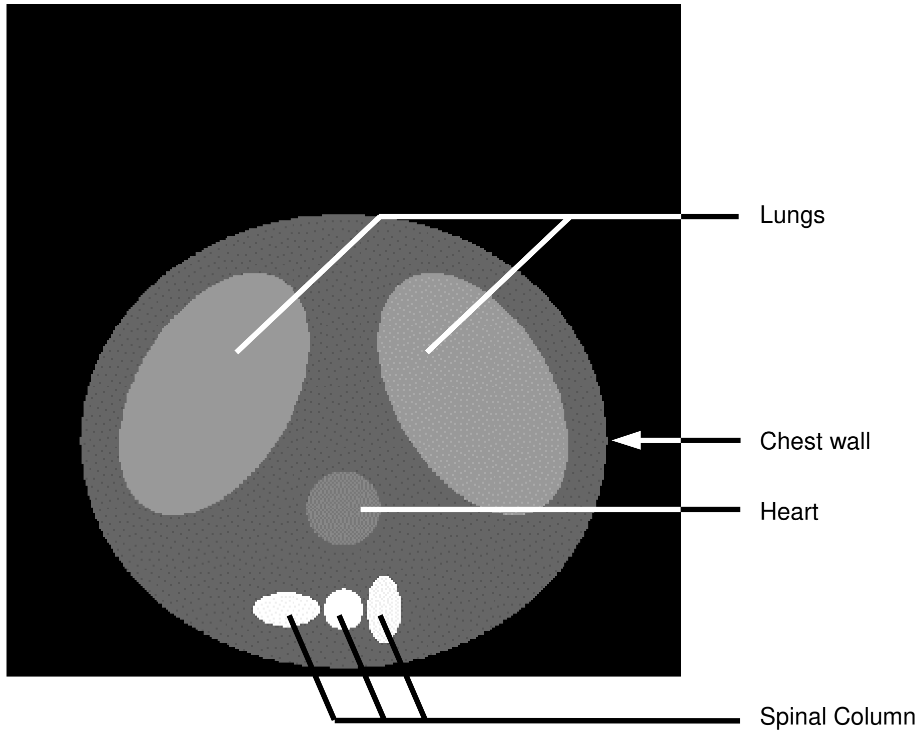

Fig. 2. The basic thoracic cavity phantom used for the simulations. The phantom captures three main anatomical structures: the lungs, heart, and spinal column without the rib cage. The lungs, heart, and chest wall are able to expand and contract to simulate these region's function. Also, the entire model can be shifted horizontally to simulate patient movement.[8] The phantom resides within a unit square.

practical limit in its digital implementation due to the cyclic nature of the discrete Fourier transform. The filter must be band limited. Practical concerns about attenuating high frequency noise also warrant a limited bandwidth.

After computing the filtered projections, we sample at $(x\cos\omega + y\sin\omega)$ for each ω and integrate over ω. In our implementation, we have only a finite number of sampled projections. Since the samples of the filtered projection required for back-projection may lie between the available sampling instants, we linearly interpolate.[6]

We use a thoracic phantom similar to the Shepp-Logan head phantom[1] to simulate a cross section (See Fig. 2). Our basic model captures three main anatomical structures: the lungs, heart, and spinal column without the rib cage. Regions of like functional activity (e.g., right and left lung) are assigned similar densities, although not exactly the same. Dislike functional components have widely varying densities. We used the thoracic model instead of the Shepp-Logan head phantom so that region growth (i.e., heart, lung, and chest wall expansion), and body shift (i.e., horizontal movement) during image acquisition could also be studied.[8]

Our projection-slice implementation collects 1,000 parallel rays at a detector that has a normalized length of 1.5. The detector length is greater than the unit square in which the phantom resides to allow projections from any angle without clipping the phantom's features. A total of 800 different angles are used. Our final reconstructed images have 512×512 pixels.

In Fig. 3, the spinal column geometries from Fig. 2 are enlarged. The top image is the original phantom while the bottom image is from the reconstruction. The majority of the squared error for the entire image resides in the edge features. The ellipses in our phantom, and therefore their projections, exceed the highest frequency of the $|f|$ filter in (1), which creates aliasing.

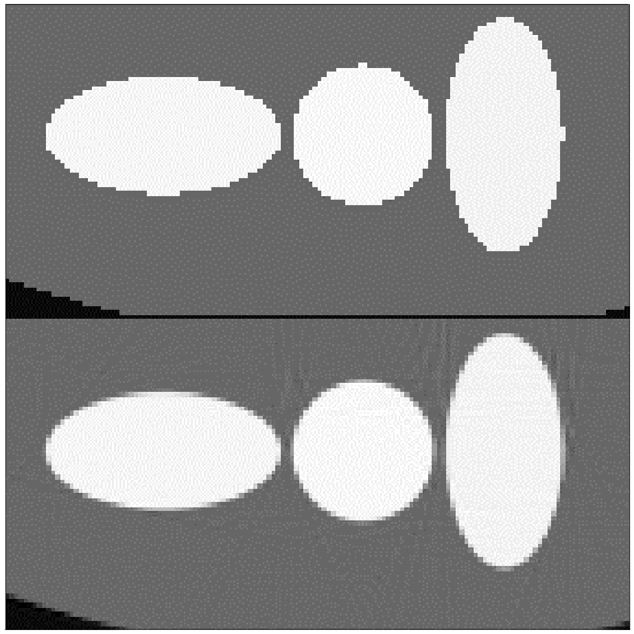

Fig. 3. The spinal column geometries from Fig. 2 are enlarged. The vast majority of the squared error for the entire reconstructed image resides in the edge features.

Improving the number of projection angles and the number of samples does reduce the squared error, but the error still resides primarily at edges. We, therefore, turn to post-reconstruction means for enhancement based on reassignment.

3. REASSIGNMENT

Any standard FIR linear convolutional filtering convolves the signal $s[n]$ with a filter kernel $h[n]$ of length M.

$$r[n] = s[n] * h[n] = \sum_{k=0}^{M-1} s[k]\, h[n-k] \tag{2}$$

Note that the location of the result $r[n]$ is placed at position n, which has a constant offset. Unfortunately, low pass, linear filters leave smooth regions relatively unaffected while areas of large change (discontinuities) are spread out or smeared. The filtered values, however, are denser surrounding the discontinuity. We require a means to rebuild these edges.

Reassignment is a nonlinear filtering approach that uses the center of mass to determine *where* $r[n]$ should be placed. The displacement function $d[n]$ calculates the center of mass of the signal, windowed with the filter:[9]

$$d[n] = n - \frac{\sum_{k=0}^{M-1} k \left| s[k]\, h[n-k] \right|}{\sum_{l=0}^{M-1} \left| s[l]\, h[n-l] \right|}. \tag{3}$$

Note that the location of the result $r[n]$ is now *reassigned* to location $d[n]$ instead of n as above (See Fig. 4). One problem with the discrete implementation is that the ratio in (3) is not necessarily an integer and could place a reassigned value between sample points. Several interpolation methods have been presented;[10] however, we constrain the location point as described below. First, let us review some of the advantageous properties of reassignment.

First, the convolution between two functions, one with length a and one with length b, results in a function that has length $a+b-1$. Conversely, reassignment results in a final function that has a length equal to the maximum of a and b. In other words, reassignment retains the same region of support. Second, Dirac delta functions remain unchanged by the reassignment operation. Finally, triangle shaped regions in the original signal result in triangle shaped regions in the result

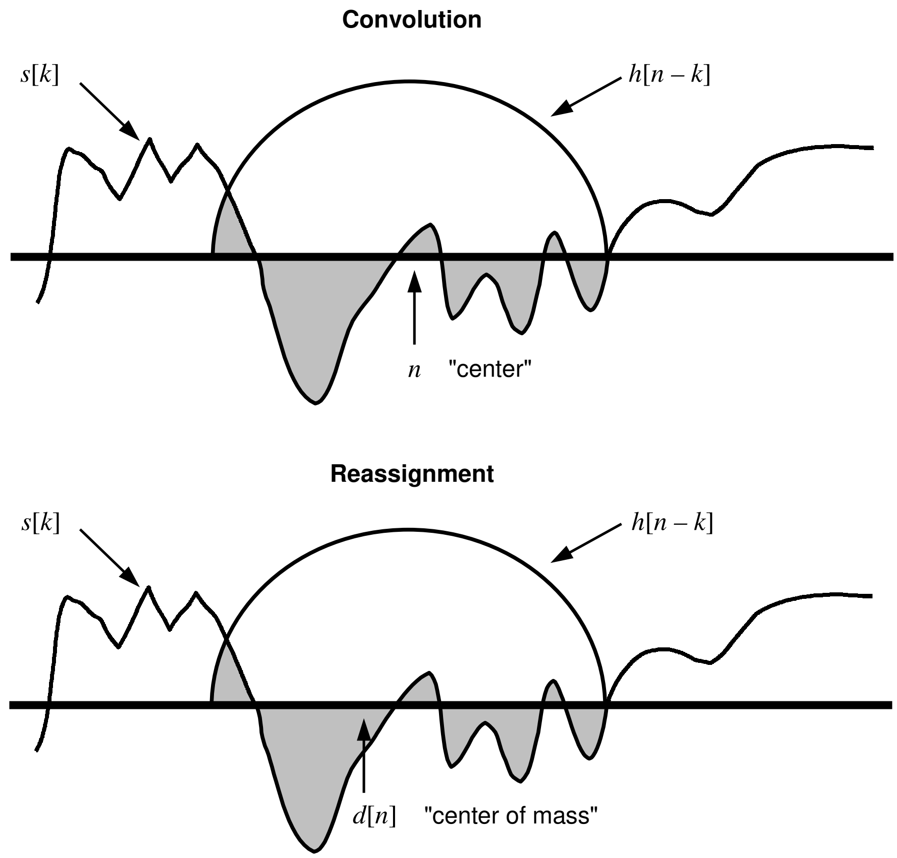

Fig. 4. Traditional convolutional filtering places the result at the center n of the operation. Reassignment differs in that the result is placed at the center of mass $d[n]$.

but with a smaller region of support after reassignment.[10] All of these properties are useful when considering our reconstructed image.

There is, however, one concern. Reassignment limits the amplitude at the center of mass to the amplitude of the filtered signal $r[n]$.[10] We alleviate the amplitude limit by *additive* reassignment. Instead of strictly relocating samples, we *rebuild* edges by adding samples to the center of mass. The details of which are provided in Section 4.

We also resolve any issue with $d[n]$ not being an integer. Since we use high and low pass filters of length two in the next section, let us focus our discussion to this specific case. Consider an edge (discontinuity) that has been reshaped by computed tomographic reconstruction. To be considered an edge, we assume that there still exist two samples separated by a relatively large jump. If a two-point linear FIR filter $h[n-k]$ is applied as in (2), a high pass (HP) filter output has a three-point triangular shape while a low pass (LP) output has a three-point smooth transition similar to the original edge with a central point in the middle. The center of mass at an edge feature is the peak of the triangular region for the HP output and the middle point of the smooth transition for the LP output. We recognize the features created by the two-point high and low pass filters, and use that information to our advantage.

4. UNDECIMATED WAVELET-DOMAIN REASSIGNMENT

We localize the application of the reassignment principle by using the undecimated discrete wavelet transform (UDWT). Wavelets have an integral edge detection property. A major drawback to wavelets is shift and rotation invariance.[13] Shift invariance is addressed by using the undecimated wavelet transform. We adequately handle diagonal edges of any orientation (rotation invariance) by using small image area operations (2×2 pixels).

The UDWT provides a high and low pass output. The single scale 1-D transform is given by:

$$LP[n] = \sum_k s[k]\ h[n-k]\,, \qquad k \in \mathbf{Z}\,,$$
$$HP[n] = \sum_k s[k]\ h'[n-k]\,, \qquad k \in \mathbf{Z} \tag{4}$$

where $LP[n]$ and $HP[n]$ are the low and high pass UDWT coefficients, and $h[n] = [1\ 1]$ and $h'[n] = [1\ -1]$ are the low and high pass Haar filters, respectively. Since our algorithm operates strictly on rows or columns, $s[n]$ is either the rows or columns of the computed tomography reconstructed image before edge enhancement. Only a single scale expansion of the wavelet transform is required to identify edge features.

We use the high pass Haar UDWT to identify the location of edges.[12] The transform operates first on each row of the image. We apply additive reassignment described below, then inverse transform. A second iteration is performed operating only on the columns, after all operations on the rows are completed. The inverse transform is given by:

$$\hat{s}[n] = \frac{1}{2}\left(LP[n] - HP[n]\right) \tag{5}$$

where $\hat{s}[n]$ is either the rows or columns of the post-reconstruction, locally edge enhanced image.

We focus first on the additive reassignment of the HP output. With reference to Fig. 5 (top), the HP signal is used to locate edges by noting the triangular regions at discontinuities. A local maxima of $|HP[n]|$ greater than a user defined threshold (*TOL*) identifies an edge. Let us call the local peak of the signal *WC*. To rebuild the edge, the samples adjacent to *WC* are added to the local peak, then reset to zero (additively reassigned), if they have the same sign as *WC*. Opposite signs indicate that the computed tomography reconstructed signal overshot the edge; therefore, it should not be reassigned. The HP reassignment allows the central point (*WC*) to exceed the original value of $HP[n]$, which is necessary to rebuild the edge. Also, our additive reassignment conserves the HP signal's energy.

We must now consider the reassignment of the LP filtered output. The LP filtered signal does not have the same spiky nature as the HP. It follows the trends in the original signal. Since we have reassigned the HP signal, the LP signal must be adjusted before we can recombine the signals through the inverse UDWT (5). Modifications at samples of the HP output require modifications at the same sample locations to the LP signal. The question now is what values do we use to modify the LP signal?

Since the peak of the reassigned HP output is the center of the transition in the LP output, we desire the "large" valued sample adjacent to the LP transition point to increase and the "small" valued adjacent sample to decrease. We adjust the LP samples adjacent to the LP transition point by the amount the HP samples were adjusted at the same positions as shown in Fig. 5. This achieves the more extreme slope desired for the LP signal.

We have one last concern. As mentioned previously, two cases exist for the HP output reassignment. We either modified the HP peak (*WC*) by both adjacent samples or only one. (A Dirac delta function in the HP signal remains unchanged after the HP reassignment.) To appropriately handle both of these cases, we additively reassign the values adjacent to *WC* to the LP transition point if they have the same sign as *WC*. Depending on the edge characteristics of the original signal, the transition point will not be altered substantially, or in the case of only one HP adjacent sample added to the LP transition point, the transition point becomes the start of the upward/downward slope of the edge.

Figure 5 shows the reassignment process graphically. On the left, a local maximum has been identified for the HP output, and both of the adjacent samples have the same sign. The adjacent samples are added or subtracted to neighboring samples as shown before the HP coefficients are reset to zero. Similarly, on the right, the same reassignment occurs even though we have a local minimum. The right most adjacent HP point is not considered, since it has the opposite sign as the local minimum. The procedure is also explained with the following pseudo-code:

1) Traverse each row of the high pass (HP) wavelet coefficients and check if the absolute value of a coefficient ≥ *TOL*. Determine if the identified sample is a local maximum or minimum by checking if its absolute value is greater than the absolute value of its adjacent HP coefficients on the same row. We label a local maximum or minimum *WC*.

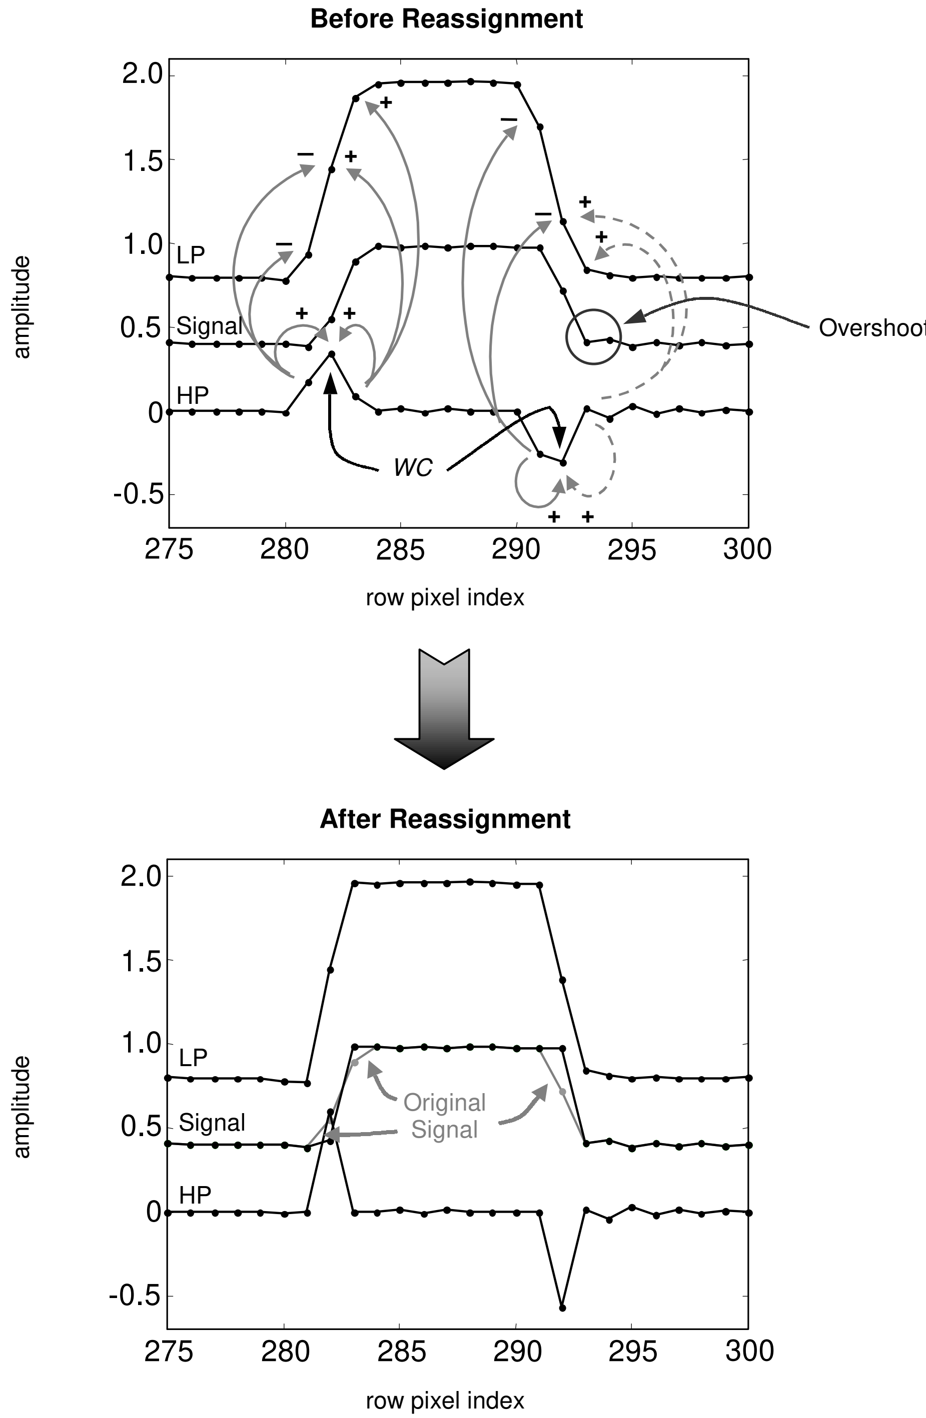

Fig. 5. The original row signal (top plot, middle) is decomposed into its high and low pass components using the Haar undecimated discrete wavelet transform (UDWT). Local maxima of the absolute value of the high pass output, greater than a user defined threshold, indicate edge features. The adjacent values to these local maximums (or local minimums), if the same sign, are used to reassign both the high and low pass adjacent values as shown in the top graph. (The right most adjacent value is not the same sign as the local minimum; hence, we display the reassignment operation with dashed lines.) The bottom graph shows the resulting waveforms for the high and low pass components. Both of these components are used to reconstruct the signal through the inverse UDWT (bottom plot, middle).

a) If the left adjacent HP coefficient has the same sign as *WC:*
   i) Add the left adjacent HP coefficient to *WC*.
   ii) Subtract the left adjacent HP coefficient from the low pass (LP) wavelet coefficient at the same location.
   iii) Subtract the left adjacent HP coefficient from the LP wavelet coefficient at the same location as *WC*.
   iv) Set the left adjacent HP coefficient to zero.
b) If the right adjacent HP coefficient has the same sign as *WC:*
   i) Add the right adjacent HP coefficient to *WC*.
   ii) Add the right adjacent HP coefficient to the low pass (LP) wavelet coefficient at the same location.
   iii) Add the right adjacent HP coefficient to the LP wavelet coefficient at the same location as *WC*.
   iv) Set the right adjacent HP coefficient to zero.

2) After all row operations are complete, apply inverse UDWT (5). Repeat the steps above for column operations on the column-wise UDWT coefficients.

The effect of the reassignments along the rows is to sharpen vertical edges and the vertical component of any diagonal edges. After all additive row reassignments are complete, we inverse transform along the rows. Next, the forward transform (4) is applied along the columns instead of the rows. Additive reassignment along the columns is completed before applying the inverse transform. In this fashion, the horizontal edges and horizontal components of any diagonal edges are also sharpened.

5. SIMULATION RESULTS

We demonstrate our procedure with the thoracic cavity phantom from Fig. 2. The phantom is contained within a unit square and is composed of 512×512 pixels. To reconstruct the phantom, we used 800 projection angles with each projection containing 1,000 sample positions on a detector of length 1.5. The projection-slice theorem and a convolution filter are used. We measure the squared error (SE) for the entire image using the following equation:

$$\mathrm{SE} = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} \left(\rho[m,n] - \hat{g}[m,n] \right)^2 \tag{6}$$

where $\rho[m, n]$ is the original phantom image, and $\hat{g}[m, n]$ is the post-reconstruction, locally reassigned image.

As shown in Table I, the squared error for the entire image is 51.9. (No noise was added in the simulation.) For analysis purposes, we define an edge as any pixel value that had a SE more than 0.0005. We based this value on the densities assigned to the phantom.[8] Even though the smooth features have a large SE (20.3) in the original reconstruction, over 95% of that error exists outside of the "patient's" body (i.e., the black region outside the cavity in Fig. 2). Errors in this region are easily removed since they are close to zero. Errors in smooth features primarily consists of Moiré patterns (visible in Fig. 3 (bottom)).

After completing the reassignment based, localized filtering by first operating on rows then on columns, the SE was lowered from 51.9 to 33.8. Reordering the process to first operate on columns then rows produced a SE of 34.1. In the former case (SE = 33.8), a dramatic reduction in error in edge features was observed while smooth region were basically unchanged.

The subjective difference can be seen by comparing the unfiltered reconstruction in Fig. 3 (bottom) against the image using post-reconstruction, localized reassignment in Fig. 6 (bottom). Figure 6 (top) contains the same enlarged region of the original phantom shown in Fig. 3 (top) for comparison.

TABLE I

SQUARED ERROR RESULTS FOR A ROW THEN COLUMN LOCALIZED REASSIGNMENT

	Original Reconstructed Image	After Localized Reassignment	Change
Edge Regions	31.61 (60.9%)	13.60 (40.2%)	18.01 (57.0%)
Smooth Regions	20.30 (39.1%)	20.24 (59.8%)	0.06 (0.3%)
Total	51.91 (100.0%)	33.84 (100.0%)	18.07 (34.8%)

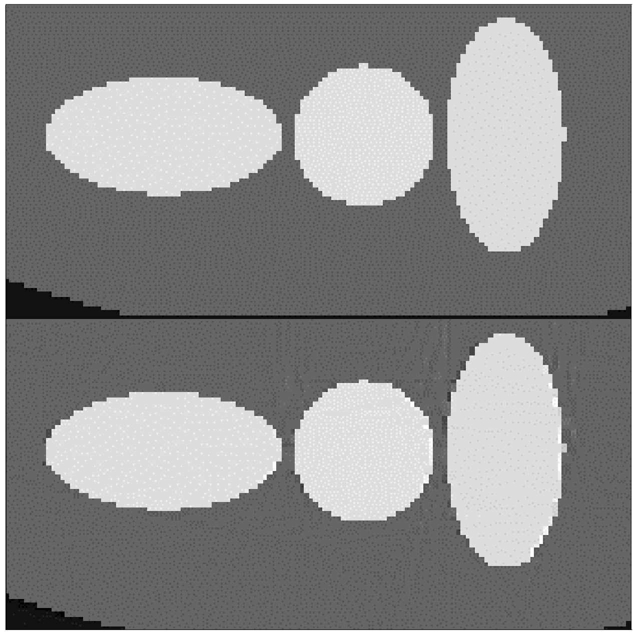

Fig. 6. The spinal column geometries from Fig. 2 are enlarged. The top image is the original phantom while the bottom image is from the post-reconstruction, localized reassignment. Our algorithm uses the undecimated discrete wavelet transform and reassignment to rebuild only the edges. Note the nearly exact pixel-for-pixel reconstruction of the ellipses' edges. (The width at the bottom of each ellipse does show some difference compared to the original.) At the same time, the Moiré patterns of the original reconstruction (Fig. 3 (bottom)) are still apparent. Non-edge features are not altered, as these may be key elements in the image and do not contain the majority of the error.

6. CONCLUSIONS

In this paper, we have proposed a new post-reconstruction, localized reassignment method. Practical application of the projection-slice theorem for tomographic reconstruction limits the amount of high frequency information. Consequently, edges are blurred while reproducing smooth regions accurately. Our method uses the undecimated discrete wavelet transform to localize the reassignment operation on the wavelet coefficients themselves. We essentially rebuild only the edge features. We focus our effort on the edge features since the vast majority of the error exists there. In addition, subtle variations in gray scale may indicate key features, such as a poorly differentiated tumor, and must not be altered.

The results of our test case are a reduction in squared error at the edges of 57%, while areas with small variations in gray scale were altered only 0.3%.

ACKNOWLEDGEMENTS

We wish to acknowledge the support of the Army Research Office, the Environmental Protection Agency, the National Science Foundation, DARPA, and CNRS.

REFERENCES

1. A. Rosenfeld and A. Kak, *Digital Picture Processing v.1*. Academic Press, New York, New York, 1982.
2. G. Herman, *Image Reconstruction from Projections*, Academic Press, New York, New York, 1980.
3. D. Dudgeon and R. Mersereau, *Multidimensional Digital Signal Processing*. Prentice Hall, Englewood Cliffs, NJ, 1984.
4. J. Lim, *Two-Dimensional Signal and Image Processing*. Prentice Hall, Englewood Cliffs, NJ, 1990.

5. K. Castleman, *Digital Image Processing*. Prentice Hall, Saddle River, NJ, 1996.
6. H. Choi and D. Munson, Jr., "Direct-Fourier Reconstruction in Tomography and Synthetic Aperture Radar", *International Journal Imaging System Technology*, **9**, pp. 1-13, 1998.
7. P. Joseph, "Sampling errors in projection reconstruction MRI", *Magnetic Resonance in Medicine*, **40**(3), pp. 460-466, Sept. 1998.
8. S. Bhashyam, A. Doran, and T. Dorney, "Parallel Beam Tomographic Reconstruction", www.dsp.rice.edu/~mit. Complete MATLAB® code for the thoracic phantom and a detailed report on parallel beam tomographic reconstruction available.
9. F. Auger and P. Flandrin, "Improving the readability of time-frequency and time-scale representations by the reassignment method", *IEEE Transactions on Signal Processing*, May, 1995.
10. J. Romberg, R. Baraniuk, and P. Flandrin, "Nonlinear Filtering by Reassignment", Rice Electrical Engineering Technical Report #TR9809.
11. I. Daubechies *Ten Lectures on Wavelets*. Springer-Verlag, Philadelphia, PA, 1992.
12. C. Burrus, R. Gopinath, and H. Guo, *Introduction to Wavelets and Wavelet Transforms: A Primer*. Prentice Hall, Upper Saddle River, NJ, 1998.
13. Y. Wan and R. Nowak, "Quasi-circular Rotation Invariance in Image Denoising", *IEEE ICIP-99*, pp. 25PP5.3, 1999.