Using Local 3D Structure for Segmentation of Bone from Computer Tomography Images

Carl-Fredrik Westin[+] Abhir Bhalerao[+] Hans Knutsson[‡] Ron Kikinis[+]

[+]Surgical Planning Laboratory
Brigham and Women's Hospital
Harvard Medical School
http://splweb.bwh.harvard.edu:8000
email: westin@bwh.harvard.edu

[‡]Computer Vision Laboratory
Dept. of Electrical Engineering
Linköping University
http://www.isy.liu.se/cvl

Abstract

In this paper we focus on using local 3D structure for segmentation. A tensor descriptor is estimated for each neighbourhood, i.e. for each voxel in the data set. The tensors are created from a combination of the outputs form a set of 3D quadrature filters. The shape of the tensors describe locally the structure of the neighbourhood in terms of how much it is like a plane, a line, and a sphere. We apply this to segmentation of bone from Computer Tomography data (CT). Traditional methods are based purely on gray-level value discrimination and have difficulties in recovering thin bone structures due to so called partial voluming, a problem which is present in all such sampled data. We illuminate the partial voluming problem by showing that thresholding creates complicated artifacts even if the signal is densely enough sampled and can be perfectly reconstructed. The unwanted effects of thresholding can be reduced by a change of the signal basis. We show that by using additional local structure information can significantly reduce the degree of sampling artifacts. Evaluation of the method on a clinical case is presented, the segmentation of a human skull from a CT volume. The method shows that many of the thin bone structures which disappear in a pure thresholding can be recovered.

1 Introduction

In this paper we focus on using local tensors as basis elements. The tensors are created from a combination of the outputs form a set of 3D quadrature filters [7]. The shape of the tensors describe locally the structure of the neighbourhood in terms of how alike it is to a plane, a line, and a sphere. Describing the signal in geometrical terms, rather than a neighbourhood of voxel values, is intuitively appealing. We have applied the tensor descriptors to the problem of segmenting bone, and in particular *thin* bone.

Segmentation of bones from Computer Tomography data has been shown to be very useful in surgery planning for procedures such as endoscopic sinus surgery and transnasal pituitary surgery. The segmentation is traditionally done by thresholding followed by some mending procedure such connectivity and manual editing. Segmentation of bone by thresholding is a fairly successful procedure since the CT values for bone are higher than that of surrounding soft tissues. Although the CT value for bone is well-known and constant between scans due to calibration, a discrimination based on these values does not work well in practice for all areas of bone. The main reason is because the acquired signal is sampled and not continuous and the effect of partial voluming becomes apparent after thresholding. Figure 6 (top left) shows the partial voluming effects of segmenting bone from CT data simply by thresholding. Large holes in the thin-bone structures occur, e.g. around the eye sockets, because of the location of the samples relative to the signal structure.

2 Signal sampling and local structure

The effect of thresholding a sampled signal is sometimes hard to anticipate, see Figure 1. Figures (a) and (b) show two gray-level images presenting a two-dimensional sine-wave pattern having different orientation. The size of these images is 32×32. The Nyquist theorem states that if a signal is sampled with at least double rate of the fastest signal variation, it can be reconstructed perfectly, i.e. it is not damaged by the sampling process. The sampling of the patterns in (a) and (b) are sampled with a rate approximately five times higher than the Nyquist limiting frequency, i.e. they are oversampled by a factor of about 5. This means

that they can in principle be perfectly reconstructed and, if required, resampled to any resolution.

Figures (c) and (d) show the effect of thresholding the signals (a) and (b) with the same thresholding value. Clearly, the thresholding has destroyed the basic structures which are apparent in the grey scale images. This also illustrates that even small transformations of the input signal may cause large unpredictable changes in the output. In (e) and (f) we have resampled the original grey-level images to 128×128 prior to thresholding with the same threshold as in (c) and (d). Again, note that resampling to any resolution is possible in this stylized example since a continuous signal can be reconstructed perfectly from the information in the top images. The thresholding artifacts in figures (e) and (f) are less severe and will be eliminated completely when the resolution reaches infinity, i.e. the continuous case.

Resampling of a signal can be regarded as a change of signal basis. The choice of basis as used above, giving a higher sampling rate is only one of many possible and may be impractical when dealing with large data sets. The success of a change of basis is, in general, related to how well the signal can be described in the new basis. Two common approaches to supersampling are a linear approximation, e.g. trilinear interpolation, or sinc interpolation, which is the optimal choice for bandlimited signals. Our choice is motivated by a need to describe signals that are not necessarily bandlimited, e.g. a plane in 3D, but nevertheless can be succinctly represented by a simple local basis. The more the descriptors resemble the signal, the lower dimensionality that is required for an adequate signal description. Furthermore, the gain in descriptive power helps reduce the noise sensitivity during segmentation.

The use of alternative signal basis to better model signal structure is not new [13]. Signal representations such as scale-space, and more recently wavelets [2], have been extensively used for this purpose. Specifically, scale space approaches have been adapted for detecting local "blob" like structures [9]. A more general approach is that of multiresolution Fourier descriptors which have been used for modeling image features such as line and texture [12].

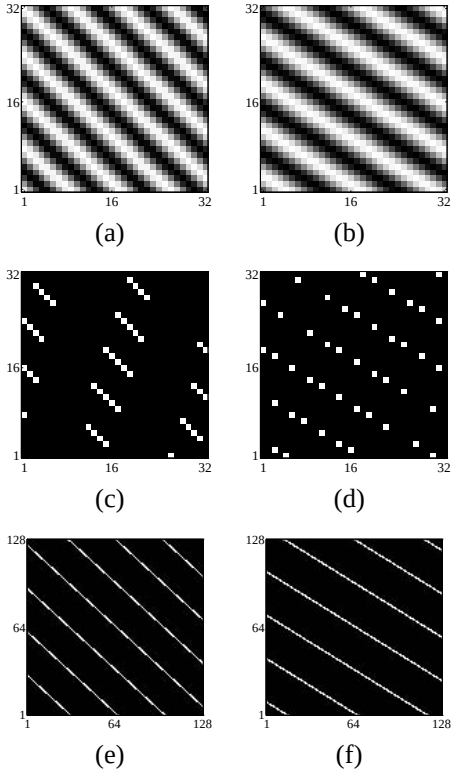

Figure 1. Effect of thresholding a sampled signal. (a) and (b) shows two gray-level 32×32 images presenting a two-dimensional sine-wave pattern having different orientation. In (c) and (d) the corresponding images are thresholded at the same threshold. In (e) and (f) the top grey-level images have been resampled to 128×128, and then thresholded with the same threshold as in(c) and (d)

3 Methodology

The technique presented in this paper is based on spatial filtering and tensor representations [3, 11]. For each neighbourhood, a set of quadrature filters are employed which are sensitive to signals at different orientations in space.

In two dimensions, the local structure is described in terms of dominant local orientation and isotropy, where isotropy means lack of dominant orientation. In three dimensions, in addition to isotropy, the tensor contains information about how planar or linear are the neighbourhoods. In higher dimensions, the basic shapes are more complicated than lines and planes, and the possible anisotropies become more complex. Previous work using these descriptors has mainly been for time varying 2D images such as optical flow estimation or enhancement of time varying images [3]

An underlying assumption for the approach taken here is that local gradient directions contain vital information. The importance of local orientation descriptions in the mammalian visual system has already been demonstrated by the physiological findings of Hubel and Wiesel and others [4]. Equally important is the *local simplicity hypothesis* [3], which is that the spatial variation of the gradient directions will be, in general, much slower than the spatial variation of the signal itself.

Examples of methods related to the tensor approach taken here in the sense that local structure is represented by a 3 by 3 matrix of numbers, are the use of local inertia matrices and Hessians [1, 8]

3.1 Local orientation in 3D

Regardless of dimensionality, a representation tensor $\mathbf{T}$ for orientation is given by

$$\mathbf{T} \equiv A\hat{\mathbf{x}}\hat{\mathbf{x}}^T \tag{1}$$

where $A > 0$ is an arbitrary constant and $\hat{\mathbf{x}}$ is a vector pointing along the orientation of maximal signal variation. The size of the tensor will depend on the dimensionality of the signal, e.g. three-dimensional signals are described by a 3×3 tensor.

A *simple neighbourhood* of n dimensions is expressed as:

$$\mathbf{S}(\xi) = \mathbf{G}(\xi^T\hat{\mathbf{x}}) \,, \tag{2}$$

where $\mathbf{S}$ is an n-variable function and $\mathbf{G}$ is a single variable function. ξ is the spatial position vector and $\hat{\mathbf{x}}$ is a constant vector oriented along the axis of maximal signal variation. In three dimensions, simple functions are constant on parallel planes. The direction perpendicular to these planes is the direction of maximum signal variation, $\hat{\mathbf{x}}$, or simply "the direction" of the signal. Two examples of three-dimensional simple signals are shown in Figure 2.

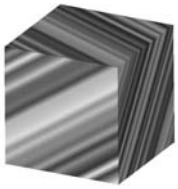

$\mathbf{S}_a(\xi) = \mathbf{G}_1(\xi^T\mathbf{x}_1)$

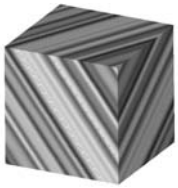

$\mathbf{S}_d(\xi) = \mathbf{G}_2(\xi^T\mathbf{x}_2)$

Figure 2. Two different three-dimensional simple neighbourhoods. The neighbourhoods are constructed using two different signal functions ($\mathbf{G}_1$ and $\mathbf{G}_2$) and two different signal orienting vectors ($\mathbf{x}_1$ and $\mathbf{x}_2$).

3.2 Geometric interpretation

In three dimensions, a visualisation of a symmetric tensor $\mathbf{T}$ can be made using an ellipsoid, i.e. an ellipsoid having eigenvectors as principle axes. The eigenvalue distributions and the corresponding tensor representations are given below for three particular cases of $\mathbf{T}$ for the three-dimensional case. The relations between the local spatial auto-correlation function and the corresponding energy distributions in the Fourier domain can be studied in Figure 3.

Let $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$ be the eigenvalues of a tensor $\mathbf{T}$ and let $\hat{\mathbf{e}}_i$ be an eigenvector corresponding to λ_i.

$$\mathbf{T} = \lambda_1\hat{\mathbf{e}}_1\hat{\mathbf{e}}_1^T + \lambda_2\hat{\mathbf{e}}_2\hat{\mathbf{e}}_2^T + \lambda_3\hat{\mathbf{e}}_3\hat{\mathbf{e}}_3^T$$

The rank, i.e. the number of non-zero eigenvalues, of the representation tensor reflects the complexity of the neighbourhood.

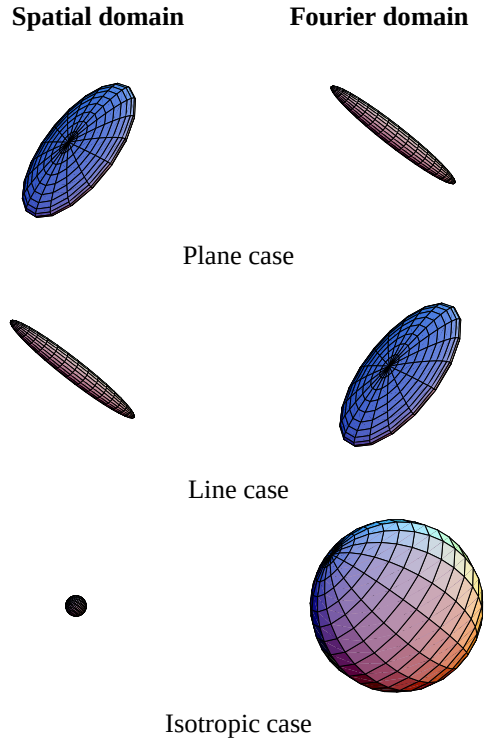

Figure 3. Iso-surface plots of spatial auto-correlation functions and corresponding energy distributions in the Fourier domain.

Plane case: The plane case corresponds to rank 1 or simple neighbourhood where $\lambda_1 \gg \lambda_2 \simeq \lambda_3$.

$$\mathbf{T} \simeq \lambda_1\mathbf{T}_1 = \lambda_1\hat{\mathbf{e}}_1\hat{\mathbf{e}}_1^T \tag{3}$$

This case corresponds to a neighbourhood that is approximately *planar*, i.e. it is constant on parallel planes of a given orientation. The orientation of the *normal vector* of the planes is given by $\hat{\mathbf{e}}_1$. A planar auto-correlation function in the spatial domain corresponds to energy being distributed on a line in the Fourier domain, see Figure 3 (top).

Line case: the line case corresponds to a rank 2 neighbourhood where $\lambda_1 \simeq \lambda_2 \gg \lambda_3$.

$$\mathbf{T} \simeq \lambda_1\mathbf{T}_2 = \lambda_1\,(\hat{\mathbf{e}}_1\hat{\mathbf{e}}_1^T + \hat{\mathbf{e}}_2\hat{\mathbf{e}}_2^T) \tag{4}$$

This case corresponds to a neighbourhood that is approximately constant on parallel *lines*. The orientation of the lines is given by the eigenvector corresponding to the smallest eigenvalue, $\hat{\mathbf{e}}_3$. An auto-correlation function concentrated on a line in the spatial domain corresponds to a planar energy distribution in the Fourier domain, see Figure 3 (middle).

Isotropic case: The isotropic case corresponds to a rank 3 neighbourhood where $\lambda_1 \simeq \lambda_2 \simeq \lambda_3$.

$$\mathbf{T} \simeq \lambda_1 \mathbf{T}_3 = \lambda_1 \, (\hat{\mathbf{e}}_1 \hat{\mathbf{e}}_1^T \; + \; \hat{\mathbf{e}}_2 \hat{\mathbf{e}}_2^T \; + \; \hat{\mathbf{e}}_3 \hat{\mathbf{e}}_3^T) \qquad (5)$$

This case corresponds to an approximately *isotropic* neighbourhood, meaning that there exists energy in the neighbourhood but no dominant orientation, e.g. in the case of noise. A spherical auto-correlation function in the spatial domain corresponds to a spherical energy distribution in the Fourier domain, see Figure 3 (bottom).

In general, $\mathbf{T}$ will be a linear combination of these cases,

$$\begin{aligned} \mathbf{T} &= \lambda_1 \hat{\mathbf{e}}_1 \hat{\mathbf{e}}_1^T + \lambda_2 \hat{\mathbf{e}}_2 \hat{\mathbf{e}}_2^T + \lambda_3 \hat{\mathbf{e}}_3 \hat{\mathbf{e}}_3^T \\ &= (\lambda_1 - \lambda_2) \hat{\mathbf{e}}_1 \hat{\mathbf{e}}_1^T + \\ &\quad (\lambda_2 - \lambda_3)(\hat{\mathbf{e}}_1 \hat{\mathbf{e}}_1^T + \hat{\mathbf{e}}_2 \hat{\mathbf{e}}_2^T) + \\ &\quad \lambda_3 (\hat{\mathbf{e}}_1 \hat{\mathbf{e}}_1^T + \hat{\mathbf{e}}_2 \hat{\mathbf{e}}_2^T + \hat{\mathbf{e}}_3 \hat{\mathbf{e}}_3^T) \end{aligned}$$

which gives that $\mathbf{T}$ can be expressed as:

$$\mathbf{T} = (\lambda_1 - \lambda_2)\mathbf{T}_1 \; + \; (\lambda_2 - \lambda_3)\mathbf{T}_2 \; + \; \lambda_3 \mathbf{T}_3 \qquad (6)$$

where $(\lambda_1 - \lambda_2)$, $(\lambda_2 - \lambda_3)$ and λ_3 may be viewed as coordinates of $\mathbf{T}$ in the tensor basis $\{\mathbf{T}_1, \mathbf{T}_2, \mathbf{T}_3\}$.

An alternative to ellipsoids is to use an object which can be seen as a sum of a line, a plate and a sphere. The line describes the principal direction of the tensor $\lambda_1 \hat{\mathbf{e}}_1 \hat{\mathbf{e}}_1^T$ where the length is proportional to the largest eigenvalue, λ_1. The plate describes the plane spanned by the eigenvectors corresponding to the two largest eigenvalues, $\lambda_2(\hat{\mathbf{e}}_1 \hat{\mathbf{e}}_1^T + \hat{\mathbf{e}}_2 \hat{\mathbf{e}}_2^T)$. The sphere with a radius proportional to the smallest eigenvalue shows how isotropic the tensor is, $\lambda_3(\hat{\mathbf{e}}_1 \hat{\mathbf{e}}_1^T + \hat{\mathbf{e}}_2 \hat{\mathbf{e}}_2^T + \hat{\mathbf{e}}_3 \hat{\mathbf{e}}_3^T)$. This type of object is shown in Figure 4.

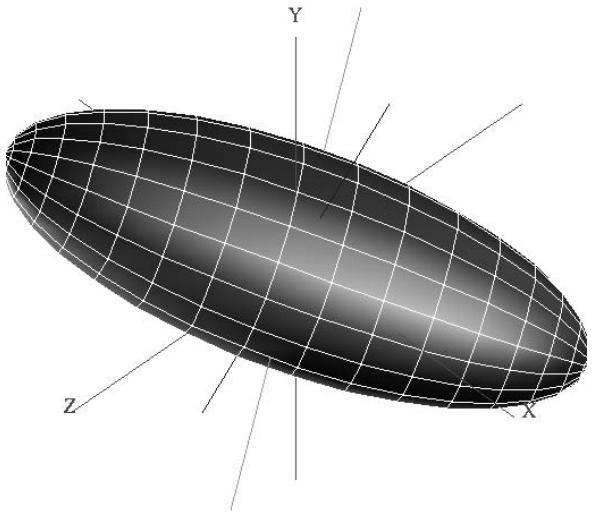

Figure 4. A way to visualise the geometry representing a symmetric tensor of rank 3.

3.3 Construction of the orientation tensor

Knutsson [7] has shown that a tensor representation of *local orientation* can be produced by combining the outputs from polar separable *quadrature* filters (see below). In three dimensions, the minimum required number of quadrature filters is six [7]. The orientation tensor $\mathbf{T}$ is obtained by a linear summation of the quadrature filter output magnitudes,

$$\mathbf{T} \; = \; \sum_k \mathbf{M}_k q_k \qquad (7)$$

were $\mathbf{M}_k$ is a *dual tensor basis* corresponding to a *tensor basis* $\mathbf{N}_k = \hat{\mathbf{n}}_k \hat{\mathbf{n}}_k^T$, where $\hat{\mathbf{n}}_k$ are the filter directing vectors and q_k is the output magnitude from the quadrature filter k. For details see [3, 11].

Acquired data are seldom simple in the sense of Equation (2). It is, however, still possible to find a best approximation to $\mathbf{T}$ corresponding to a simple neighbourhood, Equation (2). This is done by finding the rank one tensor $\mathbf{T}_s = A \hat{\mathbf{x}} \hat{\mathbf{x}}^T$ that minimizes

$$\Delta \; = \; \|\mathbf{T} - \mathbf{T}_s\| \qquad (8)$$

This expression is minimized by [7]:

$$\mathbf{T}_s \; = \; \lambda_1 \hat{\mathbf{e}}_1 \hat{\mathbf{e}}_1^T \qquad (9)$$

where λ_1 is the largest eigenvalue of $\mathbf{T}$ and $\hat{\mathbf{e}}_1$ is the corresponding eigenvector. The value of Δ indicates how well the one-dimensionality hypothesis fits the neighbourhood. The smaller the value, the better the fit.

3.4 Quadrature filters

A *quadrature filter* can be defined, independently of the dimensionality of the signal space, as a filter that is zero over one half of the Fourier space. In the spatial domain, the filter is complex: an even real part and an odd imaginary part. In two dimensions, the even part can be seen as a line filter, and the odd part as an edge filter [6, 3].

The *lognormal filter* is a spherically separable quadrature pair filter with a radial frequency function that is Gaussian on a *logarithmic* scale.

$$Q(\mathbf{u}) \; = \; R(\rho) D_k(\hat{\mathbf{u}}) \qquad (10)$$

where $\mathbf{u}$ is the multidimensional frequency variable, $R(\rho)$ and $D_k(\hat{\mathbf{u}})$ are the radial and the directional function respectively.

Directional functions having the necessary properties for directional interpolation are written [6, 7]

$$\begin{cases} D_k(\hat{\mathbf{u}}) \;=\; (\hat{\mathbf{u}}^T \hat{\mathbf{n}}_k)^2 & \text{if } \mathbf{u}^T \hat{\mathbf{n}}_k \;>\; 0 \\ D_k(\hat{\mathbf{u}}) \;=\; 0 & \text{otherwise} \end{cases} \tag{11}$$

where $\hat{\mathbf{n}}_k$ is the filter directing vector, i.e. $D(\hat{\mathbf{u}})$ varies as $\cos^2(\varphi)$, where φ is the angle between $\mathbf{u}$ and the filter direction $\hat{\mathbf{n}}_k$.

If the input signal is simple with the direction $\hat{\mathbf{x}}$, the output from a quadrature filter in direction k is

$$f_k \;=\; d(\hat{\mathbf{x}}^T \hat{\mathbf{n}}_k)^2 \;=\; d\langle \hat{\mathbf{x}}\hat{\mathbf{x}}^T, \hat{\mathbf{n}}_k \hat{\mathbf{n}}_k^T \rangle \tag{12}$$

where d is independent of the filter orientation and depends only on radial distribution of the signal spectrum and the radial filter function. The radial function $R(\rho)$ can thus be chosen arbitrarily without violating the directional interpolation properties of the filter, and the choice of $R(\rho)$ is subject to considerations similar to those traditionally found in one-dimensional filter design. In equation 7 the magnitude of the (complex) filter responses is used.

$$q_k = |f_k| \tag{13}$$

4 Local structure and segmentation

There are many ways of incorporating the knowledge about local structure given by the tensor descriptor introduced above in the process of segmentation. A field of estimated tensors could for example be used for driving active contour models [5]. These types of methods are, in general, fairly insensitive to local artifacts and noise since the models are global. They are currently gaining importance in interactive, semiautomatic, segmentation procedures. A good survey of deformable models used in medical images analysis can be found in [10].

CT segmentation is traditionally done by thresholding. In order avoid severe soft tissue artifacts a fairly high threshold is required. However, as seen in the stylized example in the introduction, a high threshold gives artifacts, especially in thin bone structures due to partial voluming. Here we use a straight forward example to show that important information is indeed captured by these descriptors. The segmentation is performed using local structure for adaptive thresholding. The threshold in spatial position $\mathbf{x}$ is defined by

$$t(\mathbf{x}) = t_0 - \beta\, c_{plane}(\mathbf{x}) \tag{14}$$

where t_0 is a global threshold which is modified locally by a constant β times the spatially varying planar measure $c_{plane}(\mathbf{x})$. A planar measure is chosen since thin bone strutures are problematic to segment. By adapting the threshold using the certainty value c_{plane}, the thin bone regions can be better recovered. In other cases e.g. segmentation of vessels, a linear measure would be preferable.

As mentioned, by examining the relations between the eigenvalues of the orientation tensor it is possible to determine to which shape category the neighbourhood belongs. For the plane case, the following certainty estimate was used:

$$c_{plane} = \frac{\lambda_1 - \alpha\lambda_2}{\lambda_1} \tag{15}$$

where α is a constant > 1. Negative certainty estimates in the formula above are set to zero. A large α sets the certainty to zero for all estimates not corresponding to the plane case.

In all practical cases, noise from various sources is present in the input signal. One way of removing noisy estimates is to remove all estimates below a certain threshold. Thresholding the output field based on the norm of the tensor removes some of the "good" estimates without removing all the noisy ones. Instead of doing this, it has been useful to add a constant tensor, $\epsilon\mathbf{I}$, to the estimated ones and then remove tensors based only on their shape.

$$\mathbf{T}_r = \mathbf{T} + \epsilon\mathbf{I} \tag{16}$$

where $\mathbf{I}$ is the identity tensor. The $\epsilon\mathbf{I}$ may be viewed as a *regularization* term. Small tensors will become proportionally "rounder" than larger ones when adding the tensor $\epsilon\mathbf{I}$, and will consequently be removed by a planar shape constraint. In the experiments below, $\epsilon\mathbf{I}$ has been added to the output field using an ϵ equal to 1 percent of the norm of the largest output tensor (globally).

4.1 Results

In evaluating the results, we have focused on areas of the human skull where the bone is thinnest: around the eye sockets and inside the nose. In fact, the sinus region consists of many fine bones clustered together.

Figure 5 shows the segmentation on a CT slice in the sinus (top left). The top right part show the grey-level values of this slice. The two lower images in this figure show the segmentation results. The bottom left image is the result from simple thresholding where much of the thin bone structure is missing. Decreasing the threshold to get the thin bone region introduces severe artifact since a we begin to include grey-level regions of soft tissues as well as bone.

The bottom right images is the result of incorporating the knowledge of how close the local signal corresponds to a plane. Comparing left and right images, there is better continuity of the fine thin bone structure in the processed data.

In Figure 6 we have surface rendered the segmented CT volume. The left column contains the result from segmentation using traditional thresholding. The right column contains the result obtained by incorporating local structure.

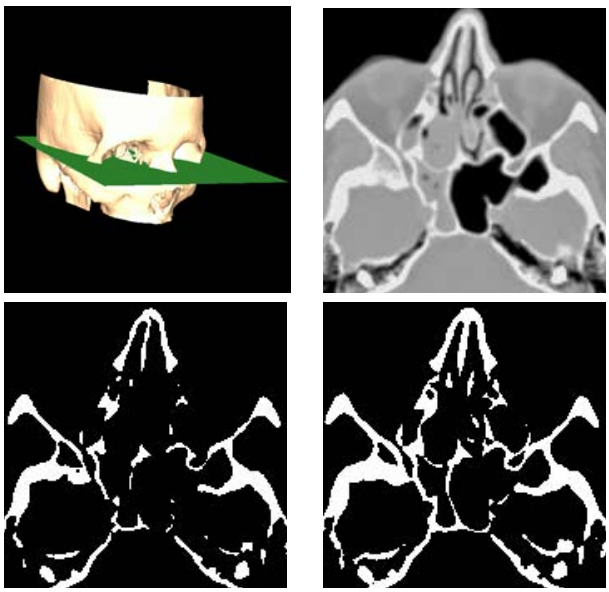

Figure 5. Result of segmentation of CT. Top left shows a a cut through the skull indicating the location of the slice of interest. Top right shows the grey-level image. Lower left shows the segmentation result from simple thresholding. Lower right shows the result from adaptive thresholding using local shape information.

The middle left figure shows that the algorithm was successful in filling in the holes in the orbital plate. Compare with the same region in the left image where the entire lacrimal bone is missing. That this bone is missing in the left images is not surprising since it is one of the thinnest bones in the skull. In the right image, the orbital surface of the frontal bone is also much improved (in the superior orbital region). There is also better definition of the vomer in the nasal cavity which is otherwise fragmented by simple thresholding. In the left cheek bone, the spurious hole below the infraorbital foramen is not present with knowledge of local structure. Note that this is accomplished without destroying the neighbouring foramen.

5 Conclusions

In this work we have concentrated on estimating to what degree a given 3D image neighbourhood is planar. Sampling theory shows us that partial voluming artifacts can be overcome by resampling the image using a new signal basis - one which more closely resembles the signal. A tensor description formed by combining the outputs of a set of orientation selective quadrature filters provides us with just such a signal decomposition. In 3D, we can interpret three simple neigbourhoods from the symmetric tensor: a planar, a linear and an isotropic case. This analysis when used to perform adaptive thresholding of CT data has been very effective in recovering thin bone structure which would be otherwise lost. We believe that a tensor description in general, and its geometric interpretation in 3D, is at the same time intuitively appealing and a powerful tool for a variety of tasks in image processing.

Acknowledgement

This work is funded by the Wenner-Gren Foundations, Sweden.

References

[1] J. Bigün, G. H. Granlund, and J. Wiklund. Multidimensional orientation: texture analysis and optical flow. In *Proceedings of the SSAB Symposium on Image Analysis*. SSAB, March 1991.

[2] I. Daubechies. The wavelet transform, time-frequency localization and signal analysis. *IEEE Trans. on Information Theory*, 36(5):961–1005, September 1990.

[3] G. H. Granlund and H. Knutsson. *Signal Processing for Computer Vision*. Kluwer Academic Publishers, 1995. ISBN 0-7923-9530-1.

[4] D. H. Hubel and T. N. Wiesel. Brain mechanisms of vision. *Scientific American*, 1979.

[5] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int J. on Computer Vision*, 1988.

[6] H. Knutsson. *Filtering and Reconstruction in Image Processing*. PhD thesis, Linköping University, Sweden, 1982. Diss. No. 88.

[7] H. Knutsson. Representing local structure using tensors. In *The 6th Scandinavian Conference on Image Analysis*, pages 244–251, Oulu, Finland, June 1989.

[8] T. Koller, G.Gerig, G.Szekely, and D. Dettwiler. Multiscale detection of curvilinear structures in 2-d and 3-d image data. In *Proc. ICCV'95*, 1995.

[9] T. Lindeberg. Scale-space behaviour of local extrema and blobs. *Journal of Mathematical Imaging abd Vision*, 1992.

[10] T. McInerney and D. Terzopoulos. Deformable models in medical image analysis, a survey. *Medical Image Analyis*, 1996.

[11] C.-F. Westin. *A Tensor Framework for Multidimensional Signal Processing*. PhD thesis, Linköping University, Sweden, S–581 83 Linköping, Sweden, 1994. Dissertation No 348, ISBN 91–7871–421–4.

[12] R. Wilson, A. D. Calway, and E. R. S. Pearson. A Generalized Wavelet Transform for Fourier Analysis: The Multiresolution Fourier Transform and Its Application to Image and Audio Signal Analysis. *IEEE Trans. on Information Theory*, 38(2):674–690, 1992.

[13] R. Wilson, H. Knutsson, and G. H. Granlund. Anisotropic non-stationary image estimation and its applications — part II: Predictive image coding. *IEEE Trans on Communications*, March 1983. Report LiTH–ISY–I–0463, Linköping University, Sweden, 1981.

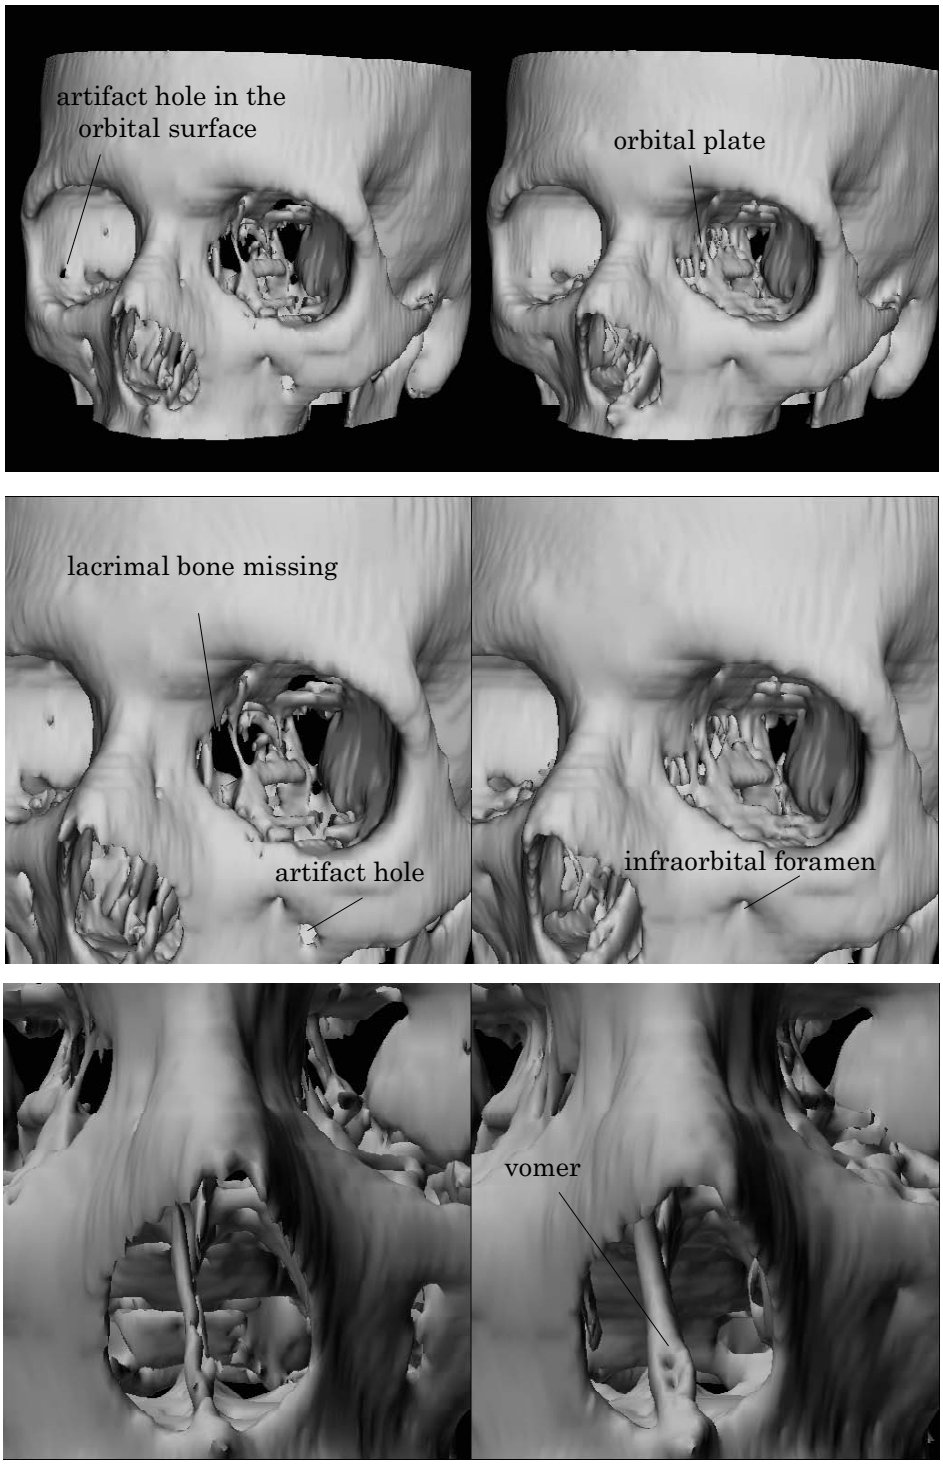

Figure 6. Left: Segmentation by thresholding CT values. Right Segmentation using additional local shape information. The top images show a overview of the surface rendered segmentation. Middle images show a close up of the left orbital region. Bottom images show a close up of the nasal cavity.