

Bone enhancement filtering: application to sinus bone segmentation and simulation of pituitary surgery

Maxime Descoteaux¹, Michel Audette², Kiyoyuki Chinzei², and Kaleem Siddiqi³

¹ Odyssee Project Team, INRIA Sophia-Antipolis / ENPC-Paris / ENS-Ulm Paris, France

Maxime.Descoteaux@sophia.inria.fr

² Surgical Assist Group, AIST, Tsukuba, Japan.

³ School Of Computer Science, McGill University, Montreal, Canada.

Abstract. We present a novel multi-scale bone enhancement measure that can be used to drive a geometric flow to segment bone structures. This measure has the essential properties to be incorporated in the computation of anatomical models for the simulation of pituitary surgery, enabling it to better account for the presence of sinus bones. We present synthetic examples that validate our approach and show a comparison between existing segmentation techniques of paranasal sinus CT data.

1 Introduction

Pituitary gland tumors represent the third most common primary intracranial tumors encountered in neurosurgical practice. In the majority, a surgical intervention is required. A neurosurgeon typically enters through the nose, has to break thin paranasal sinus bones and remove soft tissues while avoiding nerves and blood vessels to reach the pituitary gland (Fig.). This requires extensive practice and precise knowledge of the anatomy, the absence of which can have serious implications on the patient [0]. Currently, the only way to train a neurosurgery resident for such an operation is by multiple observation and by elementary maneuver attempts supervised by an expert neurosurgeon. This is why for the past several years, there has been a growing interest in building a surgical simulator to provide a tool for such training. Existing surgical simulators generally involve a generic anatomical model elaborated on the basis of extensive human supervision, interacting with a fast but constitutively limited biomechanics engine. The goal of our research is to

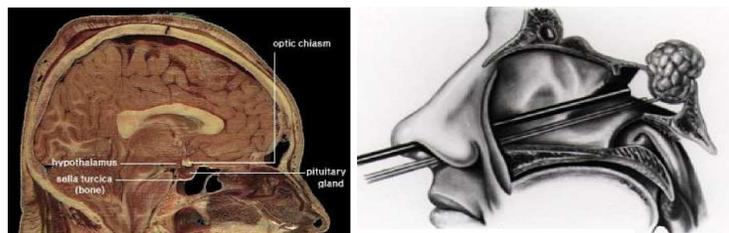


Fig. 1. Endoscopic pituitary surgery

formulate a minimally supervised method for producing a set of patient-specific anatomical models, from MR and CT data sets, in a manner that can interact with a hierarchical finite-element based biomechanics engine. To do so, we need a precise 3-dimensional (3D) partition of tissue classes into bone, air, vessel, nerve and soft-tissue.

In this paper, we focus on paranasal sinus bone enhancement and segmentation from CT data. The usual methods for segmentation of bone in CT are based on thresholding followed by some image connectivity measures or manual editing which is quite tedious and prone to human error. At a coarse scale, segmentation by thresholding is quite good due to the 3-class nature of CT images and the known Hounsfield value range for bone. Air has close to no signal and bone has much higher signal than surrounding tissues (see Fig.). However, from Fig. , we can clearly see that thin bones can have holes and diffusive boundaries. True bones do not have these features, they are due to partial volume effects and noise present in CT data sets. For such thin bones, a simple thresholding procedure gives unsatisfying results.

In this paper, we introduce a novel algorithm for bone enhancement filtering and segmentation based on the multi-scale second order local structure of the image. We exploit the eigenvalue decomposition of the Hessian matrix which is known to give valuable local information on the blob-like, tube-like and sheet-like behavior of iso-intensity level sets in the image $[0,0,0,0]$. We propose a *sheetness* measure that can be used to drive an active surface to segment bone. This is motivated in part by Frangi's tubular structure enhancement filtering measure [0] and Descoteaux et al [0] multi-scale geometric flow. We illustrate the power of the approach with segmentation results in regions with holes and low Hounsfield values in CT data and compare them to results from [0] based on local structure extraction with the structure tensor. We also validate the approach quantitatively on synthetic data. To our knowledge, our method is the first flow-based approach for paranasal sinus bone segmentation.

2 Using Local 3D Structure for filtering and segmentation

In this section, we investigate medical image filtering algorithms using the tensor matrix, \mathcal{T} and the Hessian matrix, \mathcal{H} as shape descriptors. For a 3D image \mathcal{I} , they are defined as

$$\mathcal{T} = \nabla \mathcal{I}^T \nabla \mathcal{I} = \begin{pmatrix} \mathcal{I}_x \mathcal{I}_x & \mathcal{I}_x \mathcal{I}_y & \mathcal{I}_x \mathcal{I}_z \\ \mathcal{I}_x \mathcal{I}_y & \mathcal{I}_y \mathcal{I}_y & \mathcal{I}_y \mathcal{I}_z \\ \mathcal{I}_x \mathcal{I}_z & \mathcal{I}_y \mathcal{I}_z & \mathcal{I}_z \mathcal{I}_z \end{pmatrix} \quad \mathcal{H} = \begin{pmatrix} \mathcal{I}_{xx} & \mathcal{I}_{xy} & \mathcal{I}_{xz} \\ \mathcal{I}_{xy} & \mathcal{I}_{yy} & \mathcal{I}_{yz} \\ \mathcal{I}_{xz} & \mathcal{I}_{yz} & \mathcal{I}_{zz} \end{pmatrix}.$$

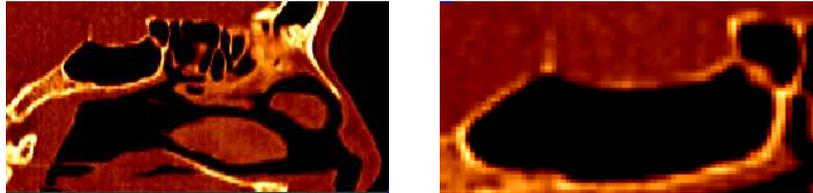


Fig. 2. Sagittal slice of CT data set in region of interest for pituitary gland surgery (see Fig.).

In both cases, an eigenvalue analysis is performed to extract the local behavior of iso-intensity level sets in the image. Many methods have been developed for blood vessel segmentation using these models [0,0,0,0,0]. We review only a selection of the representative techniques for our bone segmentation problem.

2.1 Modeling sheet-like structures with the tensor descriptor

As mentioned before, the usual method for bone segmentation for CT data is simple thresholding which fails on thin bone structures. Recently, Westin et al. [0] have introduced an adaptive thresholding approach using the structure tensor to segment thin bones around the eye socket and in the paranasal sinus area. The idea is to evaluate the structure tensor at every voxel of the data and determine the degree to which its shape resembles a line, a plane or a sphere. Letting $\lambda_1, \lambda_2, \lambda_3$ ($0 \leq \lambda_1 \leq \lambda_2 \leq \lambda_3$) be the eigenvalues of the structure tensor, the interest is when the structure tensor can be approximated by a plane (see [0] for details). The authors describe a planar measure $c_{plane} = \frac{\lambda_3 - \alpha \lambda_2}{\lambda_3}$ such that in theory, it has a value of 1 for plane structures and 0 for others. This measure is used to adaptively threshold the input data set, $t(x) = t_0 - \alpha c_{plane}(x)$, where t_0 is a global threshold modified locally and α is a weight factor for the planar measure.

In Section , we demonstrate several properties of this approach. In particular, since the tensor is based on first order variation, the c_{plane} measure is strong at boundaries (where the gradient is strong) and weak inside the bone structure. For our application, we seek a measure that is high at the center of the structure with a fall off at boundaries where a priori, our confidence in a voxel being strictly bone or strictly soft tissue is weak. Such a confidence index can guide the subsequent surface and volume meshing of bone tissue relevant to the simulation of pituitary surgery. We explore the properties of the Hessian shape operator to define such a measure.

2.2 Modeling sheet-like structures using the Hessian operator

The Hessian matrix encodes important shape information. An eigenvalue decomposition measures the maximum changes in the normal vector (gradient vector) of the underlying intensity iso-surface in a small neighborhood. Hence, it can differentiate between tube-like, sheet-like and blob-like structures. The classification of Table was first explored by Sato et al. [0], and Lorenz et al. [0] separately. In [0], Frangi defines three ratios using tube-like properties of the eigenvalues of Table to separate blood vessels from other structures. Derivatives of the Hessian matrix are computed with multi-scale γ -parametrized

eigenvalue conditions	local structure	examples
$\lambda_1 \approx \lambda_2 \approx 0, \lambda_3 \gg 0$	sheet-like	bone, skin
$\lambda_1 \approx 0, \lambda_2 \approx \lambda_3 \gg 0$	tube-like	vessels, nerves
$\lambda_1 \approx \lambda_2 \approx \lambda_3 \gg 0$	blob-like	nodule
$\lambda_1 \approx \lambda_2 \approx \lambda_3 \approx 0$	noise-like	background, noise

Table 1. Local structure classification assuming $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$.

Gaussian kernels and the three quantities are integrated in a *vesselness* measure designed to be maximum when computed at the scale corresponding to the radius of the tubular objects. The vessel index is thus maximum nearby the vessel center and is zero outside. In [0] a vesselness measure is used to find putative centerlines of tubular structures along with their estimated radii and is then distributed to create a vector field which is orthogonal to vessel boundaries so that the flux maximizing flow algorithm of [0] can be applied to recover them. This method can recover low contrast and thin vessels from standard anatomical proton density weighted data sets.

Inspired by these approaches, we propose a multi-scale *sheetness* measure that enhances bone structures and then use it to drive a deformable surface that stops at bone boundaries. At every voxel, we determine whether the underlying iso-intensity surface behaves like a sheet. In this case, we know that the eigenvectors corresponding to the null eigenvalues span the plane of the plate structure and the other eigenvector is perpendicular to it. We define three ratios, R_{sheet} , R_{blob} , R_{noise} , to differentiate sheet-like structures from others. Their behavior is described in Table . Then, just as in [0], we can define the

Ratios	sheet	tube	blob	noise
$R_{sheet} = \lambda_2 / \lambda_3 $	0	1	1	<i>undefined</i>
$R_{blob} = (2 \lambda_3 - \lambda_2 - \lambda_1) / \lambda_3 $	2	1	0	<i>undefined</i>
$R_{noise} = \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}$	λ_3	$\sqrt{2}\lambda_3$	$\sqrt{3}\lambda_3$	0

Table 2. Theoretical properties of the ratios defined to construct the sheetness measure.

sheetness measure, S , as the maximum response over all scales σ at which the derivatives of the Hessian are computed,

$$S = \max_{\sigma \in \Sigma} S(\sigma) = \begin{cases} 0 & \text{if } \lambda_3 > 0 \\ \left(\exp\left(\frac{-R_{sheet}^2}{2\alpha^2}\right) \right) \left(1 - \exp\left(\frac{-R_{blob}^2}{2\beta^2}\right) \right) \left(1 - \exp\left(\frac{-R_{noise}^2}{2c^2}\right) \right), & \end{cases} \quad (1)$$

where Σ is a finite set of scales chosen by the user in the range of smallest to thickest possible bone structure ($0.5 \leq \sigma \leq 3.0$) in the data and parameters α, β, c are set to 0.5, 0.5 and half the maximum Frobenius norm (R_{noise}) respectively as suggested in [0,0].

Each term of the equation has a function and depends on the characteristics of Table . To avoid division by a null λ_3 in the case of noise, the *undefined* can be set to obtain the desired behavior. Breaking down the terms of Eq. 1, we have

1. $\exp\left(\frac{-R_{sheet}^2}{2\alpha^2}\right)$ is a sheet enhancement term, where the maximum occurs for sheet-like structures and the minimum for others. We set *undefined* to 1.
2. $\left(1 - \exp\left(\frac{-R_{blob}^2}{2\beta^2}\right)\right)$ is a blob and noise elimination term since it is zero for both. Moreover, the term is high for a sheet and lower for a tube. We set *undefined* to zero.
3. $\left(1 - \exp\left(\frac{-R_{noise}^2}{2c^2}\right)\right)$ is a background or noise reduction term. R_{noise} is known as the Frobenius norm and is high only in the presence of structure.

Note that we do not define a tube elimination term as the curved ends of bone structures have a behavior that is both tube-like and sheet-like. Thus, the sheetness measure is designed to be maximum for sheet-like voxels, less important for tube-like regions and zero for other structures. The power of this approach resides in the fact that after this sheetness computation, we have a confidence sheet-like score at each voxel and in addition, for high score locations, we have the scale estimate of the radius of the sheet as well as the normal vector to the plane.

3 A Geometric Flow for Segmenting Bone Structures

There have been a few deformable model methods proposed for bone segmentation which are quite different from our approach because they are suited to 2D images from different modalities and different bones. In [0], the segmentation of carpal bones of the hand in CT images is faced with similar challenges as in our sinus bone CT data sets. A skeletally coupled curve evolution framework is proposed that combines probabilistic growth with local and global competition. Promising results are shown on 2D images with gaps and diffused edges. However, the method is based on skeletal points which would be difficult to determine in one or two voxel wide bone structures such as those in the paranasal sinuses.

In our application, we propose to use the bone enhancement measure of Eq. 1 to drive a 3D surface evolution. We construct a vector field that is both large in magnitude and orthogonal to bone structures. The key idea is to *distribute the sheetness measure, which is concentrated on the center sheet, to the bone boundaries* implied by the local scale and direction estimates coming from the multi-scale sheetness measure. At each voxel where the sheetness index is high, we consider a disc or flat ellipsoid with its two semi-minor axes aligned with the estimated plane orientation and its semi-major axis equal to the approximated radius. The sheetness measure is then distributed over every voxel on the boundary of the disc. We define the addition of the extensions carried out independently at all voxels to be the ϕ *distribution*. The extended vector field is now defined as the product of the normalized gradient of the *original* image with the above ϕ distribution, $\vec{\mathcal{V}} = \phi \frac{\nabla I}{|\nabla I|}$. This vector field is an ideal candidate for the static vector field in the flux maximizing geometric flow $S_t = \text{div}(\vec{\mathcal{V}})\vec{\mathcal{N}}$. The flow evolves the surface S to converge to the zero-crossing of the speed term to a $\text{div}(\vec{\mathcal{V}})$ (see Fig.). The flow is topologically adaptive due to its implementation with levelset methods, is computationally efficient and requires minimal user interaction. We refer the reader to [0] for further details.

4 Quantitative Validation on Synthetic Objects

In order to validate the extension of the sheetness measure to boundaries and to evaluate the effectiveness of the speed term driving the geometric flow, we constructed several binary synthetic objects of varying widths and centerline curvatures. Each volume was then smoothed using 2 iterations of mean curvature smoothing to simulate partial volume effects at their boundaries. Ground truth surface points were obtained as the 0.5 crossings of

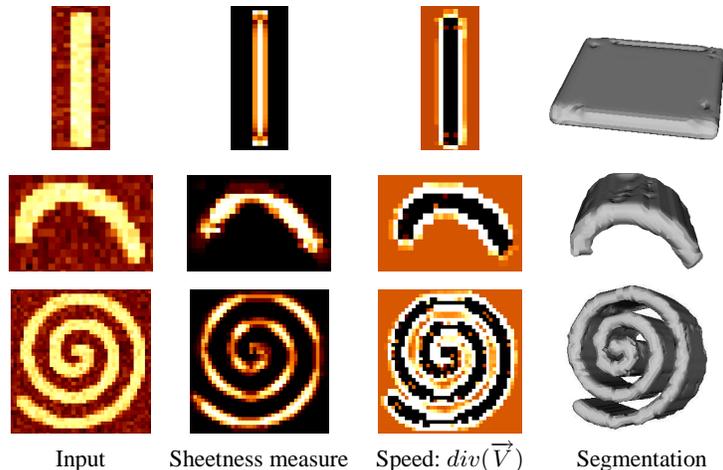


Fig. 3. Experiment on synthetic objects. We show slices of the input volumes, the sheetness measure, the speed term and a surface rendering of the segmentation.

each object, obtained by linear interpolation on the voxel grid. We then added white Gaussian noise to each voxel to simulate typical noise levels in a CT acquisition process (we used a 10% signal-to-noise ratio) and followed the steps detailed above to construct the sheetness measure, obtain the ϕ distribution and then build the vector field \vec{V} , using the same parameters throughout. Empirical surface points were defined as the zero-crossings of the speed term $div(\vec{V})$ in the geometric flow.

A visualization of the important terms is shown in Fig. . We evaluated the accuracy of the zero-crossings of the speed term by computing the average and maximum Euclidean distance errors between each empirical surface voxel and its closest ground truth surface voxel. We also computed the ratio between classified voxels in the segmentation and bone voxels in the binary volume. The results, shown in Table , indicate that the average error is typically less than 0.3 voxels and that the agreement between the reconstructed and original volume is above 90%. We have determined empirically that the maximum errors occur at the two ends of each synthetic object, which is to be expected since mean curvature smoothing causes the most shrinkage there.

object	average distance error (voxels)	maximum error (voxels)	ratio (%)
plate	0.22	1.00	95
rib	0.25	1.12	95
spiral	0.26	1.22	91

Table 3. Agreement level between reconstruction and binary ground truth

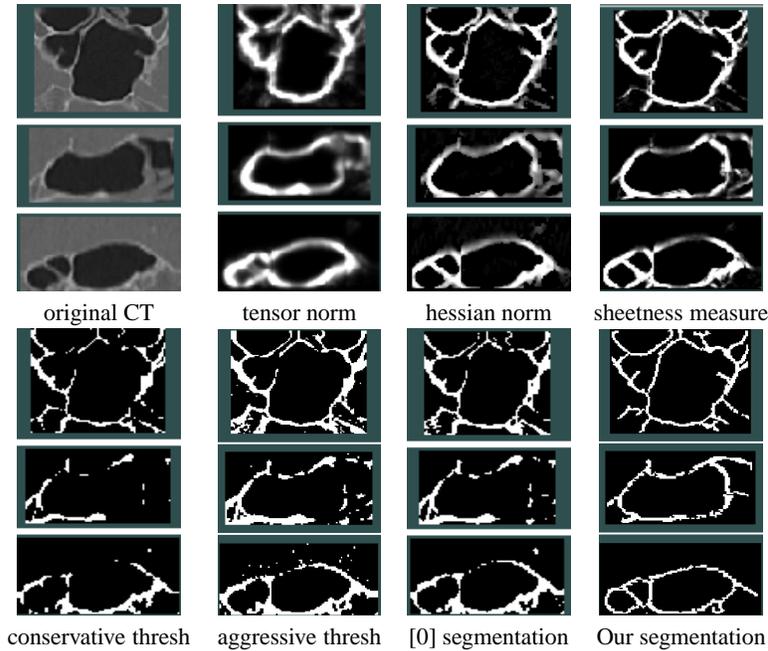


Fig. 4. Comparison between different segmentation methods. Our method connects most of the thin bone structure and reconstructs more bone than the approach in [0] and than thresholding.

5 Bone enhancement and segmentation on CT data

In order to compare our algorithm with Westin et al.’s adaptive thresholding method and simple thresholding, we cropped a 53 mm x 89 mm x 98mm region of a CT data sets around the paranasal sinuses and ran the segmentation methods. The original volume was first resampled to a $0.468mm^3$ isotropic grid.

We make several qualitative observations from Fig. 4. First, we see that both the algorithm of [0] and our method work better than simple thresholding. Second, our method appears to better exploit tissue contiguity and connect more bone structure than Westin’s segmentation. This is to be expected since our geometric flow method is designed to evolve and connect as much sheetness information as possible whereas [0] remains a thresholding method, albeit one that accounts for local structure. We have chosen adaptive thresholding (c_{plane}) parameters which we thought gave the best results and it is clear that, as the authors suggest, if incorporated in an deformable model framework a better segmentation could be obtained. In our experiments with the approach, we have found that the structure tensor picks out the direction of maximum change in intensity but spatially, it does not behave well for our application. The c_{plane} measure is strong mostly at boundaries and tends to thicken the edges as seen in the tensor norm response in Fig. 4. Our method has the advantage of extracting locations at the center of bone structures where the underlying iso-intensity level set behaves like a plane. The sheetness measure

combined with a flow designed to evolve and stop at boundaries performs better as it is able to propagate along low sheetness regions. Most importantly, on its own, the sheetness measure has suitable characteristics to be incorporated in a surgical simulator for pituitary intervention. The measure is high on the center plane of bones and decreases towards the boundaries where there is an uncertainty about the tissue classification of a voxel. Hence, it can allow a surface mesh model to account for uncertainty in determining triangulated bone boundaries, to better model sinus bones in the simulation. For our application, it is thus more useful to have a confidence measure for all voxels than a binary segmentation of bones which certainly still has missing segments and no notion of distance to the boundaries.

6 Conclusions

We have presented a general multi-scale bone enhancement measure that can be used to drive a geometric flow to segment sheet-like structures. The key contribution is the introduction of a sheetness measure based on the properties of the multi-scale Hessian shape operator which has well-founded differential geometric theory and is stable under the presence of noise. The measure gives a confidence index on the presence of bones and for voxels with high values, the scale and direction of the local bone structure is directly obtained. We are able to detect and connect very thin and diffusive bone structure boundaries. In current work, we are trying to use both the bone segmentation as well as the sheetness, vesselness and other tissue classification cues from the CT and corresponding MR data to use a global class competition levelset framework.

References

1. I. Ciric et al. Complications of transsphenoidal surgery: Results of a national survey, review of the literature, and personal experience. *Neurosurgery*, 40(2):225–236, 1997.
2. M. Descoteaux, L. Collins, and K. Siddiqi. A multi-scale geometric flow for segmenting vasculature in mri. In *MICCAI'2004*, September 2004.
3. A. Frangi, W. Niessen, K. L. Vincken, and M. A. Viergever. Multiscale vessel enhancement filtering. In *MICCAI'98*, pages 130–137, 1998.
4. K. Krissian, G. Malandain, and N. Ayache. Model-based detection of tubular structures in 3d images. *Computer Vision and Image Understanding*, 80(2):130–171, November 2000.
5. C. Lorenz, I. Carlsen, T. Buzug, C. Fassnacht, and J. Weese. Multi-scale line segmentation with automatic estimation of width, contrast and tangential direction in 2d and 3d medical images. In *CVRMED-MRCAS'97*, volume 1205, pages 233–242, 1997.
6. Y. Sato et al. and R. Kikinis. 3d multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. *Medical Image Analysis*, 2(2):143–168, 1998.
7. T. B. Sebastian et al. and B. B. Kimia. Segmentation of carpal bones from 3d ct images using skeletally coupled deformable models. In *MICCAI'98*, pages 1184–1194, 1998.
8. A. Vasilevskiy and K. Siddiqi. Flux maximizing geometric flows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(12):1–14, 2002.
9. C.-F. Westin, A. Bhalerao, H. Knutsson, and R. Kikinis. Using local 3D structure for segmentation of bone from computer tomography images. In *CVPR'97*, pages 794–800, 1997.