

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

*Maxime Descoteaux^a, Michel Audette^b, Kiyoyuki Chinzei^b,
and Kaleem Siddiqi^c*

^aOdyssee Team, INRIA Sophia-Antipolis / ENPC-Paris / ENS-Ulm Paris, France

^bSurgical Assist Group, AIST, Tsukuba, Japan

^cSchool of Computer Science, McGill University, Montreal, Canada

Correspondence to:

Maxime Descoteaux
Odyssee Team, INRIA Sophia Antipolis
2004 route des Lucioles
06902 Sophia Antipolis, France
Phone: +33 4 92 38 7589
Fax: +33 4 92 38 7845
Email: Maxime.Descoteaux@sophia.inria.fr

Bone Enhancement Filtering

Abstract

Simulation of pituitary gland surgery raises many interesting challenges since a precise classification of all tissues (soft tissues, vessels, bones) in the pituitary gland area is needed to create a realistic simulator for endoscopic surgery. In particular, bone structures are very thin and have diffusive edges in the CT dataset. Thus, the common method of thresholding the data is insufficient and produces incomplete segmentations that have many holes. In this paper, we focus on a better segmentation of paranasal sinus bones. We present a novel multi-scale bone enhancement measure that can be used to drive a geometric flow to segment any sheet-like structure. The measure is designed to capture sheet-like structure using the local shape information from the eigenvalue decomposition of the Hessian matrix. It is maximum on the center sheet of the bone structure and provides local estimates of the width and orientation of the bone structure. These can be used to create a vector field which is orthogonal to the bone boundaries so that the flux maximizing flow algorithm can be applied to recover the bone structure. Hence, the sheetness 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 quantitatively and show a comparison between the existing segmentation techniques of paranasal sinus CT data.

Key words: Bone segmentation, structure tensor, Hessian matrix, geometric flows, CT, multi-scale analysis, validation.

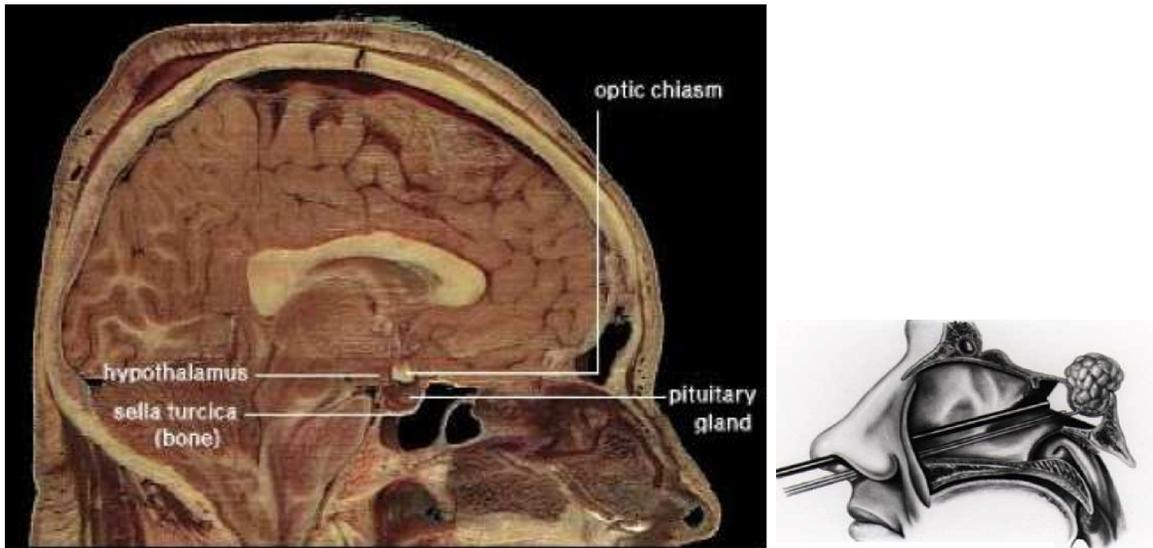


Figure 1: Pituitary gland area anatomy Dr. Jastrow [1] and schematic drawing demonstrating endoscopic pituitary surgery Dr. Jho [2].

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.1). This requires extensive practice and precise knowledge of the anatomy, the absence of which can have serious implications on the patient [3]. 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 formulate a minimally supervised method for producing a set of patient-specific anatomical models, from MR and CT datasets, 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. However, from Fig. 2, 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 datasets. For such thin bones, a simple thresholding procedure gives unsatisfying results.

We introduce a novel algorithm for bone enhancement filtering and segmentation. It is important to mention that even though the approach is motivated by paranasal sinus bone segmentation, it can be used to enhance and segment *any* plate-like structure since it is based on a general multi-scale second order local shape operator. 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 [4–8]. We propose a *sheetness* measure that can be used to drive an active surface to segment bone.

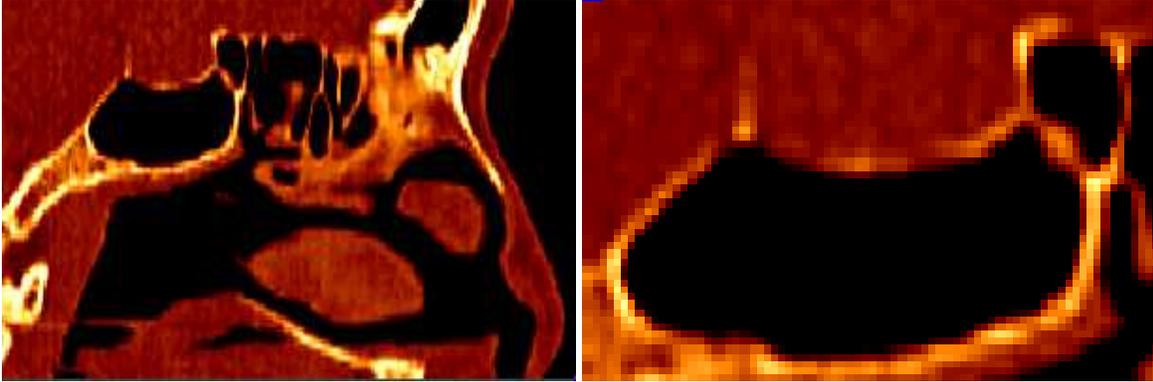


Figure 2: Sagittal slice of CT dataset in region of interest for pituitary gland surgery (see Fig. 1).

This is motivated in part by Frangi's tubular structure enhancement filtering measure [5] and Descoteaux et al [4] 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 [9] 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.

The paper is outlined as follows. In Section 2 we review relevant background material on the use of local 3D structure for segmentation as well as existing active contour algorithms for bone segmentation. We then develop our sheetness measure that can be incorporated in the flux maximizing flow [10], in Section 3. We present several synthetic reconstructions and segmentations for a high resolution CT dataset in Section 4. Finally, we conclude by discussing ongoing work and the way to fit this approach in the tissue classification engine of our surgical simulator.

2 Using Local 3D Structure for filtering and segmentation

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

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

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 [4–8]. 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. [9, 11] 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 positive semi-definite structure tensor \mathcal{T} 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. In that case, tensor \mathcal{T} has rank 1 which means that the two smallest eigenvalues are null, i.e. $\lambda_1 = \lambda_2 = 0$

and $\lambda_3 \gg 0$ (see [9] for more details). The authors then describe a planar measure,

$$c_{plane} = \frac{\lambda_3 - \alpha\lambda_2}{\lambda_3}, \quad (1)$$

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 dataset \mathcal{I} . The threshold at each voxel x is then defined as

$$t(x) = t_0 - \alpha c_{plane}(x), \quad (2)$$

where t_0 is a global threshold manually selected depending on the dynamics of the input image and α is a weight factor on the planar measure.

In Section 5, we demonstrate several properties of this approach. In particular, \mathcal{T} is positive semi-definite which means that all its eigenvalues are positive, i.e. $\lambda_1, \lambda_2, \lambda_3 > 0$. Moreover, 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.

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 $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$.

2.2 Modeling sheet-like structures using the Hessian operator

The Hessian matrix encodes important shape information. Whereas the structure tensor models the first order changes in intensity, the Hessian, looks at second order variations of the data, An eigenvalue decomposition measures the maximum changes in the normal vector (gradient vector) of the underlying intensity iso-surface in a small neighborhood.¹ Hence, it can differentiate between tube-like, sheet-like and blob-like structures. The classification of Table 1 was first explored by Sato et al. [8], and Lorenz et al. [7] separately. In [5], Frangi defines three ratios using tube-like properties of the eigenvalues of Table 1 to separate blood vessels from other structures,

$$R_B = \frac{|\lambda_1|}{\sqrt{|\lambda_2\lambda_3|}} \quad , \quad R_A = \frac{|\lambda_2|}{|\lambda_3|} \quad , \quad R_{noise} = \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}. \quad (3)$$

From Table 1, it can be seen that R_B is non zero only for blob-like and noisy structures. The R_A ratio differentiates sheet-like from tube-like structures. Finally, R_{noise} , the Frobenius norm, is used to ensure that random noise effects are suppressed from the response. For a particular scale σ , the intensity image is

¹Details on the mathematical justification between differential geometry of surfaces and the Hessian operator can be found in [13].

convolved with derivatives of γ -parametrized Gaussian kernels [12] with standard deviation σ to compute the Hessian matrix and the following vesselness response function, $V(\sigma)$, is computed:²

$$V(\sigma) = \begin{cases} 0 & \text{if } \lambda_2 < 0 \text{ or } \lambda_3 < 0 \\ (1 - \exp\left(-\frac{R_A^2}{2\alpha^2}\right)) \exp\left(-\frac{R_B^2}{2\beta^2}\right) (1 - \exp\left(-\frac{R_{noise}^2}{2c^2}\right)) & \end{cases} \quad (4)$$

Due to the multi-scale approach, the *vesselness* measure is 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 [4] 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 [10] can be applied to recover them. This method can recover low contrast and thin vessels from standard anatomical proton density weighted datasets.

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. Fig. 3 shows a sheet-like structure with eigenvectors corresponding to non-zero eigenvalues perpendicular to the sheet. We define two new ratios, R_{sheet} , R_{blob} and reuse R_{noise} defined in Eq. 3, to

²The vesselness expression is given for the case of a dark tubular structure on a brighter background (as in proton density MRI). In the case of angiographic data, the signs in condition 1 must be changed, i.e. $V(\sigma) = 0$ if $\lambda_2 > 0$ or $\lambda_3 > 0$.

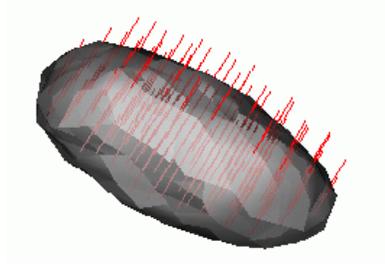


Figure 3: Local sheet-like structure. In red, the eigenvectors corresponding to non-zero eigenvalue are perpendicular to the sheet structure.

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.

differentiate sheet-like structures from others,

$$R_{sheet} = \frac{|\lambda_2|}{|\lambda_3|} \quad , \quad R_{blob} = \frac{|(2|\lambda_3| - |\lambda_2| - |\lambda_1|)|}{|\lambda_3|}. \quad (5)$$

Their behavior is described in Table 2. Then, just as in [5], we can define the sheetness measure, M , as the maximum response over all scales σ at which the derivatives of the Hessian are computed,

$$M = \max_{\sigma \in \Sigma} S(\sigma) = \begin{cases} 0 & \text{if } \lambda_3 > 0 \\ (\exp\left(\frac{-R_{sheet}^2}{2\alpha^2}\right))(1 - \exp\left(\frac{-R_{blob}^2}{2\beta^2}\right))(1 - \exp\left(\frac{-R_{noise}^2}{2c^2}\right)), & \end{cases} \quad (6)$$

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 [4, 5].

Each term of the equation has a function and depends on the characteristics of Table 2. 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. 6, 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.

2.3 Hybrid Modeling with the Structure Tensor and the Hessian operator

One can imagine combining the respective strengths of the structure tensor and Hessian operators into a hybrid operator. Krissian et al [14] do so with a new positive semi-definite and scale invariant descriptor to infer the local shape behavior of the aorta from low quality ultrasound images. Although it is a different problem and another modality, the underlying model uses similar ideas and is also based on a multi-scale eigenvalue decomposition at each voxel to obtain the local orientation and the cross-sectional plane of the aorta. However, the proposed hybrid operator is essentially a generalization of the structure tensor which has similar properties and advantages but also, for our application, has the same limitations as the structure tensor, introduced at the end of Section 2.1 The main problem for our application is the stronger response of the operators involving the structure tensor at bone boundaries. The other problem is the positive semi-definiteness of the operators. Having all eigenvalues with the same sign is actually a disadvantage in our application because one cannot differentiate between grey-to-white or white-to-grey change of intensity in the signal. In the paranasal sinus area, where bone structures are thin and very close to one another, the response from the shape operator must be able to make that distinction. Hence, in this paper, we only use the Hessian operator to perform the eigenvalue decomposition and define the sheetness measure M of Eq. 6.

3 A Geometric Flow for Segmenting Bone Structures

The deformable model is commonly used for segmentation in many computer vision applications and the corresponding literature is very rich, motivated in large part by the classical parametric snakes introduced

by Kass et al. [15]. These models have also been extended to handle changes in topology due to the splitting and merging of contours [16]. However, there have been a few deformable model methods proposed for bone segmentation which are quite different from our approach because they are suited for 2D images from different modalities and different bones ([17–19]). In [17], the segmentation of carpal bones of the hand in CT images is faced with similar challenges as in our sinus bone CT dataset. 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. 6 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 of Eq. 6. 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,

$$\vec{\mathcal{V}} = \phi \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|}. \quad (7)$$

The extended vector field of Eq. 7 explicitly models the scale at which bone boundaries occur, due to the multi-scale nature of the sheetness measure M (Eq. 6) as well as the expected gradient in the direction normal to bone boundaries. Thus, it is an ideal candidate for the static vector field in the flux maximizing geometric flow of Vasilevsky et al [10], $S_t = \text{div}(\vec{\mathcal{V}})\vec{\mathcal{N}}$. The flow evolves the surface S in time t to converge to the zero-crossing of the speed term, $\text{div}(\vec{\mathcal{V}})$. The surface evolution equation works out to be

$$\begin{aligned}
S_t &= \text{div}(\vec{\mathcal{V}})\vec{\mathcal{N}} \\
&= \left[\left\langle \nabla\phi, \frac{\nabla\mathcal{I}}{|\nabla\mathcal{I}|} \right\rangle + \phi \text{div} \left(\frac{\nabla\mathcal{I}}{|\nabla\mathcal{I}|} \right) \right] \vec{\mathcal{N}} \\
&= \left[\left\langle \nabla\phi, \frac{\nabla\mathcal{I}}{|\nabla\mathcal{I}|} \right\rangle + \phi\kappa_{\mathcal{I}} \right] \vec{\mathcal{N}}.
\end{aligned} \tag{8}$$

Here, $\kappa_{\mathcal{I}}$ is the Euclidean mean curvature of the iso-intensity level set of the image. Note that this is a hyperbolic partial differential equation since all terms depend solely on the vector field and not on the evolving surface. We now enumerate several properties of this geometric flow.

1. The first term $\left\langle \nabla\phi, \frac{\nabla\mathcal{I}}{|\nabla\mathcal{I}|} \right\rangle$ acts like a doublet. $\nabla\phi$ has a zero-crossing at bone boundaries and $\nabla\mathcal{I}$ does not change sign. Hence, when the evolving surface overshoots the boundary slightly, this term acts to push it back toward the boundary. Such doublet terms have also shown to be beneficial in earlier geometric flows for segmentation [20–22].
2. The second term behaves like a geometric heat equation since $\kappa_{\mathcal{I}}$ is the mean curvature of the iso-intensity level set of the original intensity image. This equation has been extensively studied in the mathematics literature and has been shown to have remarkable anisotropic smoothing properties [23, 24]. It is also the basis for several nonlinear geometric scale-spaces such as those studied in [25, 26].

3. Combining both terms, it is clear that the flow cannot leak in regions outside vessels since both ϕ and $\nabla\phi$ are zero there. Hence, when seeds are placed at locations where the sheetness measure M is high the flow given by Eq. 8 will evolve toward the closest zero level set of the divergence of the vector field $\vec{\mathcal{V}}$. This is clearly illustrated in the synthetic segmentation examples of Fig. 4.

3.1 Implementation Details

Below we review some of the details of the implementation of our multi-scale geometric flow (Eq. 8), which is based on level set methods [27]. The flow is topologically adaptive due to its implementation with levelset, is computationally efficient and requires minimal user interaction.

1. The ϕ distribution is carried out from voxels at bone structure center sheets since at such locations one has strong confidence in the scale and orientation estimate from sheetness measure. This is done using the following procedure

$$\text{if } (M > \text{threshold} \ \&\& \ \frac{M}{\text{local_max}} > \text{percentile})$$

Distribute sheetness over flat ellipsoid

We use a sheetness threshold of 0.01 and a percentile of 0.75 and local_max is the maximum sheetness response in a 3x3x3 neighborhood of the voxel. The voxels inside the ellipsoidal regions give the seeds used to initialize the geometric flow.

2. The derivatives in the doublet term $\left\langle \nabla\phi, \frac{\nabla\mathcal{I}}{|\nabla\mathcal{I}|} \right\rangle$ are computed using central differences for $\nabla\phi$ and a second-order essentially non-oscillatory (ENO) scheme for the normalized gradient of the input

image [28]. This is essentially because ϕ is by construction relatively smooth since it is constructed by superposition. On the other hand, the gradient of the original intensity image can be sharp and is better captured by an ENO method.

3. $\kappa_{\mathcal{I}}$, the mean curvature of each intensity iso-surface is computed using a 3-neighbor central difference scheme for all derivatives:

$$\kappa_{\mathcal{I}} = \frac{(\mathcal{I}_{yy} + \mathcal{I}_{zz})\mathcal{I}_x^2 + (\mathcal{I}_{xx} + \mathcal{I}_{zz})\mathcal{I}_y^2 + (\mathcal{I}_{xx} + \mathcal{I}_{yy})\mathcal{I}_z^2 - 2(\mathcal{I}_x\mathcal{I}_y\mathcal{I}_{xy} - \mathcal{I}_x\mathcal{I}_z\mathcal{I}_{xz} - \mathcal{I}_y\mathcal{I}_z\mathcal{I}_{yz})}{(\mathcal{I}_x^2 + \mathcal{I}_y^2 + \mathcal{I}_z^2)^{\frac{3}{2}}} \quad (9)$$

4. A first-order in time discretized form of the level-set version of the evolution equation is given by

$$\Psi_n = \Psi_{n-1} + \Delta t * \mathcal{F} * \|\nabla \Psi_{n-1}\|$$

where $\mathcal{F} = \left\langle \nabla \phi, \frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right\rangle + \phi \operatorname{div} \left(\frac{\nabla \mathcal{I}}{|\nabla \mathcal{I}|} \right)$, Ψ is the embedding hypersurface and Δt is the step size. This is now a standard numerical approach for solving partial differential equations of this type since it allows topological changes to occur without any additional computational complexity and can be made efficient using a narrow band implementation.

5. The evolving surface S is obtained as the zero level set of the Ψ function. The numerical derivatives used to estimate $\|\nabla \Psi\|$ must be computed with up-winding in the proper direction as described in [27].

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 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 $\vec{\mathcal{V}}$, using the same parameters throughout. Empirical surface points were defined as the zero-crossings of the speed term $div(\vec{\mathcal{V}})$ in the geometric flow.

A visualization of the important terms is shown in Fig. 4. 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. This measure indicates the degree to which the segmented data explains the ground truth data, but in terms of an average distance error. We also computed the ratio between classified voxels in the segmentation and bone voxels in the binary volume. This measure indicates the degree to which the ground truth data is accounted for by the segmented data. The results, shown in Table 3, 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.

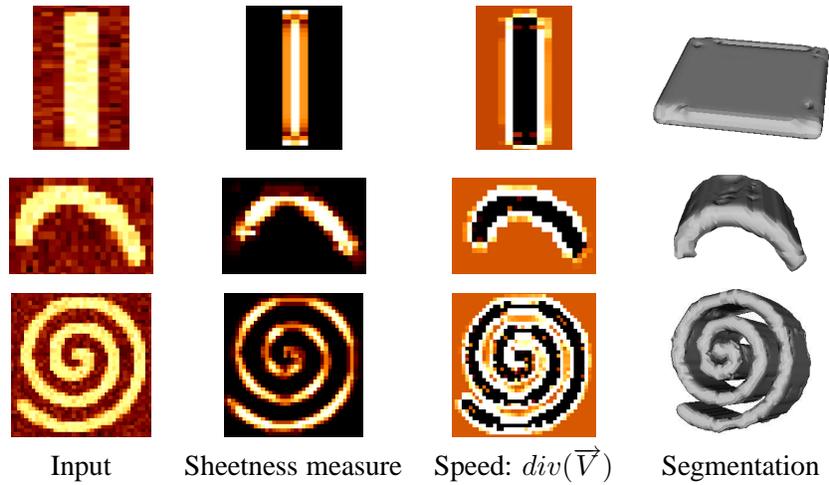


Figure 4: Experiment on synthetic objects. We show slices of the input volumes, the sheetness measure, the speed term and a surface rendering of the segmentation.

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

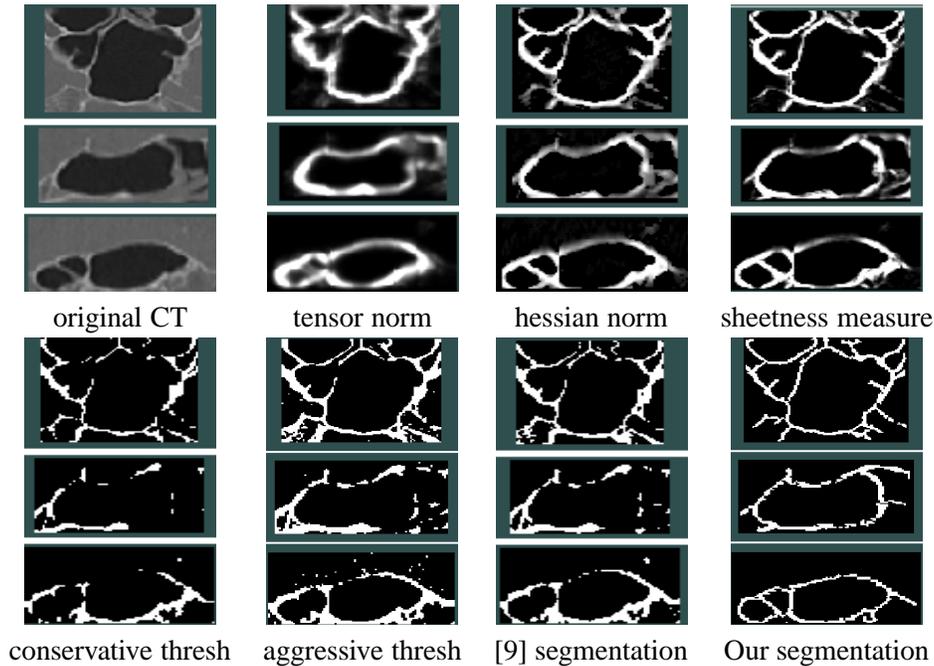


Figure 5: Comparison between different segmentation methods. Our method connects most of the thin bone structure and reconstructs more bone than the approach in [9] and than thresholding.

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 dataset 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 [9] 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 [9] 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 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 any sheet-like structures. The key contribution is the introduction of a the 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. We showed synthetic examples that validate our approach quantitatively and we compared the existing segmentation techniques on paranasal sinus from a CT dataset.

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. Although we have applied our method to paranasal sinus bone segmentation from CT data, it is important to note that the algorithm can be applied to other plate-like bone structures and other modalities. In fact, there is interest in our automatic method for skull segmentation from CT or MRI data.

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 MRI data to use a global class competition levelset framework. This would provide a complete characterization of the tissues in the pituitary gland region which could then employed in a virtual endoscopic surgery simulator.

Acknowledgements

We would like to thank the JSPS and NSERC summer program in Japan for funding this research. A special thank you to Alexandre Thinnes for valuable input in this work.

References

- [1] Dr. Jastrow H. <http://www.uni-mainz.de/FB/Medizin/Anatomie/workshop/VH/female/Filme/Filme.html>
- [2] Dr. Jho HD. http://drjho.com/pituitary_surgery.htm
- [3] Ciric I, Ragin A, Baumgartner C, Pierce D. Complications of transsphenoidal surgery: Results of a national survey, review of the literature, and personal experience. *Neurosurgery*, 40(2):225–236, 1997.
- [4] Descoteaux M, Collins L, Siddiqi K.. A multi-scale geometric flow for segmenting vasculature in mri. In *MICCAI'2004 Part 1*, pp 500–507, 2004.
- [5] Frangi A, Niessen W, Vincken K.L, Viergever M.A. Multiscale vessel enhancement filtering. In *MICCAI'98*, pp 130–137, 1998.
- [6] Krissian K, Malandain G, Ayache N. Model-based detection of tubular structures in 3d images. *Computer Vision and Image Understanding*, 80(2):130–171, November 2000.
- [7] Lorenz C, Carlsen I, Buzug T, Fassnacht C, Weese J. 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, pp 233–242, 1997.

- [8] Sato Y, Nakajima S, Atsumi H, Koller H, Gerig G., Yoshida Y, Kikinis R. 3d multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. *Medical Image Analysis*, 2(2):143–168, 1998.
- [9] Westin C-F, Bhalerao A, Knutsson H, Kikinis K. Using local 3D structure for segmentation of bone from computer tomography images. In *CVPR'97*, pp 794–800, 1997.
- [10] Vasilevskiy A, Siddiqi K. Flux maximizing geometric flows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(12):1–14, 2002.
- [11] Westin C-F, Warfield S, Bhalerao B, Mui L, Richolt J, Kikinis K. Tensor Controlled Local Structure Enhancement of CT Images for Bone Segmentation. In *MICCAI'97*, pp 1205–1212,1998.
- [12] Lindeberg T. Edge detection and ridge detection with automatic scale selection. *International Journal of Computer Vision*, 30(2):77–116, 1998.
- [13] Descoteaux M. A Multi-Scale Geometric Flow for Segmenting Vasculature in MRI: Theory and Validation *School Of Computer Science, McGill University, Msc Thesis*, June 2004.
- [14] Krissian K, Ellsmere J, Vosburgh K, Kikinis R, Westin CF. Multiscale segmentation of the aorta in 3D ultrasound images. In *Engineering in Medicine and Biology Society*. 2003; 638–641.
- [15] Kass M, Witkin A, Terzopoulos D. Snakes: Active contour models. *International Journal of Computer Vision* 1987;1:321–331.

- [16] McInerney T, Terzopoulos D. T-snakes: Topology adaptive snakes. *Medical Image Analysis* 2000; 4:73–91.
- [17] Sebastian T.B, Tek H, Crisco J.J, Kimia B.B. Segmentation of carpal bones from 3d ct images using skeletally coupled deformable models. In *MICCAI'98*, pp 1184–1194, 1998.
- [18] Ballerini L, Bocchi L. Bone segmentation using multiple communicating snakes. In *SPIE Medical Imaging: Image Processing*, volume 5032. 2003; 1621–1628.
- [19] Lorigo LM, Faugeras OD, Grimson WEL, Keriven R, Kikinis R. Segmentation of bone in clinical knee mri using texture-based geodesic active contours. In *MICCAI*. 1998; 1195–1204.
- [20] Kichenassamy S, Kumar A, Olver P, Tannenbaum A, Yezzi A. Gradient flows and geometric active contour models. In *International Conference On Computer Vision*. 1995; 810–815.
- [21] Caselles V, Kimmel R, Sapiro G. Geodesic active contours. In *International Conference On Computer Vision*. 1995; 694–699.
- [22] Siddiqi K, Lauzière YB, Tannenbaum A, Zucker SW. Area and length minimizing flows for shape segmentation. *IEEE Transactions on Image Processing* 1998;7(3):433–443.
- [23] Gage M, Hamilton R. The heat equation shrinking convex plane curves. *Journal of Differential Geometry* 1986;23:69–96.
- [24] Grayson M. The heat equation shrinks embedded plane curves to round points. *Journal of Differential Geometry* 1987;26:285–314.

- [25] Alvarez L, Guichard F, Lions PL, Morel JM. Axiomes et équations fondamentales du traitement d'images. *C R Acad Sci Paris* 1992;315:135–138.
- [26] Kimia BB, Tannenbaum A, Zucker SW. Shape, shocks, and deformations I: The components of two-dimensional shape and the reaction-diffusion space. *International Journal of Computer Vision* 1995; 15:189–224.
- [27] Osher SJ, Sethian JA. Fronts propagating with curvature dependent speed: Algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics* 1988;79:12–49.
- [28] Osher S, Shu CW. High-order essentially non-oscillatory schemes for Hamilton-Jacobi equations. *SIAM Journal of Numerical Analysis* 1991;28:907–922.