Anatomical Medial Surfaces with Efficient Resolution of Branches Singularities

Debora Gil[a,c], Sergio Vera[b], Agnés Borràs[a], Albert Andaluz[a], Miguel A. González Ballester[d,e]

[a] *Computer Vision Center, Computer Science Dept., Campus UAB, 08193 Bellaterra, Barcelona, Spain*
[b] *Alma IT Systems, C/ Vilana 4B, 4-1, Barcelona 08022, Spain*
[c] *Serra Hunter Fellow*
[d] *Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain*
[e] *ICREA, Barcelona, Spain*

Abstract

Medial surfaces are powerful tools for shape description, but their use has been limited due to the sensibility of existing methods to branching artifacts. Medial branching artifacts are associated to perturbations of the object boundary rather than to geometric features. Such instability is a main obstacle for a confident application in shape recognition and description. Medial branches correspond to singularities of the medial surface and, thus, they are problematic for existing morphological and energy-based algorithms. In this paper, we use algebraic geometry concepts in an energy-based approach to compute a medial surface presenting a stable branching topology. We also present an efficient GPU-CPU implementation using standard image processing tools. We show the method computational efficiency and quality on a custom made synthetic database. Finally, we present some results on a medical imaging application for localization of abdominal pathologies.

Keywords: Medial Representations, Shape Recognition, Medial Branching Stability, Singular Points

1. Introduction

The presence of anatomic abnormalities in organs shape, like protrusions or deformations, is an indicator of a variety of pathologies, such as cardiac hyperthrophy (Rohilla et al. (2012)) or hippocampal abnormalities in schizophrenia (Csernansky et al. (2002)). A main field in medical imaging methods is oriented to the analysis of the anatomy by the means of the shape deviation of the shape of the organs to provide better decision support systems for diagnosis (Liu et al. (2010); Yao and Summers (2009)) and treatment planning (Stough et al. (2007); Pizer et al. (2005)).

Medial structures completely determine the geometry of the boundary volume (Gray (2004)). Thus, they could be used to characterize pathological abnormalities (Styner et al. (2004)) and provide more interpretable representations of complex organs (Yao and Summers (2009)). Further, medial structures are the basis for the definition of tubular coordinates (Gray (2004)) which can model, both, the organ boundary as well as its interior (Blum (1967)). Finally, by setting equal values to equivalent anatomical sites tubular coordinates allow the location of specific anatomical regions across patients and time (Garcia et al. (2010); Vera et al. (2014); Wang et al. (2005)). This constitutes a main advantage over more conventional boundary-based and volumetric representations of the data.

In spite of their big potential to help in diagnosis and treatment planning, the use of medial structures in systematic clinical practice is still very limited. In our opinion, the main barrier for a systematic application is the presence of artifacts arising in their digital computation (Yushkevich (2009); N.Faraj et al. (2013)).

A main requirement for a confident representation of shapes is the stability of medial manifolds under perturbations of the object boundary (Giblin et al. (2009)). Existing methods for computation of medial surfaces often generate spikes or loose connectivity at main branches. This lack of stability prevents from using medial structures directly in

most applications (not only medical (N.Faraj et al. (2013))). In the particular context of shape modelling and description in medical applications, extra branches complicate statistical modelling of patient populations. Furthermore, complex branching medial geometry also hinders the definition of the medial tubular coordinate (Sun and et al. (2010); Terriberry and Gerig (2006)). Finally, medial lack of stability implies that not all branches are meaningful from neither the application point of view nor the object boundary geometrical features.

Stability of medial properties depends on the domain on which the medial manifold is computed. Existing methods compute medial structures on either a tetrahedral mesh of the volume boundary or directly on the volumetric voxel domain.

Mesh methods are based on the Voronoi tetrahedral mesh of a set of points sampled on the object boundary (Dey and Zhao (2002); Amenta et al. (2001); Giesen et al. (2009)) and can naturally resolve branching medial surfaces. However, they introduce 1 dimensional spikes associated to boundary irregularities that have to be further pruned (Amenta et al. (2001); Giesen et al. (2009)). Although some recent methods (Giesen et al. (2009)) are capable of efficiently dealing with surface perturbations, they are prone to introduce medial loops that distort the medial topology (N.Faraj et al. (2013)) in a way that could be erroneously considered a pathology. Also their computational cost and quality depend on the number of vertices defining the volume boundary mesh and, thus, on the volume resolution (Dey and Zhao (2002)). Finally, in the context of medical applications, the voxel discrete domain is the format in which medical data are acquired from medical imaging devices and, thus, it is the natural domain for the implementation of image processing (Khalifa et al. (2010); Dinguraru et al. (2010)) and shape modelling (Park et al. (2003)) algorithms.

Although it is possible to convert from the discrete voxels to continuous meshes and viceversa, several pre-processing steps (such as smoothing and decimation) are required in order to apply Voronoi methods to clinical volumetric data of large size. These processing steps add computational complexity and inaccuracies due to data interpolation and round-offs, which advises against the use of surface methods in medical data.

Volumetric approaches can be classified into two big types: morphological thinning and energy-based methods. Morphological methods compute medial manifolds by iterative thinning of the exterior layers of the volumetric object until more thinning breaks surface topology (Bouix et al. (2005); Siddiqi et al. (2002); Ju et al. (2007); Svensson et al. (2002)). Meanwhile, energy-based approaches define medial structures as singular points of energy maps, usually given by ridges of the distance map to the object boundary (Vera et al. (2012a); Bouix et al. (2005)).

Morphological methods, while simple, generate completely different results based on the connectivity, but also depending on the different simplicity and ordering criteria used to determine what pixels can be removed or not. Often the resulting manifold contains many spurious branches that are of little use to many applications and need to be removed using pruning methods (Pudney (1998); Amenta et al. (2001)). There are numerous different techniques that deal with spurious branches, such as pruned Voronoi skeletons, PDE-based methods, or template-based methods like M-Reps. However, existing medial simplifications have the following disadvantages for a satisfactory applicability to medical applications. In the case of skeleton pruning, spike/branch removal is controlled by some filtering over some geometric conditions, including ratio of geodesic distances (Ogniewicz and Ilg (1992)), angle between generating points (Dey and Zhao (2002)), distance to boundary (Chazal and Lieutier (2005)) or spike size (Ju et al. (2007); Pudney (1998)). Although all these criteria are well suited for describing and removing medial surface noise (Amenta et al. (2001); Dey and Zhao (2002)), none of them is able to identify the relevance of the branch in the boundary geometric description (Giesen et al. (2009)). The relation between branch simplification and volume reconstruction accuracy is of primary importance for shape modelling applications and, thus, has experienced an increasing research interest in the last years (Giesen et al. (2009); N.Faraj et al. (2013); Vera et al. (2012b)). Concerning template-based methods like CM-Reps (Yushkevich (2009)), they maintain branching topology and, thus, they can only fit target anatomy approximately (Sun and et al. (2010)). Although for some simple structures (like the hippocampus), the approximation error is quite small (Sun et al. (2008)), they are prone to fail at properly modelling and detecting pathological deformations. Finally, PDE-based methods (Siddiqi et al. (2002)) also implement a morphological thinning and, thus, undergo the same surface test

than morphological thinning.

In the case of energy-based approaches the definition of the ridge map and its further binarization play a prominent role in the stability and quality of medial manifolds. The usual approach is thresholding a ridge operator based on image intensity (Bouix et al. (2005)). The definition of the threshold is a delicate issue that strongly depends on the application and object geometry. Our recent methods (Vera et al. (2013)) using a Non-Maxima Suppression (NMS) scheme for extracting medial surfaces are more independent of the thresholding value, especially when a ridge detector that combines the advantages of steerable filters and level sets geometry is used (Vera et al. (2013)). Another advantage of ridge-based NMS methods compared to thinning approaches is their capability for producing more stable medial surfaces without spikes (Vera et al. (2012a)). A main disadvantage is that these methods might suffer from common pitfalls of the ridge detection algorithms (Vera et al. (2013)), which are prone to fail at medial self-intersecting branches as the direction of the ridge is not properly defined. In another work (Bouix et al. (2005)) a method based on ridges and thinning partially corrected these problems, although it inherits the flaws of thinning methods instead.

This paper presents a new approach based on algebraic geometry to describe the complexity of level sets singular points in any dimension using generic filters that codify changes in a single direction. In the context of algebraic geometry, singular points are of special interest and there is a rich literature on their properties and resolution. Algebraic description of singular loci has already been used to classify medial branches transitions under boundary deformations (Giblin et al. (2009)). The novelty of our paper consists in applying Blowups theory (Hironaka (1964)) for algebraic curves desingularization to improve the performance of energy-based approaches using a scheme that admits efficient implementation using a CPU-GPU parallelization scheme easy to reproduce. Our algebraic strategy is used to improve a ridge-based computation of medial surfaces recently published by the authors (Vera et al. (2013)), to provide stable branch geometry that properly describes the anatomy of the volume boundary.

The contents are organized as follows. Section 2 summarizes state-of-art ridge based operators. Section 3 presents the resolution of medial singularities using blowup, as well as, its efficient implementation. Section 4 describes the experimental set-up for validation of medial branch stability and applicability to medical applications. Section 5 presents our experiments on medial stability and medical applicability. Finally, conclusions and future work are exposed in Section 6.

2. Medial Surfaces Based on Ridge Operators

The computation of medial manifolds based on an energy map splits into computation of a medial map from the original volume and binarization of such map. In the case of NMS ridge-based approaches (Vera et al. (2013)), the medial map is given by the ridges of the distance to the object boundary and binarization is obtained by applying NMS to the ridge map.

Ridge detectors are based either on level sets geometry (Lopez et al. (1999)) or image intensity profiles (Freeman and Adelson (1991)). Geometric ridge detectors define ridges as lines joining points of maximum curvature of the distance map level sets. Some techniques (Lopez et al. (1999)) compute it using the divergence of the maximum eigenvector of the structure tensor of the distance map, D, to the shape boundary:

$$\mathrm{NRM} := \mathrm{div}(\tilde{V}) = \partial_x \tilde{V}_1 + \partial_y Q \tilde{V}_2 + \partial_z \tilde{V}_3 \qquad (1)$$

for $\tilde{V} = (\tilde{V}_1, \tilde{V}_2, \tilde{V}_3)$ the primary eigenvector of the structure tensor reoriented along ∇D:

$$\tilde{V} = \mathrm{sign}(< \vec{V} \cdot \nabla D >) \cdot \vec{V}$$

being $< \cdot >$ the scalar product. A main advantage is that NRM $\in [-\mathcal{N}, \mathcal{N}]$, being $\mathcal{N}$ the dimension of the embedding space, so medial points can be detected using a common threshold. A main concern (Vera et al. (2013)) is its significant drop at branches and the generation of internal holes in medial surfaces.

Intensity-based maps are computed by convolution with a bank of steerable filters (Freeman and Adelson (1991)) given by derivatives of oriented anisotropic 3D Gaussian kernels. Let $\sigma = (\sigma_x, \sigma_y, \sigma_z)$ be the scale of the filter and $\Theta = (\theta, \phi)$ its orientation given by the unitary vector $\eta = (cos(\phi)cos(\theta), cos(\phi)sin(\theta), sin(\phi))$, then the oriented anisotropic 3D Gaussian kernel, g_σ^Θ, is given

by:

$$g_{\sigma}^{\Theta} = g_{(\sigma_x,\sigma_y,\sigma_z)}^{(\theta,\phi)} = \frac{1}{(2\pi)^{3/2}\sigma_x\sigma_y\sigma_z} e^{-\left(\frac{\tilde{x}^2}{2\sigma_x^2}+\frac{\tilde{y}^2}{2\sigma_y^2}+\frac{\tilde{z}^2}{2\sigma_z^2}\right)} \tag{2}$$

for $(\tilde{x},\tilde{y},\tilde{z})$ the change of coordinates given by the rotations of angles θ and ϕ that transform the z-axis into the unitary vector η. The second partial derivative of g_{σ}^{Θ} along the $\tilde{z}$ axis constitutes the principal kernel for computing ridge maps:

$$\partial_z^2 g_{\sigma}^{\Theta} = (\tilde{z}^2/\sigma_z^4 - 1/\sigma_z^2) g_{\sigma}^{\Theta} \tag{3}$$

We note that by tuning the anisotropy of the Gaussian, we can detect independently medial surfaces and medial axes. For detecting sheet-like ridges, the scales should be set to $\sigma_z < \sigma_x = \sigma_y$, while for medial axes they should fulfill $\sigma_z < \sigma_x < \sigma_y$.

The maximum response for a discrete sampling of the angulations and the scales gives the medial map:

$$SGR := \max_{i,j,k} \left(\partial_z^2 g_{\sigma_k}^{\Theta_{i,j}} * D\right) \tag{4}$$

for $\Theta_{i,j}$ given by $\theta_i = \{i\frac{\pi}{N}, \forall i = 1,\ldots,N\}$ and $\phi_j = \{j\frac{\pi}{M}, \forall j = 1,\ldots,M\}$ and $\sigma_k = (\sigma_x^k, \sigma_y^k, \sigma_z^k) = (2^{k+1}, 2^{k+1}, 2^k), k = [0, K]$. A main advantage of using steerable filters is that their response does not decrease at self-intersections. Their main counterpart is that their response is not normalized, so setting the threshold for binarization becomes a delicate issue (Bouix et al. (2005)).

Given that geometric and intensity methods have complementary properties, some authors propose (Vera et al. (2013)) to combine them into a Geometric Steerable Medial Map (GSM2):

$$GSM2 := SGR(NRM) \tag{5}$$

GSM2 generates medial maps with good combination of specificity in detecting medial voxels while having good characteristics for NMS binarization, which does not introduce internal holes.

NMS consists in checking the two neighbors of a pixel in a specific direction, $V = (V_x, V_y, V_z)$, and deleting pixels if their value is not the maximum one. Let $\mathcal{M}$ be a generic medial map, then its NMS map along the direction V is given by:

$$NMS_{\mathcal{M}}(x,y,z) = \begin{cases} \mathcal{M}(x,y,z) & \text{if } \mathcal{M}(x,y,z) > \\ & > \max(\mathcal{M}_{V+}, \mathcal{M}_{V-}) \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

for $\mathcal{M}_{V+} = \mathcal{M}(x+V_x, y+V_y, z+V_z)$ and $\mathcal{M}_{V-} = \mathcal{M}(x-V_x, y-V_y, z-V_z)$. A thresholding of $NMS_{\mathcal{M}}$ produces 1-pixel wide surfaces. The search direction for local maxima is given by the primary eigenvector of the medial map structure tensor, $ST_{\rho,\sigma}(\mathcal{M})$ and the optimal threshold value can be set by Otsu thresholding in the case of the normalized ridge operators, like NMR and GSM2 (Vera et al. (2013)).

3. Medial Surfaces Preserving Branches

Even if the medial map achieves a uniform response at branches, binarization using NMS is likely to break branch connectivity. The NMS step keeps points achieving a local maximum along a direction that represents the normal to medial surfaces. It follows that NMS is consistent as far as surfaces have a well-defined unique normal vector that can be computed by means of the structure tensor. Branches are loci of surface self-intersections and, thus, their normal space is generated by the normal vectors of the surface intersecting folds. This singular feature influences the computation of NMS from, both, a theoretical and a practical point of view. On one hand, from a theoretical point of view, the definition of NMS should take into account multiple search directions at branching points. On the other hand, the primary eigenvector of the structure tensor used to compute NMS is an average of the folds normal vectors prone to be oriented in a non-maximal direction.

Figure 1 shows NMS artefact at branches for a 2D cut of a liver. We show the response to GSM2 in the left image, the $ST_{\rho,\sigma}(GSM2)$ primary eigenvector for computation of NMS and the NMS map before thresholding. Even if the medial response preserves branch connectivity, NMS breaks it, as the right image close-up shows. This is due to a deficient search direction given by the $ST_{\rho,\sigma}(GSM2)$ primary eigenvector, which fails to be oriented along neither of the branch intersecting segments.

3.1. Resolution of Singularities by Blowups

From a geometric point of view, branching points are singular points in the sense that the codimension of the variety at that point is higher than at the remaining regular points. Singular points commonly arise in the context of algebraic varieties (i.e. zero-sets of a polynomial) as points of multiplicity higher than one. Given that their geometric properties are particular and standard tools of differential geometry do not apply, resolution of singularities has been extensively studied Kollár (2007).

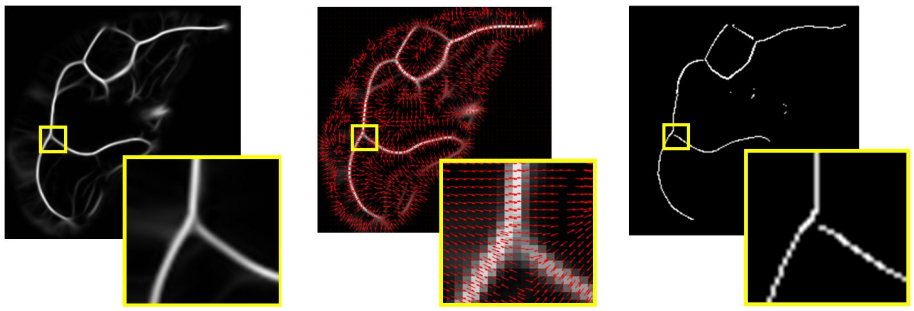

(a) (b) (c)

Figure 1: Rupture of branch connectivity arising from NMS binarization: (a) GSM2 response. (b) Structure tensor orientations. (c) NMS binarization with broken connectivity.

In the context of algebraic geometry, resolution of singularities for a given algebraic variety X tackles with the problem of finding a non-singular variety X' such that there is a surjective differential map $\pi : X' \to X$. In plain words, the desingularization of X is a variety regular (smooth) everywhere probably living in a higher dimensional space such that its projection onto X produces its singularities (Hironaka (1964)).

Blowups are the algebraic tool for constructing resolution of singularities. Blowups are transformations that iteratively untie each singularity of the variety and it is guaranteed (by Hironaka theorem (Hironaka (1964))) that in finite number of steps the variety is resolved. Each blowup is formulated around a singular submanifold of X (called center of the blowup) in terms of its normal space as follows (Hironaka (1964)). Let us assume that $X \subset W(=\simeq \mathbb{R}^n)$ and let $Z \subset X$ be a submanifold of singular points of $\dim Z = d$. For any $a \in Z$, the ambient space W can be locally decomposed as $W = Z \times U \simeq \mathbb{R}^d \times \mathbb{R}^{n-d}$, for U a submanifold transversal to Z of complementary dimension $n-d$. Write points $w \in W$ as pairs $w = (w_1, w_2)$ with $w_1 \in Z$ and $w_2 \in U$ and consider the projective space of lines through a in U, $\mathbb{P}(U) = \mathbb{P}^{n-d-1}$. Then, the blowup is the closure of:

$$\Lambda = \{(w, l_w),\ w \in W \backslash Z, l_w \in \mathbb{P}(U)\} \subset W \times \mathbb{P}^{n-d-1}$$

and $\pi : W \times \mathbb{P}^{n-d-1} \longrightarrow W$ is given by the projection onto the first factor. Around the singular point $a \in Z \subset X$, $\pi^{-1}(X)$ assigns points $(a, l_w(a))$ for $l_w(a)$ a line normal to X at a. Intuitively, the blowup unties X by adding coordinates that represent the normal space at each point of the variety. An important remark from a practical point of view is that in the case of self-intersection of folds with distinct normal vectors, a single blowup is enough to resolve each singularity.

Figure 2 illustrates the blowup of the curve $X = \{-x^3 + 3x^2 = y^2\}$ in the plane $W = \mathbb{R}^2$. Since in this case the center $Z = \{(0,0)\}$ is of codimension 2, the blowup space is $\mathbb{R}^2 \times \mathbb{P}^1 = \mathbb{R}^2 \times \mathbb{R} \cup \{\infty\}$ and can be identified with $X' = \{xz - y = 0\}$. The z-axis of $\mathbb{R}^2 \times \mathbb{P}^1$ represents the directions of all lines through $a = (0,0)$. In particular, the 2 normal lines of X at $(0,0)$ (plotted in red dashed) are lifted to the points $(0, 0, \tan(\theta_1))$ and $(0, 0, \tan(\theta_2))$ for θ_1 and θ_2 the directions of the 2 normal lines. Since $\theta_1 \neq \theta_2$, the two branches crossing at Z are completely unfolded in the blowup space and, thus, X' is regular.

3.2. Computation of Blowups using Image Processing Tools

For algebraic varieties given explicitly as zero-sets of a polynomial, there exist generic computational algorithms for implementing blowups (Bodnár and Schicho (2000)). However in the case of computation of medial surfaces from anatomical volumes, surfaces are local maxima of functions given in digital format by medial maps and without a polynomial formulation. In this section, we explain how to implement blowups for medial surfaces implicitly given by local maxima of a medial map, which we will denote by $\mathcal{M}$.

A blowup is completely specified by defining its center Z and the projective space of directions, $\mathbb{P}^{n-d-1}$, of the complementary manifold. The center of each blowup is given by a connected submanifold of singular points of X. Singular points arise from X self-intersections and are characterized by an increase of the variety codimension (nor-

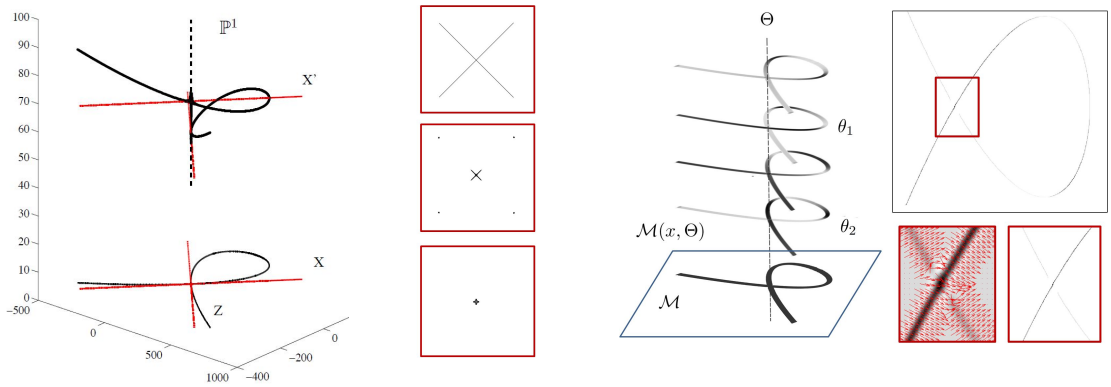

(a) (b) (c)

Figure 2: Scheme of the blowup of a curve in the plane: (a) Analytic blowup. (b) Singularity localization using corner detector: Blowup center computation (c) Blowup space implementation using steerable filters: Blowup Steerable Space and NMS in θ_2 blowup level

mal space dimension) at Z. The eigenvalue decomposition of the structure tensor $ST_{\rho,\sigma}(\mathcal{M})$ describes the geometry of the image level sets and it is commonly used to detect their singular points (corners, junctions and boundary points). Eigenvectors are aligned with the direction of extreme contrast change and eigenvalues are given by the amount of contrast change. At a regular point, the primary eigenvector is normal to level-sets and its eigenvalue is significantly larger than the other two. Meanwhile, self-intersecting level-sets are characterized by similar eigenvalues and corresponding eigenvectors normal to each of the branches. It follows that the ratio between eigenvalues has been exhaustively used in image processing for detecting singular points. Given that, our centers are given by intersection of two or more regular surfaces, we have chosen an the following corner detector (Shi and Tomasi (1994)):

$$\frac{\lambda_1\lambda_2}{\lambda_1+\lambda_2+\epsilon} \tag{7}$$

for $\lambda_1 \geq \lambda_2 \geq \lambda_3$ the eigenvalues of the structure tensor. The operator (7) gives a high response at branches, as well as surface boundary end-points. In order to remove response at surface boundaries, we iterate (7) twice. Given that the response to (7) is regular at boundaries, while it keeps the branching structure at self-intersections, one iteration suffices to remove non-branching responses. Therefore, the set of singular points of X is given by applying twice the operator (7) to the distance map, D, to the object boundary:

$$Z := T(T(D)) \tag{8}$$

The central images in figure 2 illustrate the computation of the set of singular points using the double operator (8). The top image shows a cross with 4 end-points, the central image the response $T(D)$ and the bottom image $Z = T(T(D))$. Gray intensities are shown inverted (the darkest, the largest) for visualization purposes. The first iteration of the corner detector has response at the cross central branching point and the cross end-points of a similar value. Their only difference is in their shape, which is round at end points and like the cross itself at the branching point. It follows that a second iteration of the corner operator removes end-point responses while still preserving a clear response at the branch site.

Concerning the definition of surface orientations, streerable filters constitute a useful tool since they decouple the space of possible orientations for medial surfaces. Given that medial surfaces are local maxima of $\mathcal{M}$, the convolution:

$$\mathcal{M}(x,\Theta) := \partial_z^2 g_\sigma^\Theta * \mathcal{M} \tag{9}$$

defines a function:

$$\begin{array}{ccc} \mathbb{R}^n \times \mathbb{P}^{n-1} & \longrightarrow & \mathbb{R} \\ (x,\Theta) & \mapsto & \mathcal{M}(x,\Theta) \end{array}$$

that restricted to the medial manifold implicitly defines its blowup.

For each orientation, Θ_0, the convolution (9) is only sensitive (i.e. attains its maximum values)

to points of the medial surface oriented perpendicularly to the direction given by Θ_0. Therefore, the blowup volume, $\mathcal{M}(x, \Theta_0)$, is a regular surface with a well defined tangent space. Consequently, the NMS operator given by (6) applied separately to each level $\mathcal{M}(x, \Theta_0)$ properly restores all medial branches perpendicular to Θ_0. This leads to the following implementation of Blowup Maps (which we will note $\mathcal{BUM}$) based on steerable filter:

$$\begin{aligned}\mathcal{BUM}(x) :=& max_{i,j}\left(NMS(\mathcal{M}(x, \Theta_{ij}))\right) = \\ =& max_{i,j}\left(NMS(\partial_z^2 g_\sigma^{\Theta_{ij}} * \mathcal{M})\right) \end{aligned} \quad (10)$$

for $\partial_z^2 g_\sigma^{\Theta_{ij}}$ given as in (4).

The rightmost image in figure 2 illustrates the blowup defined using steerable filters for the case of a shape having as medial surface the lace curve X shown in fig.2 left hand side. Like blowup localization using corner operators, gray intensities are shown inverted (the darkest, the largest) for visualization purposes. The bottom image shows a 2D medial map having X as maximal variety. Its blowup variety X' is given by the 3D volume $\mathcal{M}(x, \Theta)$ shown above the medial image. We show several cuts of the volume corresponding to some orientations of the steerable filters Θ_{ij}. The orientations achieving maximum response at the curve singular point are labelled θ_1, θ_2 and correspond to the curve branch normal directions at 45 and -45 degrees, respectively. First we note that they are at different levels of the image volume $\mathcal{M}(x, \Theta)$ and, thus, the blowup manifold projecting onto X is regular. Second, the corresponding image cuts (shown in side images) have curves of maximal response regular without branches, so that, the direction defining NMS is well-defined (as the bottom close-ups illustrate).

We would like to note that by using the well-stablished mathematical theory of Blowups, our proposal goes beyond existing approaches in several aspects.

A main strength is the capability to describe and reconstruct singular points of arbitrary dimension and complexity using only generic directional filters. In this sense, the theory of blowups guarantees that singular points resolution can be achieved in any case by iterative transformations of the original space into a higher dimensional spaces that add level sets tangent orientations to their x-y coordinates to achieve a regular level set representation. Being tangent spaces represented by derivatives of oriented anisotropic filters, blowup theory implies that convolution with such a generic bank of filters will describe and reconstruct singular points regardless of their complexity and level set dimension. Aside, Blowups theory ensures that the capability of a method for singular point reconstruction is independent of input image resolution since it is mainly concerned with the local geometry of image level sets.

An implementation using generic filters is a main advantage over methods using also steerable filters to construct ad hoc filters that match singular points local structure, like the pioneer work for 2D images described in (Perona (1992)). First, if the filter bank does not contain the specific local configuration of the junctions that we want to detect inside an image, the filter bank will not produce a response strong enough to be kept after a later NMS binarization. Another drawback of strategies based on specific filters that describe the local geometry of singular points is that they are not scalable to higher dimensions. This follows by the increasing complexity and variety of singularities local structure in higher dimensions, which makes the construction of specific filter banks covering all possibilities computationally prohibitive.

Another added value for the medical diagnose assessment domain is that the method admits an efficient implementation in volumetric domains easy to reproduce using the CPU-GPU parallelization scheme described in next Section.

3.3. Efficient Parallelization of BlowUps

Our implementation of blowups can be split in 3 stages (see table 1): blowup location, space decoupling and NMS computation. The most demanding steps are space decoupling and NMS computation. Space decoupling requires convolution of the volume with a filter bank of gaussians covering all possible orientations in 3D. Binarization using NMS has suboptimal performance due to a voxel-to-voxel eigenvalue decomposition of the structure tensor computed for each response to the filter bank. Given that convolution is independent for each filter and eigenvalue decomposition is a voxel-wise operation, both steps can be easily parallelized. In order to make the most of the computational resources available in a standard PC, we have chosen a combined CPU-GPU parallelization for MatLab.

Task	Implementation	Cost	Parallelized
Blowup location	Tomasi corner	Low	No
Space decoupling	Filter bank	High	Yes (GPU)
NMS computation	Structural tensor	High	Yes (CPU)

Table 1: Blowup implementation stages.

3.3.1. GPU Space Decoupling

The simplest approach for parallelization is to perform each space decoupling in a separate CPU core. Using Matlab's Parallel computing toolbox, we distribute each decoupling computational task into independent child worker processes. Matlab then schedules workloads, issues execution request and retrieves the results with no user intervention. This approach, while simple, might require an enormous amount of host RAM memory for its execution because each worker is an almost complete copy of the program, which shares some data with its parent. Therefore, its use becomes impractical as soon as the size of the input volumes increases, as it is the case of 3D anatomical volumes.

An alternative to Matlab CPU parallelization is the use of GPU programming. Consumer grade graphic cards feature a high number of computing threads. Individually, these threads are slower in comparison to a CPU core. However, GPUs are good at parallelism since there are more cores available for computation (Satish et al. (2009)). Thus, for large amounts of data, the benefits of GPU convolution become apparent. Finally, fast tools for GPU computing in Matlab exist nowadays (Pryor et al. (2011)), so experimentation is now possible.

Since the GPU is a separate processor, with dedicated memory, we have to copy the input tensor volumes from the host into the GPU before computation starts. Moreover, after execution the results must be transferred back into the host. Given enough device memory available, it would be possible to convolve the image with the whole filter bank in a single GPU computation operation. However, since our GPU has less memory than our host RAM, we can only convolve the input image with a single filter at each time. This introduces idle computation stages due to these transfers but the lower global execution time in GPU computation validates this approach.

3.3.2. NMS computation

The eigenvalue decomposition is based on an existing implementation (Hicklin et al.). Due to the algebraic properties of the decomposition, there is no spatial dependency between voxels. For this reason, it is possible to divide the original tensor volume in separate data blocks that the CPU can process in parallel. Next, we reassemble the results into a final volume. Figure 3 shows an overview of the proposed parallelization stages.

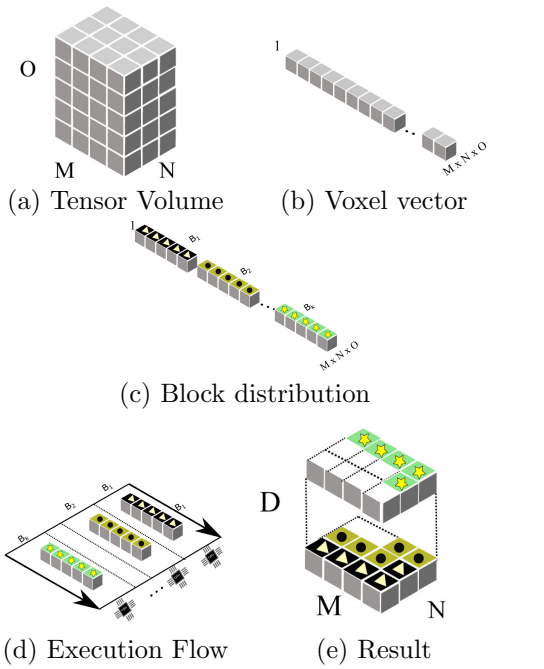

(a) Tensor Volume (b) Voxel vector

(c) Block distribution

(d) Execution Flow (e) Result

Figure 3: Parallel decomposition for NMS computation

Given a tensor of size $M \times N \times O$ voxels, we can represent this volume as a one-dimensional vector of voxels like in figure 3. Assuming K_{cores} CPU computational cores, we can distribute up to K_{cores} blocks, one for each core, with a size for each block, B_i $i = 1, \ldots, K_{cores}$, given by:

$$B_i = \begin{cases} [\frac{v}{K_{cores}}] & i = 1..K_{cores} - 1 \\ v - (K_{cores} - 1)[\frac{v}{K_{cores}}] & \text{otherwise} \end{cases} \quad (11)$$

Where $[\cdot]$ is the mathematical operation of rounding a rational number to the nearest integer. In the reference figure 3, tensor size is $M = 4, N = 3, D = 5$, thus $v = 60$. For a $K_{cores} = 12$ cores CPU, we split data into 12 blocks, namely $B_1, B_2, ...B_{12}$. In this case, all blocks would have the same amount of voxels, that is, 5 voxels. In another example, a tensor with size $M = 5, N = 5, D = 5$ would be split into 11 blocks of 6 voxels $(B_1, ..., B_{11})$ and a final block of $B_{12} = 9$ voxels.

To execute each data block in parallel, we use the *Parallel computing toolbox*. Matlab spawns K_{cores} child processes,each one receiving a corresponding B_i data block. During execution, Matlab handles all task scheduling automatically. Finally, we reassemble the results with the *reshape* utility of Matlab.

4. Experimental Setup

Our experiments focus on two different tests for assessing the stability of medial branching and applicability to medical decision support systems. The first test evaluates the stability of the medial surface branches for known volumes undergoing a controlled deformation. The second test explores the capability of medial branches for describing anatomical deformations and the computational efficiency of the medial algorithm.

Our validation protocol for medial surface stability has been applied to the method described in Section 3 computed for $\mathcal{M} = GSM2$. In order to assess the improvement in branch recovery, we have also considered $GSM2$ alone. Since a main property of medial surfaces to ensure a faithful representation using medial coordinates is preservation of original volume homotopy, for both cases, the largest connected component was selected. The parameters used to compute $GSM2$ were set taking into account that sigma should achieve a compromise between smoothing for accurate computations and accurate spatial localization of medial structures. Since in our case, filters are applied to a distance map, to ensure the highest accuracy in medial surface localization sigma is taken (Vera et al. (2013)) to the minimum possible scale, $\sigma_x = \sigma_y = 1$, $\sigma_z = 2$.

To compare to other methods, we have also tested two thinning methods with pruning and two Voronoi approaches. The thinning methods are a pruned version of the 26-connected neighborhood method (labelled ThP_{26}) (Pudney (1998)) and the 6-connected scheme (labelled Tao_6) (Ju et al. (2007)) that alternates thinning with pruning stages. The Voronoi approaches are the public domain CM-Rep with pruning (labelled $CMRep$) available at [1] and the Scale Axis Transform[2] (labelled SAT) (Giesen et al. (2009)).

Our test machine featured an Intel Core i7 3970K hexa-core Processor working at 4.2Ghz with 24 GB of RAM. For the GPU experiments, we used a Geforce GTX-550 Ti video card with 1GB VRAM featuring 192 CUDA cores. Blowups were computed using the parallelization described in Section 3.3, while $GSM2$ used a serial implementation. Concerning morphological approaches we used a set of code implementations (Pudney (1998)) that are essentially serial by the nature of thinning approaches (see the Discussion section).

4.1. Medial Branch Stability

Stability of medial branches has been tested on our own synthetic database [3] for medial surface quality evaluation (Vera et al. (2013)). Our synthetic volumes cover different aspects of medial surface geometry (including different degrees of medial branching) and volume medial representation (uniform and varying radii for the inscribed spheres).

Our synthetic experiments have been designed to measure the compromise between detection stability (for a given deformation rate) and spurious spikes arising from smaller perturbations. To such end, the original volumes have been deformed in order to generate branches at specific sites selected on a mesh of the volume boundary. Perturbations of synthetic volumes increase from none to a percentage of the original volume thickness given by the synthetic radial map used to create objects (D_{max}). The position of each selected point, $\overrightarrow{P}$, is modified by a translation δ_P along the boundary normal direction at that point, $\overrightarrow{N}_P$:

$$\overrightarrow{P} \rightarrow \overrightarrow{P} + \delta_P \overrightarrow{N}_P$$

for $\delta_P \in [0, D_{max}]$. Given that for $\delta_P = 0$, we have the original volumes, the connected components of the difference between volumes for $\delta_P = 0$ and $\delta_P > 0$ is the collection of volume spikes, namely $\{VS_i\}_{i=1}^{N_{VS}}$ generated by the deformation process.

Medial surfaces for $\delta_P = 0$ give the baseline accuracy by comparison to the database ground truth surfaces (Vera et al. (2013)). For $\delta_P > 0$, computed medial surfaces should generate new branches for each volume spike if the deformation size δ_P is large enough to introduce a significant change in volume curvature. Branches not arising from volumetric spikes changing boundary convexity profile

[1] cmrep.cvs.sourceforge.net/viewvc/cmrep/cmrep/src/VoronoiSkeletonTool.cxx

[2] http://code.google.com/p/mesecina/

[3] available at http://iam.cvc.uab.es/downloads/medial-surfaces-database

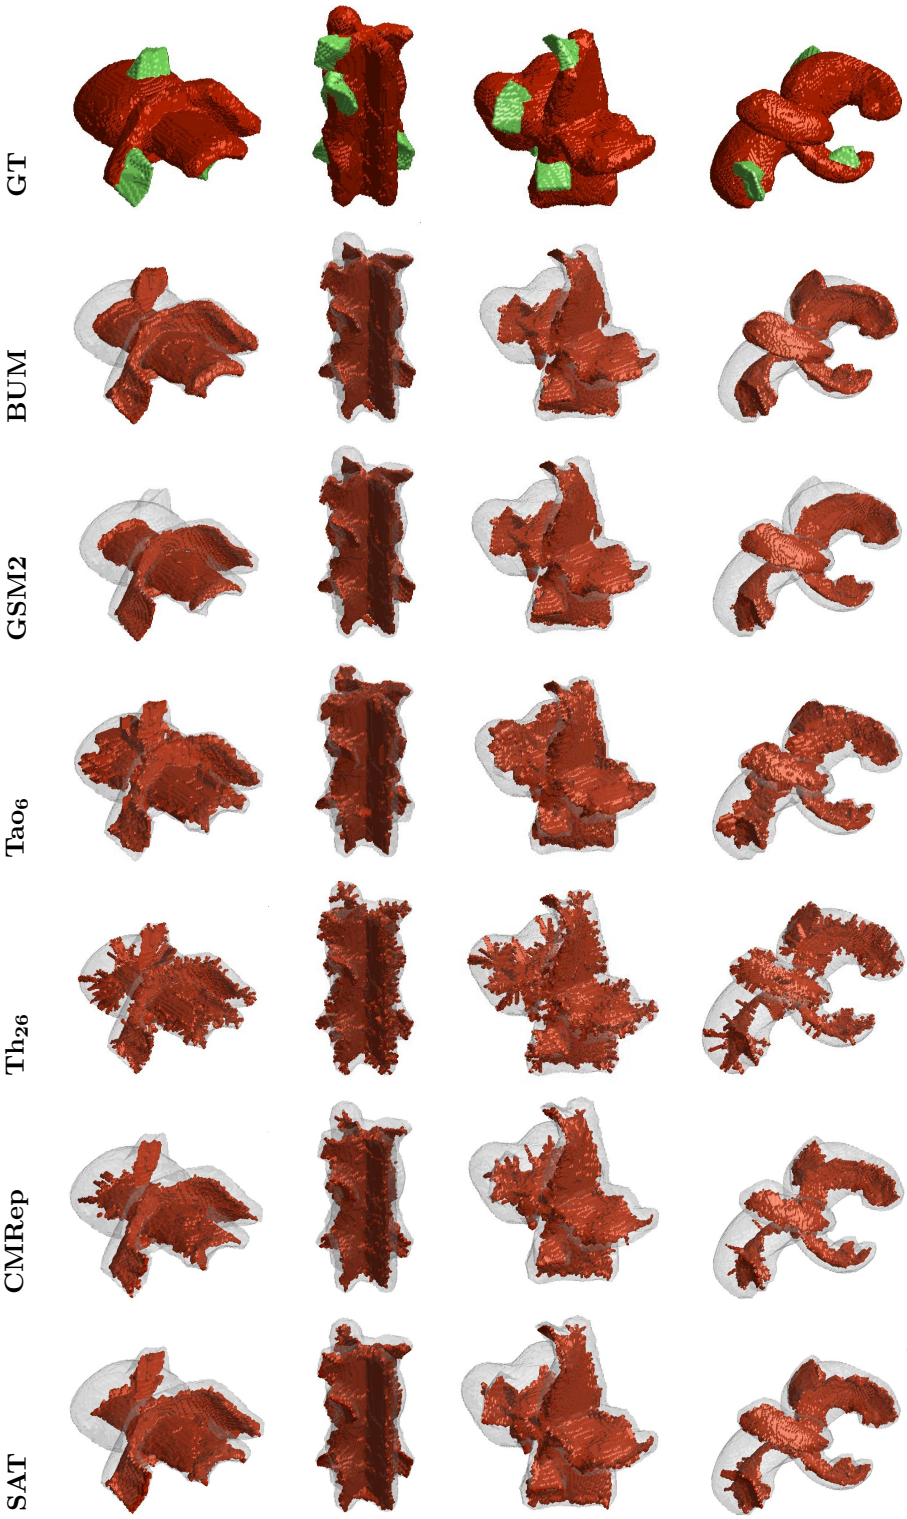

Figure 4: Assessment of medial stability for the different methods compared to the ground truth (GT).

are useless and should be as least as possible. The quality of medial branching arising from volumetric spikes has been assessed in terms of spike detection and its accurate localization. Branches arising from the volume deformation are given by the connected components of the difference between medial surfaces for $\delta_P = 0$ and $\delta_P > 0$. We will note them by $\{B_j\}_{j=1}^{N_B}$.

Spike detection rate has been measured in terms of medial branch false and true positives. A branch is considered a true positive if it intersects any of the volume spikes VS_i. In order to measure the impact of false branches arising during volume deformation (i.e. detection capability), we have also considered the percentage of area that true positives and all medial branches represent over ground truth medial surfaces:

1. *True Branches*:

$$TB = \frac{\#\{VS_i \text{ s. t. } \exists B_j,\ B_j \bigcap VS_i \neq \emptyset\}}{N_{VS}}$$

   the detection is optimal if $TB = 1$.

2. *Detection Capability*: We measure detection capability by considering the percentage of medial branches area, MBA, as well as, the percentage of true branches area, TBA. If X denotes the ground truth medial surface and B_j^{TB} denotes a true medial branch, then such area percentages are given by:

$$TBA = 100\frac{\sum \|B_j^{TB}\|}{\|X\|} \qquad MBA = 100\frac{\sum_{j=1}^{N_B} \|B_j\|}{\|X\|}$$

   for $\|\cdot\|$ denoting the area of a surface.

The score MBA measures the total amount of spikes and, in particular, for $\delta_P = 0$ the ones that are not related to any volume deformation. The score TBA indicates the percentage of surface branches that really correspond to volume deformation. The ideal quality plots for TB should present an asymptotic profile converging to 100%. Concerning TBA and MBA plots, they should be equal and have an increasing pattern. Any divergence between TBA and MBA ranges is due to the presence of spikes not linked to the introduced perturbation.

We define spike localization in terms of the distance to the volume spikes, namely d_{VS}, and the ground truth medial surfaces, namely d_X. For each point in computed medial surfaces $y \in Y$, we have that the minimum between $d_X(y)$ and $d_{VS}(y)$ reflects a compromise between between medial branches size and its proximity to a volume spike. Let $DL(y) := min(d_X(y), d_{VS}(y))$ denote such minimum. Then, our localization scores are given by the average (ADL) and maximum (MDL) values of DL over the computed surface:

1. *Spike Localization*:

$$ADL = \frac{1}{\#Y}\sum_{y\in Y} DL(y) \qquad MDL = \max_{y\in Y} DL(y)$$

4.2. Clinical Applicability

In medical imaging applications the aim is to generate the simplest medial surface that allows recovering the original volume without losing significant voxels. Besides, for its application to diagnosis and treatment planning, algorithms should reach a compromise between accuracy and computational cost.

Reconstruction capabilities have been assessed by comparing volumes recovered from surfaces generated with the different methods to ground truth volumes. Volumes are reconstructed by computing the medial representation (Blum (1967)) with radius given by the values of the distance map on the computed medial surfaces. Let A, B be, respectively, the original and reconstructed volumes and ∂A, ∂B, their boundary surface. Completeness of reconstructed volumes is assessed using the following volumetric and distance measures:

1. *Volume Overlap Error:*

$$VOE(A, B) = 100 \times \left(1 - 2\frac{\|A \cap B\|}{\|A\| + \|B\|}\right)$$

2. *Maximum Volume Boundary Difference:*

$$MVD = \max\left(\max_{x\in\partial A}(d_{\partial B}(x)), \max_{y\in\partial B}(d_{\partial A}(y))\right)$$

In order to provide a real scenario for the reconstruction tests we have used 15 livers from the SLIVER07 challenge (Heimann et al. (2009)) as a source of anatomical volumes. To compare to the Voronoi skeletonization methods SAT and $CMRep$, liver CT meshes were smoothed and decimated to 80% of the triangles using the VTK implementation of the decimation of triangle meshes taking into account the mesh error (Schroeder et al. (1992)).

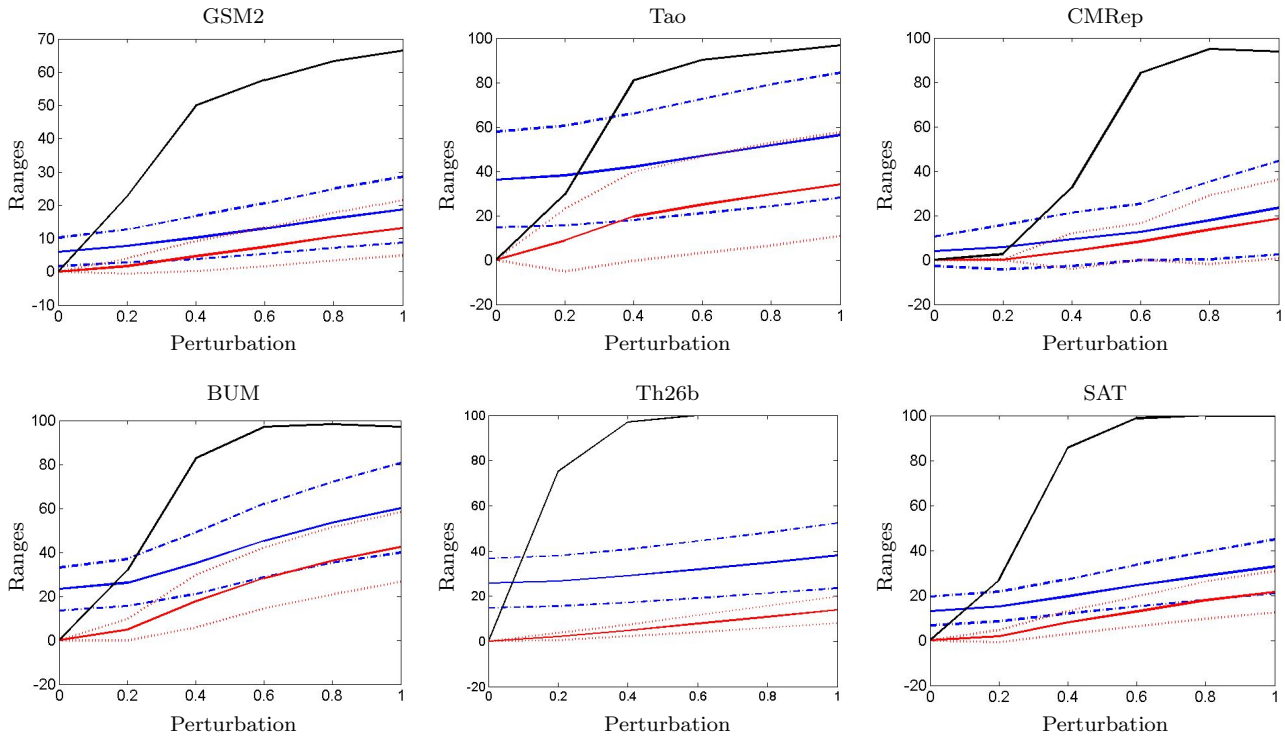

Figure 5: Medial stability scores and detection rates, energy-based methods on left images, morphological thinning on middle images and Voronoi-like on right images. Perturbation axis ranges from 0 (no perturbation) to 1, indicating maximum perturbation (D_{max})

	BUM	$GSM2$	Tao_6	ThP_{26}	$CMRep$	SAT
ADL	3.503 ± 1.1768	3.4335 ± 1.047	4.8687 ± 1.2491	5.2317 ± 1.0263	3.7637 ± 1.1504	3.918 ± 0.97782
MDL	7.8612 ± 0.95109	8.0367 ± 0.84547	10.5166 ± 2.9167	11.8142 ± 3.0043	9.1912 ± 2.7196	9.3808 ± 1.2645

Table 2: Localization ranges (mean and standard deviation) for the synthetic volumes

5. Validation Experiments

5.1. Medial Branch Stability

Figure 4 illustrates the computed medial surfaces of several volumes. In the 1st row, we show the deformed volume with its spikes in green and the volume for $\delta_P = 0$ in red. In the remaining rows, we show computed medial surfaces in solid meshes and the synthetic volume in semi-transparent color. The shape of surfaces produced using morphological thinning strongly depends on the connectivity rule used. In the absence of pruning, surfaces, in addition, have extra medial axes attached. Although they detect all volume spikes, Th_{26P} produces a bunch of medial axes instead of a clear surface for each spike. The impact of spikes is substantially reduced for Voronoi surfaces, especially for SAT, although medial surfaces still have some irregularities on their boundaries. Finally, ridge-based medial surfaces have a well defined shape matching the original synthetic surface without extra structures. The benefits of BUM for branch recovery are clearly illustrated by the 1st and 4th cases, which upper branch is missing for $GSM2$.

Figure 5 plots branch detection rates across volume perturbation. We show average curves computed across all medial surfaces for TP, MBA and TBA. In the case of MBA and TBA, their ranges given by ± standard deviation are also shown in colored bands. Both energy methods reach a similar compromise between extra structures and true branches. However the detection rate significantly drops without blowup (from 100% to less than 70%. The impact of extra structures is less than 30%.

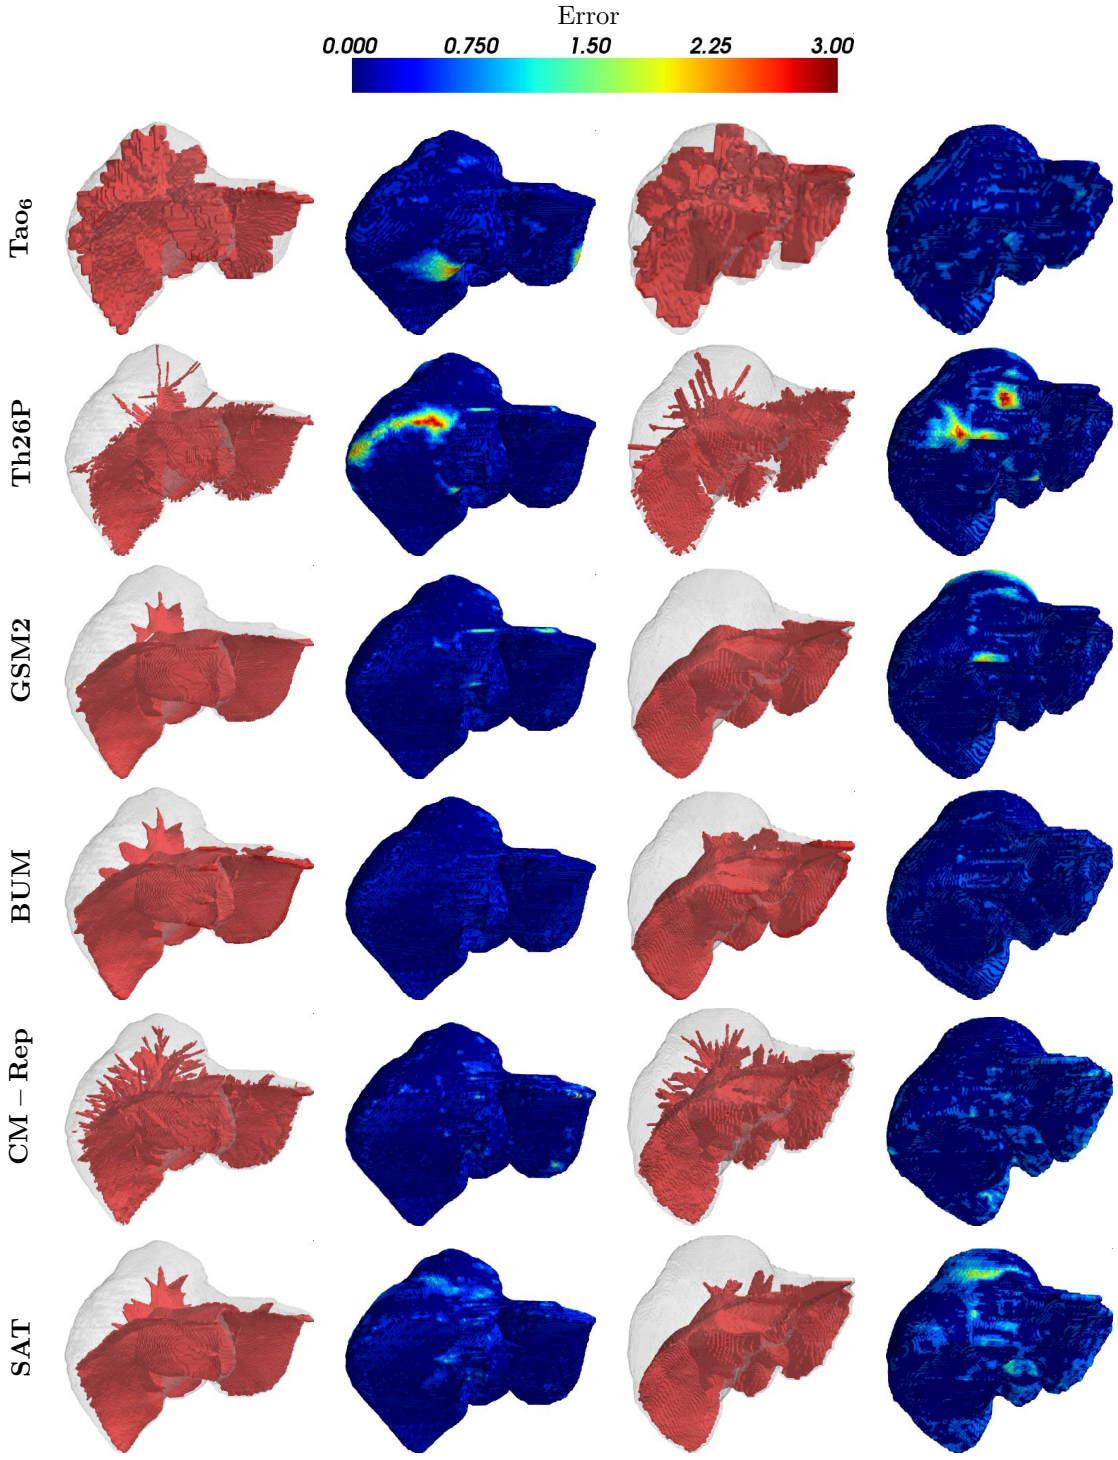

Figure 6: Reconstruction Power for Clinical Applications. Volumes reconstructed using computed medial surfaces and distance error (pixels) between reconstructed and original liver volumes. The size of the liver bounding box is shown in table 4 for reference.

It is worth noticing that such extra structures do not alter the geometry of medial surfaces(see fig.4). They rather correspond to an over detection at surface endings, which increases for blowup surfaces due to some artifacts in the localization of the blowup sites. The behavior of the morphological approaches is substantially different. The capability for proper detection of volume spikes is higher for Tao surfaces, thanks to a higher connectivity (see 2nd row in fig.4). However, it introduces a larger area of branches alien to volume deformation (about 60%) and also misses some important branches that should arise for large volume deformation (TP does not reach 100%). Although Th_{26P} has a 100% of detection rate, it presents the lowest ratio of true branches due to a fragmentation in single medial axis instead of a surface branch. Concerning Voronoi approaches, $CMRep$ is the one achieving a best detection capability as MBA and TBA ranges coincide at the cost of some missing branches. SAT performance is close to BUM in terms of branch detection capabilities and presence of extra structures.

Finally, Table 2 reports spike localization ranges computed over all shapes and deformation degrees for the four methods. There are not significant differences between the two ridge-based methods, which demonstrates that blowups improve branch without introducing other geometric artifacts. Comparing to morphological methods, we observe a significant increase in localization errors due to extra branches. Finally, both Voronoi methods have a localization capability between morphological thinnings and ridge-based surfaces.

Figure 7: SAT generated artifacts in liver natural folds. Left: medial surface (red) lying outside of liver boundary (white) near the falciform ligament. Right: sliced volume with artifacts shown in blue.

5.2. Clinical Applicability

Figure 6 illustrates assessment of reconstruction power for clinical applications for two different liver anatomical shapes. Volumes reconstructed using the computed medial surfaces (colored in red) are shown in transparent blue over true anatomical volumes shown in transparent gray. Differences between the reconstructions and the original volumes are better appreciated in volumes colored according to distances between reconstructed and original anatomical volumes.

Reconstruction accuracy of morphological methods varies upon the pruning scheme. The morphological Tao_6 achieves a good reconstruction power at the cost of a high density of branches in medial surfaces. Meanwhile, Th_{26} drops reconstruction power at some parts due to a higher branch sparsity (as illustrated by the first volume). CMReps shows multiple branching artifacts despite the pruning step, although the density of the branches produces acceptable reconstruction power. Increased pruning values also affects the reconstruction power as triangles on the edges of main surfaces are lost. SAT produces results close to BUM but with reduced reconstruction capabilities. A problem of SAT for medical applications is the introduction of medial structures that distort medial topology. This is due to the smoothing multiplicative factor applied to the object boundary distance map, which at boundary areas presenting narrow concavities can generate medial structures that partially lie outside of boundary. This artifact is exemplified in fig. 7, which shows a liver volume in gray and its SAT medial surface in red. The medial surface is visible partially outside the liver boundary near the falciform ligament concavity (surrounded by a blue ellipse). Such site is the division of the liver in two of the anatomical segments which size and shape is relevant for assessing living donor liver transplants. Similar artifacts also appear at the lower surface of the liver (also highlighted with a blue ellipse), where there is a number of similar concavities in the liver boundary. We note that in the absence of branch blowup, the energy-based $GSM2$ has a poorer reconstruction power. This drop especially increases in the presence of prominent convexities, like the one shown in the right liver.

The benefit of branch blowup is better illustrated in the images of fig.8 that present the common pitfall of ridge-based branch missing and its negative effects in restoration of volume finest details. As before, we show medial surfaces as well as volumes

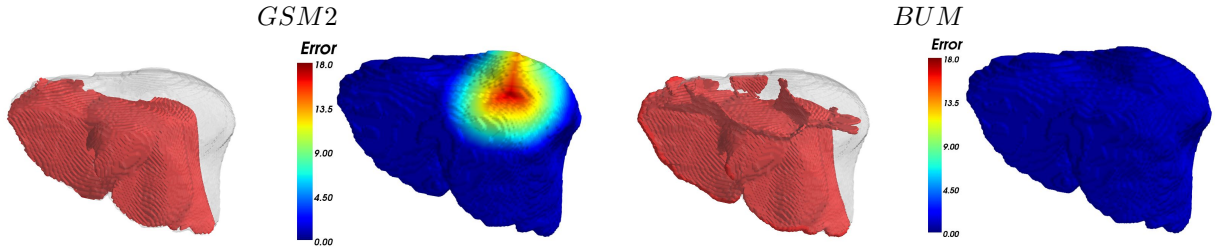

Figure 8: Benefits of BUM for restoring medial junctions. Medial surfaces and distance errors for $GSM2$, leftmost images, and BUM, rightmost images.

colored according to distance errors for $GSM2$ (left volumes) and BUM (right volumes). Small medial branches at the top of the liver are completely missing for $GSM2$ surface. This introduces a large error (up to 18 pixels) in $GSM2$ reconstruction at such area. Meanwhile, BUM surfaces present a continuous connected junction profile everywhere and, in particular, top branches are preserved. Such branches contribute to the reconstruction of the top finest anatomical details, as shown in the colored volumes. Figure 9 better shows that BUM junctions are completely connected to the main medial surface and not just linked by few voxels.

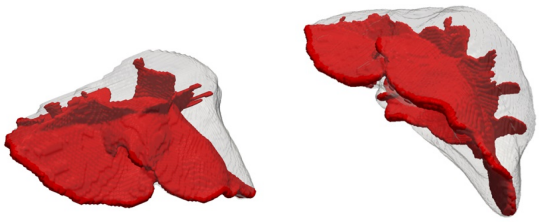

Figure 9: BUM connectivity at medial junctions.

Table 3 reports the statistical ranges for all methods and measures computed for the 15 livers. Gross differences between volumes are detected by VOE and, in spite of errors at prominent lobes, none of the cases seem to be significantly better. In medical applications, restoring local deformations can be important for early diagnosis. In this context, the surface distance score MVD is suitable for detection of local differences. Our approach BUM results is the best performer and the thinning Tao_6 the worst one. Although a bit better than morphological thinning, both Voronoi methods are behind energy-based methods and present the largest variability across liver volumes.

Computational times required for each method are given in Table 4. We report execution times for each algorithm, as well as, the average (μ) and standard deviations (σ) for the whole data-set in last rows. For Voronoi methods working with surface meshes, times include the conversion from volume voxels to mesh surfaces and we also report the number of mesh triangles. Although all algorithms scale linearly in relation to the size of the input volume, their computational execution time ranges are vastly different. In fact, both Tao_6 and $Th26P$ are slower by order of thousands of seconds in comparison to $GSM2$ and BUM. This is the expected behavior, as $Th26P$ and Tao_6 depend on the volume thickness, unlike filter-bank approaches ($GSM2$, BUM). Moreover, in the case of Tao_6, the algorithm must go through a set of voxel-suppression tests to ensure to enforce valid topologies, which further decreases overall performance. Finally, both $Th26P$ and Tao_6 depend on the surface angular topology, which means that for complex surfaces algorithm speed will be affected. Voronoi-based methods are faster than methods based on morphological thinning but slower than energy-based methods. It is worth mentioning that SAT is significantly better than CMRep and achieves a competitive computational time comparable to the one achieved by energy-based methods. These last methods produce medial surfaces of full resolution volumes in the fastest time and achieving the best reconstruction power in the case of BUM.

6. Final Remarks

A main limitation for the use of medial surfaces in applications oriented to clinical diagnosis is the presence of spurious branches or unwanted media surface manifolds which do not actually convey any anatomic description. On one hand, a complex medial medial geometry complicates modelling organ

	BUM	$GSM2$	Tao_6	ThP_{26}	$CMRep$	SAT
VOE	2.3 ± 0.5	2.2 ± 0.5	2.1 ± 0.4	4.4 ± 0.9	4.4 ± 9.8	3.7 ± 6.9
MVD	6.5 ± 4.3	8.3 ± 3.9	8.9 ± 3.5	12.3 ± 5.6	8.5 ± 13.2	9.3 ± 14.0

Table 3: Mean and standard deviation of errors in volume reconstruction for each metric.

Volume Id	Volume Size	Triangles	Tao_6	$Th26P$	$CMRep$	SAT	$GSM2$	BUM
Liver01	208 x 210 x 163	14038	120960	10944	2985	233	56	233
Liver02	139 x 172 x 138	15038	35780	3654	2994	204	28	112
Liver03	139 x 149 x 132	11770	28050	2860	2686	211	21	107
Liver04	244 x 244 x 179	27172	183630	16200	7577	282	79	193
Liver05	264 x 197 x 239	30056	259670	17345	8909	283	95	211
Liver06	204 x 131 x 144	11802	37320	3718	2697	196	33	139
Liver07	255 x 222 x 203	34452	214970	16445	9968	268	101	202
Liver08	261 x 249 x 242	43868	349850	23369	14599	336	139	246
Liver09	258 x 213 x 213	29844	222660	14388	8674	269	87	204
Liver10	212 x 227 x 245	32428	262570	17826	9456	280	104	205
Liver11	215 x 227 x 212	25896	241760	16153	6958	262	84	189
Liver12	145 x 164 x 175	13368	42300	4289	2832	209	41	152
Liver13	208 x 232 x 250	33488	236190	15114	9269	269	89	124
Liver14	275 x 364 x 246	45748	368720	26547	15458	382	195	340
Liver15	162 x 185 x 122	13950	46200	4263	2856	201	27	119
μ			176710	12874	7194	259	79	185
$\pm\sigma$			117130	7579	4303	52	48	62

Table 4: Computational Cost: size (voxels in volume of interest, number of triangles for mesh based methods), time (seconds), μ stands for mean and $\pm\sigma$ for standard deviation

populations for statistical analysis or correspondence mapping. On the other hand, hinders easy definition of tubular coordinates providing correspondence across different volumes, which constitutes a main advantage of medial representations over other volumetric models. Any method that reaches a compromise between the number of medial structures, their stability against noise in the boundary and reconstruction capability would constitute an excellent basis for using medial representations in the medical imaging field.

We describe a new method easy to reproduce based on blowup of singularity points in volumetric data that shows a good balance between the above characteristics. Aiming at its use in medical applications, we provide a hybrid CPU-GPU parallel implementation for standard PC's. In addition, we present a comparison of the stability of several medial surface methods, including morphologic approaches and Voronoi mesh methods, against the presence of spurious branches and their capability for reconstructing the original volume.

Our experiments show the benefits of using blowup resolution of medial branches in ridge-based approaches. Morphological methods present a complex medial branching with high instability. Voronoi methods achieve competitive results but require several pre-processing steps prone to introduce errors when dealing with volumetric medical data. In addition, some of them are prone to distort medial topology. Our approach achieves a good balance between branching stability, simplicity and capability for reconstructing volume boundary finest details. Such details might be good indicators of an illness first stages and, thus, a medial map able to capture such variations in shape without introducing unnecessary complexity will be a solid basis for early diagnosis and treatment progress assessment. In this context, we have used our medial representation for defining parametric maps of anatomical volumes assigning equal coordinates to equivalent anatomical feature points (Vera et al. (2014)).

An important concern for the use of methodologies in real clinical settings is their computational

cost. Morphological methodologies are based on iterative removal of simple voxels after checking that their removal does not alter volume topology. Since the removal of a voxel from the object, changes the local topology on its neighborhood, voxels can only be removed one by one, which limits the parallelization of the method without deep changes in the algorithm. Concerning mesh methods, CMRep performance depends on the number of mesh triangles and, thus, the decimation factor applied to the mesh is a key parameter for achieving a good compromise between time and accuracy. Regarding SAT, it scales well with triangle size, with its performance depends on the parameters used in the computation. Given that such parameters relate to the degree of simplification, computational time might increase if the application requires high order of detail.

The presented hybrid approach featuring GPU convolution has similar performance in comparison to a pure CPU multicore implementation. However, this has a cost in the form of considerable high memory usage. For instance, the multicore GPU approach peaks at 5GB RAM on average, whereas the pure multicore method requires up to 14 GB RAM to run. Moreover,we face restrictions in the form of limited device memory, together with data transfers overhead. For this reason, the GPU implementation is not as optimal as it could be. Finally, not all volumes fit on device memory, so this approach is not always available. Still, we feel that the hybrid approach is a good compromise between speed and memory requirements and its improvement is a topic in our future research.

7. Acknowledgements

Work supported by Spanish Projects TIN2012-33116, DPI2015-65286-R, TIN2013-47913-C3-1-R and Generalitat de Catalunya, 2014-SGR-1470.

References

Amenta, N., Choi, S., Kolluri, R., 2001. The power crust, unions of balls, and the medial axis transform. Computational Geometry: Theory and Applications 19(2-3), 127–153.

Blum, H., 1967. A transformation for extracting descriptors of shape. MIT Press.

Bodnár, G., Schicho, J., 2000. Two computational techniques for singularity resolution. Journal of Symbolic Computation 32, 39–54.

Bouix, S., Siddiqi, K., Tannenbaum, A., 2005. Flux driven automatic centerline extraction. Med. Imag. Ana. 9, 209–21.

Chazal, F., Lieutier, A., 2005. The λ-medial axi. Graphical Models 67(4), 304–331.

Csernansky, J., Wang, L., Jones, D., Rastogi-Cru, D., Posener, J., 2002. Hippocampal deformities in schizophrenia characterized by high dimensionalbrain mapping. Am. J. Psychiatry 159, 1–7.

Dey, T.K., Zhao, W., 2002. Approximate medial axis as a voronoi subcomplex, in: Proceedings of the seventh ACM symposium on Solid modeling and applications, ACM. pp. 356–366.

Dinguraru, L., Pura, J., et al., 2010. Multi-organ segmentation from multi-phase abdominal CT via 4D graphsusing enhancement, shape and location optimization, in: MICCAI, pp. 89–96.

Freeman, W., Adelson, E., 1991. The design and use of steerable filters. IEEE Trans. Pattern Analysis and Machine Intelligence 13, 891–906.

Garcia, J., Gil, D., Badiella, L., et al., 2010. A normalized framework for the design of feature spaces assessing the left ventricular function. IEEE Trans Med Imag 29(3), 733–745.

Giblin, P., Kimia, B., Pollitt, A., 2009. Transitions of the 3d medial axis under a one-parameter family of deformations. PAMI 31(5), 900–918.

Giesen, J., Miklos, B., Pauly, M., Wormser, C., 2009. The scale axis transform, in: SCG, pp. 106–115.

Gray, A., 2004. Tubes. Birkhäuser.

Heimann, T., van Ginneken, B., Styner, M.A., Arzhaeva, Y., Aurich, V., 2009. Comparison and evaluation of methods for liver segmentation from CT datasets. IEEE Trans. Med. Imag. 28(8), 1251–1265.

Hicklin, J., Moler, P., et al., . Jama: A java matrix package. URL: `http://math.nist.gov/javanumerics/jama`.

Hironaka, H., 1964. Resolution of singularities of an algebraic variety over a field of characteristiczero. Ann. Math. 79, 109–326.

Ju, T., Baker, M.L., Chiu, W., 2007. Computing a family of skeletons of volumetric models for shape description, in: In Geometric Modeling and Processing, pp. 235–247.

Khalifa, F., El-Baz, A., Gimel'farb, G., Ouseph, R., 2010. Shape-appearance guided level-set deformable model for image segmentation, in: ICPR.

Kollár, J., 2007. Lectures on Resolution of Singularities (AM-166). Princeton University Press. URL: `http://www.jstor.org/stable/j.ctt7rptq`.

Liu, X., Linguraru, M., Yao, J., Summers, R., 2010. Organ pose distribution model and an MAP framework for automated abdominalmulti-organ localization, Springer. pp. 393–402.

Lopez, A., Lumbreras, F., Serrat, J., Villanueva, J., 1999. Evaluation of methods for ridge and valley detection. IEEE Trans. Pat. Ana. Mach. Intel. 21, 327–335.

N.Faraj, Thiery, J.M., Boubekeur, T., 2013. Progressive medial axis filtration, in: SIGGRAPH.

Ogniewicz, R., Ilg, M., 1992. Voronoi skeletons: Theory and applications, in: CVPR, pp. 63–69.

Park, H., Bland, P., Meyer, C., 2003. Construction of an abdominal probabilistic atlas and its application insegmentation. IEEE Trans Med Imaging 22(4), 483–92.

Perona, P., 1992. Steerable-scalable kernels for edge detection and junction analysis, in: Image and Vision Computing, pp. 3–18.

Pizer, S., Fletcher, P., et al, 2005. A method and software for segmentation of anatomic object ensembles by deformableM-Reps. Medical Physics 32, 1335–1345.

Pryor, G., Lucey, B., Maddipatla, S., McClanahan, C., Melonakos, J., Venugopalakrishnan, V., Patel, K., Yalamanchili, P., Malcolm, J., 2011. High-level GPU computing with jacket for MATLAB and C/C++ , 806005–806005–6.

Pudney, C., 1998. Distance-ordered homotopic thinning: A skeletonization algorithm for 3Ddigital images. Comp. Vis. Imag. Underst. 72(2), 404–13.

Rohilla, A., Kumar, P., Rohilla, S., Kushnoor, A., 2012. Cardiac hypertrophy: A review on pathogenesis and treatment. Int. J. Pharma. Sci. Drug Res. 4(3), 164–167.

Satish, N., Harris, M., Garland, M., 2009. Designing efficient sorting algorithms for manycore GPUs, in: Intern. Symp. on Parallel and Distributed Processing, pp. 1–10.

Schroeder, W.J., Zarge, J.A., Lorensen, W.E., 1992. Decimation of triangle meshes, in: Computer Graphics (SIGGRAPH '92 Proceedings, pp. 65–70.

Shi, J., Tomasi, C., 1994. Good features to track, in: CVPR '94, pp. 593–600.

Siddiqi, K., Bouix, S., Tannenbaum, A., Zucker, S., 2002. Hamilton-Jacobi skeletons. Int. J. Comp. Vis. 48, 215–231.

Stough, J., Broadhurst, R., Pizer, S., Chaney, E., 2007. Regional appearance in deformable model segmentation 4584, 532–543.

Styner, M., Lieberman, J.A., D., P., Gerig, G., 2004. Boundary and medial shape analysis of the hippocampus in schizophrenia. MediMA 8, 197–203.

Sun, H., et al., A.F., 2010. Automatic cardiac mri segmentation using a biventricular deformable medialmodel, in: MICCAI, Springer. pp. 468–475.

Sun, H., Avants, B., Frangi, A., Ordas, S., Gee, J., Yushkevich, P., 2008. Branching medial models for cardiac shape representation, in: ISBI, pp. 766–773.

Svensson, S., Nyström, I., di Baja, G.S., 2002. Curve skeletonization of surface-like objects in 3d images guided by voxelclassification. Pattern Recognition Letters 23(12), 1419–1426.

Terriberry, T., Gerig, G., 2006. A continuous 3-d medial shape model with branching, in: MICCAI Math Found Comp Ana.

Vera, S., Gil, D., Borràs, A., Linguraru, M.G., González Ballester, M.A., 2013. Geometric steerable medial maps. Machine Vision and Applications 24(6), 1255–1266.

Vera, S., González Ballester, M.A., Gil, D., 2012a. A medial map capturing the essential geometry of organs, in: International Symposium on Biomedical Imaging, IEEE. pp. 1691–1694.

Vera, S., González Ballester, M.A., Gil, D., 2014. Anatomical parameterization for volumetric meshing of the liver, in: Proc. SPIE.

Vera, S., González Ballester, M.A., Linguraru, M.G., Gil, D., 2012b. Optimal medial surface generation for anatomical volume representations, in: MICCAI Workshop on Abdominal Imaging. Computational and Clinical Applications, pp. 265–273.

Wang, Y., Gu, X., Hayashi, K., Chan, T., Thompson, P.M.and Yau, S., 2005. Brain surface parameterization using riemann surface structure, in: MICCAI05, pp. 657–665.

Yao, J., Summers, R., 2009. Statistical location model for abdominal organ localization, in: MICCAI, Springer. pp. 9–17.

Yushkevich, P., 2009. Continuous medial representation of brain structures using the biharmonic PDE. NeuroImage 45, 99–110.