

SCHOOL OF COMPUTER SCIENCE
PHD STUDENT PROGRESS REPORT FORM
by **Emmanuel Piuze**, PhD candidate at McGill University
January 2011

1 Introduction

Cerebral blood vessels can help differentiate normal from severely pathological anatomy. The identification of vessel lesions and malformations in medical images strongly correlates with risk factors for arterial embolisms (vascular blockage), ischemic strokes, and aneurysms, all of which are relatively common conditions that can be debilitating and fatal [17]. Complications in the study of medical images arise from the fact that they typically are large and dense volumes, part of vast data sets. They are often cluttered with a variety of structures and tissue types (for example blood vessels, skin, bone, cerebrospinal fluid, background, white and grey matter tissue) that can vary greatly from one patient to another. Moreover, blood vessels can be dimmed or rendered invisible as the result of a limited acquisition resolution, sampling artifacts, spatial aliasing, and noise, and vasculature networks can thereby be broken. These difficulties make the identification and study of blood vessels laborious, and suggest the involvement of image processing techniques to assist clinicians in the reconstruction, analysis, and interpretation of cerebrovascular images. [17]

In computer vision, *image segmentation* is defined as the process of partitioning an image into regions of similarity. Segmentation methods can support the clinician in extracting neurovasculature from medical images. These vascular representations are helpful in the development of safer treatments for cerebral conditions, for a more efficient preoperative planning, operating room (OR) and postoperative monitoring, and for intra- and inter-patient comparative studies [8, 10]. Common features of interest include vessel branch points, quantification of the spatial relationship between vessels, vessel segment extraction, vessel diameter measurements, and vessel complexity assessment. When a sequence of images is available, it can also be useful to quantify changes in these quantities over time or across patients [23], offering a lot of research potential.

The main focus of this literature review is to present a selection of key articles covering different approaches employed in the vessel segmentation literature. Basic principles of neuroanatomy and neurovascular imaging are first introduced in Section 2. The selected articles are then discussed in Section 3 and divided into four subsections: active contours in Section 3.1, curvilinear filters in Section 3.2, model-based approaches in Section 3.3, and histogram-based approaches in Section 3.4. Concluding remarks are found in Section 4.

2 Neuroanatomy and Vessel Imaging

The human brain has a dense and complex structure. Since this literature review is highly constrained in length and only concerns neurovasculature segmentation, many anatomical parts and details are omitted in order to favour the major vascular networks and related structures. These networks are largely studied and often used clinically to differentiate normal from pathological anatomy, making them interesting for applying segmentation algorithms. Neuroanatomical and neuroimaging information presented in this

section is taken from Morris' book on Practical neuroangiography [15] and Osborn's book on Diagnostic cerebral angiography [17]. Additional neuroimaging content is taken from Reimer and Landwehr [18].

2.1 Neuroanatomy

The brain is surrounded and protected by many layers. Starting from the cortex and going outwards, they are the pia mater, the arachnoid, the dura mater, the skull, the periosteum and finally the skin. Two of these layers are highly vascularized. The pia mater allows blood vessels to pass through and nourish the brain whereas the dura mater is a sac that among other functions, carries blood back from the brain towards the heart.

The human brain can be divided into three major parts: the forebrain or *prosencephalon*, the cerebellum or *mesencephalon*, and the hindbrain or *rhombencephalon*, as seen in Fig. 1. This division follows a neuroanatomist' classification of brain structures, from the most evolved to the most primitive.




Figure 1: The human brain divided into three main parts. (McGill's *Brain From Top To Bottom*)

The brain is a highly vascularized organ. The previous section indicated that its superficial layers carry blood to and out of the brain. However, a more complex neurovasculature network lies within the brain itself. Three main arteries vascularize the brain with freshly oxygenated blood coming from the heart: the anterior, the middle, and the posterior cerebral arteries. They connect with each other at the base of the brain in the circle of Willis, a well-known anatomical structure. The anterior communicating artery links the two anterior cerebral arteries, and the two posterior communicating arteries link the carotid artery to the posterior cerebral artery.

The drainage of deoxygenated blood from the brain is done through a complex venous network. There is no equivalent to the well-structured circle of Willis for this venous network. Instead, at the base of the brain, all venous systems drain towards a large venous sinus and finally to the jugular vein, where deoxygenated blood is carried back to the heart.

2.2 Brain Vessel Imaging

Invasive angiography techniques that use catheters and radioactive contrast agents and dyes has long been the gold standard. However, many of these methods have been shown to pose a health risk. For instance, originally x-ray neuroimaging was performed along with the injection of a contrast agent, Thorotrast, a highly carcinogenic radioactive compound that increased risk factors for cancer by 100 times.

The development of non-invasive vascular imaging techniques provided the medical community with powerful tools that posed a significantly lower threat to a patient's life, which in turn directly stimulated clinical research. In particular, the development of the three main noninvasive 3D neurovascular imaging

modalities: magnetic resonance imaging (MR), X-ray Computed tomography (CT) and ultrasound (US). These imaging modalities also come in different flavours but they are not described here for sake of length.

MRI Magnetic resonance imaging produces cross-sectional images (slices) of the body using magnetization and radio waves. When used for imaging vessels, the term MRA for magnetic resonance angiography is instead used. The physical principle behind the MRI is that charged particles emit radiation as their magnetization is changed. Our body is largely composed of water molecules – each containing two protons – distributed in different structures and in various densities. If there was therefore a way to magnetize the protons contained in water molecules and listen for their radiation signal, structural information could be revealed. In MRI, the brain is exposed to a strong magnetic field that aligns the magnetization of the water protons. By keeping this magnetic field constant, an electromagnetic field is turned on and then off, causing the protons to align in a new direction relative to the field and then to relax back to their initial magnetization. As they do so, they emit signals that can be detected. The electromagnetic field can be tuned to different frequencies and different emitting signals can be listened to in order to give contrast to different structures. By varying the direction of the electromagnetic field, other slices of the brain can be obtained. The resulting 3D brain volume is obtained as an assembly (sometimes requiring interpolation) of all the cross-sectional slices. MRI is arguably the most useful – it can achieve sub-millimetre accuracy – and safest noninvasive modern neurovascular imaging modality.

CT Computed tomography imaging produces cross-sectional images of the body using x-ray radiation. An x-ray beam can penetrate solid objects and the denser a material is, the more radiation will be blocked. The remaining rays that pass through the object can be captured and from them an image can be acquired. Structure in this image directly correlates with regions of different densities. For instance, bone has a greater density than most anatomical structures and thus has a high contrast in x-ray images. A CT brain scan works by shooting these x-ray beams through multiple planes surrounding the brain, which are then combined to form a 3D volume. In CT angiography (CTA), a dye is injected that binds to specific structures in the body – blood vessel, for instance – such as to manipulate the signal from the x-ray beams. The resulting images generally have a higher resolution than in MRA. However, the combination of the dye and of the x-ray imaging exposes the patient to a strong dose of radiation and potentially carcinogenic agents. The use of a CTA in neurovasculature imaging is therefore only prescribed when absolutely necessary.

US Ultrasound imaging produces produces 2D images of body structures using sound waves. High-frequency sound waves are aimed at a structure of interest and the resulting echo is recorded. The echo time for the propagating front wave gives estimates about the density of the underlying tissue, which can be reconstructed as a 2D image. US is impractical for neuroimaging since the thick layer of the skull and the numerous structures of various density and geometry that compose the brain can produce a noisy signal that is hard to interpret. As a result, the imaging must be done directly on the skin and closely follows its topology near the structure of interest, such that US devices are typically hand-manipulated by a human. Nevertheless, US is a fast and extremely safe imaging modality, and developments in US image reconstruction makes it useful for quickly assessing certain arterial pathologies, including atherosclerosis. Moreover, 3D images can be obtained by reconstructing multiple views from 2D US images.

3 Segmentation Schemes

3.1 Active Contour Segmentation

Active contours evolve a deformable model over an image according to internal – given by the contour itself – and external forces – given by the image – in order to extract object boundaries. This section presents two main subclasses of active contours: parametric and implicit active contours. Parametric active contours or *snakes* use an explicitly evolved and possibly manually assisted deformable contour. In contrast, implicit active contours define a contour as the zero level of a higher dimensional manifold. The evolution of the contour is thus carried on indirectly, through operations on the manifold itself.

3.1.1 Standard parametric snakes

Parametric snakes models are a class of methods that define a 2D contour $\mathbf{v}(s, t)$ that is time-varying over t and parameterized by $s \in [0, 1]$. The boundary condition $\mathbf{v}(0, t) = \mathbf{v}(1, t)$ is typically used to produce a closed snakes. An internal energy term $\mathcal{S}(\mathbf{v})$ is defined to characterize the deformation of the snakes as that of an elastic contour $\mathcal{S}(\mathbf{v}) = \int_0^1 w_1(s) \left| \frac{\partial v}{\partial s} \right| + w_2(s) \left| \frac{\partial^2 v}{\partial s^2} \right| ds$, where the physical weight function w_1 controls stretching and w_2 controls bending at the parametric coordinate s . An external energy term $\mathcal{P}(\mathbf{v})$ is constructed based on *a priori* knowledge of the image and of the application at hand. \mathcal{P} is defined by the scalar potential function $\mathcal{P} = \int_0^1 P(v(s, t)) ds$. The minima of the potential function corresponds with the image feature of interest. For a snakes that clings to edges, a potential of the form $P(x, y) = -\|\nabla [G_\sigma * I(x, y)]\|$ is used, where G_σ is a Gaussian smoothing filter of standard deviation σ . Combined together, the internal and external energy terms yield the energy functional

$$\mathcal{E}(\mathbf{v}) = \mathcal{S}(\mathbf{v}) + \mathcal{P}(\mathbf{v}) \tag{1}$$

$$= \int_0^1 \left(w_1(s) \left| \frac{\partial v}{\partial s} \right| + w_2(s) \left| \frac{\partial^2 v}{\partial s^2} \right| + P(v(s, t)) \right) ds. \tag{2}$$

Although the energy functional \mathcal{E} can be minimized statically, a dynamic system formulation is more flexible and allows a user to directly interact with the snakes as it evolves to equilibrium. The boundary of the snake is therefore evolved by the Euler-Lagrange equations of the energy functional:

$$\mu \frac{\partial^2 v}{\partial t^2} + \gamma \frac{\partial \mathbf{v}}{\partial t} - \frac{\partial}{\partial s} \left(w_1 \frac{\partial \mathbf{v}}{\partial s} \right) + \frac{\partial^2}{\partial s^2} \left(w_2 \frac{\partial^2 \mathbf{v}}{\partial s^2} \right) = -\nabla P(\mathbf{v}(s, t)). \tag{3}$$

which can be solved for numerically using for instance a semi-implicit Euler method.

3.1.2 Topology Adaptive Snakes (McInerney and Terzopoulos [14])

McInerney and Terzopoulos [14] present an explicit active contour algorithm that extends classical parametric snakes and can segment tubular structures in medical images. This algorithm has been subsequently cited and used extensively in the medical image processing literature for segmenting neurovasculature [10]. The topology adaptive snakes or *T-snakes* are developed to overcome limitations of the standard snakes: 1) their oversensitivity to contour initialization and 2) the topological inflexibility such that long vessels with branching, protrusions, merging and fragmentation are difficult to extract.

T-snakes are implemented as a discretized form of the standard snakes described in Section 3.1.1. At time t , the T-snake is represented as a connected set of N nodes $\mathbf{x}_i(t)$ with $\mathbf{x}_1(t) = \mathbf{x}_N(t)$. The parameterization is embedded in a regular spatial grid composed of triangular cells. The resulting space

is an affine cell image decomposition or *ACID*; interestingly, McInerney and Terzopoulos [14] call the resulting combination *snakes in ACID*. The size of the grid cells can be decreased in order to better approximate the contour of the object of interest. The T-snake is evolved according to a discrete and simplified version of Eq. 3, resulting in a set of first-order ordinary differential equations of motion

$$\gamma_i \dot{\mathbf{x}}_i + a\boldsymbol{\alpha}_i + b\boldsymbol{\beta}_i = \boldsymbol{\rho}_i + p\mathbf{f}_i \quad (4)$$

where $\dot{\mathbf{x}}_i$ is the velocity of node i and γ_i is a damping coefficient. The left hand side of Eq. 4 represent internal forces whereas the right hand side represents external forces. $a\boldsymbol{\alpha}_i$ and $b\boldsymbol{\beta}_i$ respectively control the resistance of the contour to stretching and to bending deformations. $\boldsymbol{\rho}_i$ is an inflation force that pushes the node towards the normal to the contour, by a magnitude proportional to the image intensity at \mathbf{x}_i . $p\mathbf{f}_i$ is a force proportional to the gradient of the scalar potential defined in Section 3.1.1 that stops the contour at significant edges. Eq. 4 is solved for using a simple forward Euler approach, $\mathbf{x}_i(t + \Delta t) = \mathbf{x}_i(t) + \frac{\Delta t}{\gamma} (a\boldsymbol{\alpha}_i(t) + b\boldsymbol{\beta}_i(t) - \boldsymbol{\rho}_i(t) - \mathbf{f}_i(t))$. Although this form is typically unstable, McInerney and Terzopoulos [14] mention that a large range of step sizes produce a stable behaviour.

As the T-snake is evolved, topological transformations need to be performed under certain conditions. This is the case when a T-snake collides with itself or with another T-snake, or when a T-snake breaks into multiple parts. The ACID framework allows to handle these transformations consistently. Whenever collisions or fractures of the T-snake are detected, a topological transformation is performed and the correspondence of the T-snake is re-established with the ACID grid. The equilibrium state for the evolution of a T-snake is reached when all of its nodes have been inactive for a given number of deformation steps.

3.1.3 Implicit Curve Evolution (Lorigo et al. [13])

The topological invariability of standard snakes described in Section 3.1.1 was addressed by McInerney and Terzopoulos [14] in the previous section. Although they work reasonably well in practice, T-snakes are based on heuristics and depend on the choice of the embedding space partition to work properly. Noise tends to generate spurious local minima that hinder the contour evolution. Geodesic active contours, on the other hand, offer a mathematically elegant paradigm that handles topological changes naturally and is independent of parameterization.

Geodesic active contours consider a curve as an interface between two media. When using a level-sets formulation, instead of evolving an n -dimensional curve $\mathbf{C}(q)$ to segment object boundaries, a manifold of codimension one ($(n + 1)$ -dimensional space) of which $\mathbf{C}(q)$ is the zero level-set is used. This higher dimensional manifold offers “room” for the contour deformation to undergo topological changes smoothly, without requiring a spatial reparameterization, as shown in Fig. 2.

Let $\mathbf{C}(q) : [0, 1] \rightarrow \mathbb{R}^2$ be a parameterized planar curve that is deformed to match an object boundary. An energy functional can be designed that controls the smoothness of the contour and makes it cling to the boundary:

$$\int_0^1 g(|\nabla I(\mathbf{C}(q))|) |C'(q)| dq \quad (5)$$

where g is a strictly decreasing function of the image gradient such that $g \rightarrow 0$ as $r \rightarrow 0$, e.g. $g = \frac{1}{1+|\nabla I|^2}$. Since \mathbf{C} is deformed over time t , the Euler-Lagrange equations can be computed and yield the curve evolution equation [2]

$$\mathbf{C}_t = g\kappa\mathbf{N} - (\nabla g \cdot \mathbf{N})\mathbf{N}, \quad (6)$$




Figure 2: Geodesic active contours naturally handle topological changes. a) shows two approaching and merging curves (in red) as the zero level set of a higher dimensional manifold. b)-d) show the evolution of multiple contours as merging zero level-sets. (Adapted from Deschamps [4])

where κ is the Euclidean curvature and \mathbf{N} is the unit inward normal to the curve. By embedding the curve as the zero level-set of a two-dimensional surface $u : \mathbb{R}^2 \rightarrow \mathbb{R}$, the signed distance function to \mathbf{C} , the following evolution equation is equivalent to Eq. 6:

$$u_t = g\kappa |\nabla u| + \nabla g \cdot \nabla u. \quad (7)$$

Given a cost function g , the surface u in Eq. 7 can be evolved to segment object boundaries. The resulting contour can be obtained by extracting the zero level-set of the final surface.

Lorigo et al. [13] use a codimension-two level-set scheme for segmenting blood vessels. That is, rather than evolving a planar curve in 2D (codimension-one), it is evolved in 3D. The geodesic active contour approach does not directly hold and a more general form of manifold evolution is therefore required. First, Eq. 6 is generalized to [13]

$$\mathbf{C}_t = \kappa \mathbf{N} - \Pi \mathbf{d} \quad (8)$$

where Π is the projection operator onto the normal space of \mathbf{C} , and \mathbf{d} is a cost vector field in \mathbb{R}^3 analogous to g in Eq. 6. Lorigo et al. [13] define an embedding space v that is the distance function to \mathbf{C} . They construct an isolevel set $\Gamma_\epsilon = \{x | v(\mathbf{x}) = \epsilon\}$ where ϵ is small and positive, that is Γ_ϵ is a thin tube around \mathbf{C} , made to represent a blood vessel. The larger principal curvature of this tube depends on ϵ while the smaller principal curvature depends on \mathbf{C} . The smaller principal curvature $\lambda(\nabla v(\mathbf{x}, t), \nabla^2 v(\mathbf{x}, t))$ is therefore used as a measure of manifold curvature. The evolution of the embedding space in Eq. 7 then becomes $v_t = \lambda(\nabla v, \nabla^2 v) + \nabla v \cdot \mathbf{d}$.

A cost vector field $\mathbf{d} = \mathbf{H} \frac{\nabla I}{|\nabla I|}$, where ∇I is the image gradient and \mathbf{H} is the Hessian, is defined. This vector field stops the contour evolution at specific curvature maxima. The equation of the embedding space then becomes

$$v_t = \lambda(\nabla v, \nabla^2 v) + \frac{g'}{g} \nabla v \cdot \mathbf{H} \frac{\nabla I}{|\nabla I|}. \quad (9)$$

The resulting contour is then obtained by extracting the zero level-set of the embedding space v_t when it reaches equilibrium, that is when the volume of the segmented region changes less than some specified percentage of total volume across a specified number of iterations. The radius of the vessel can be estimated as the inverse of principal curvature, i.e. $r = 1/\kappa$.

3.1.4 Shape Driven Flow (Nain et al. [16])

In practice, the active contours in Lorigo et al. [13] need to be initialized near the neurovasculature network. Moreover, the choice of ϵ in the definition of a tube as an isolevel set embedding determines the width of blood vessels, making the curve evolution largely dependent on the scale at which vessels appear. The region growing can fail when attempting to model vessels that exist at the inner of scale of the data and can be difficult to control.

Nain et al. [16] address some of these issues by proposing a level set technique that uses a soft shape prior rather than a rigid model of tubular structure. The level set formulation is similar to the one introduced in Eq. 8, where $\Pi = \phi$ is a cost or *speed* term based on the underlying image, alike g in Lorigo et al. [13], and $\mathbf{d} = \mathbf{N}$. The novelty resides in the introduction of *shape filters* for determining areas where the contour is widening and potentially leaking. The image gradient being highly sensible to noise, they suggest to use a local filter that is computed at a scale larger than that of the derivative. Filtering is carried out in a spherical domain, a ball $\mathcal{B}(\mathbf{x}, r)$ of radius r at each point \mathbf{x} of the contour \mathbf{C} , where r is an upper bound on the expected width of the vessel. The filter is defined as the percentage of the N points \mathbf{y}_i in $\mathcal{B}(\mathbf{x}, r)$ that fall inside the contour \mathcal{R} :

$$\epsilon_1(\mathbf{x}) = \frac{1}{N} \sum_{\mathbf{y}_i \in \mathcal{B}(\mathbf{x}, r)} \chi(\mathbf{y}_i) \quad (10)$$

where $\chi(\mathbf{y}) = \begin{cases} 1 & \text{for } \mathbf{y}_i \in \mathcal{R} \\ 0 & \text{for } \mathbf{y}_i \notin \mathcal{R} \end{cases}$. For a radius r is close to the width of the vessel, most neighbouring points

on the contour will have the same response to the ϵ_1 filter since locally the same percentage of neighbours fall within the filter radius. On the other hand, in a location \mathbf{x} where the contour starts leaking, points in the neighborhood of \mathbf{x} will have a large ϵ_1 value since the measure is high for widening regions. These regions correspond to non-tubular structures and should be penalized. These requirements are incorporated in a combined filter

$$\epsilon_2(\mathbf{x}) = \left(\frac{1}{N} \sum_{\mathbf{y}_i \in \mathcal{B}(\mathbf{x}, r)} \chi(\mathbf{y}_i) \right)^2 + \frac{2}{N} \sum_{\mathbf{y}_i \in \mathcal{B}(\mathbf{x}, r)} \chi(\mathbf{y}_i) \left(\sum_{\mathbf{y}_j \in \mathcal{B}(\mathbf{y}_i, r)} \chi(\mathbf{y}_j) \right). \quad (11)$$

The soft shape combined filter ϵ_2 yields the final curve evolution equation:

$$\mathbf{C}_t = \kappa \mathbf{N} - \phi \mathbf{N} + \alpha \epsilon_2(\mathbf{x}, p) \mathbf{N}, \quad (12)$$

where α is a parameter that controls the amount of penalty desired for non-tubular structures. In practice, ϵ_2 is normalized to the range $[0, 1]$. A value of $\alpha \leq 0.65$ is shown to provide satisfactory vessel segmentations. Larger values of α tend to erode the contour near branch points, i.e. only nearly straight and idealized vessels are preserved. Nain et al. [16] provide examples where the soft shape prior prevents a traditional active contour from leaking and results in more regular segmentations that would otherwise require extensive manual corrections.

3.2 Filter-based Segmentation

Filter-based segmentation methods use the response obtained from the convolution of an image with a filter to extract blood vessels. By selecting appropriate filters, different structure types can be revealed.

This section covers a selection of segmentation approaches that utilize image filters to derive a measure of similarity to vascular structures and extract them.

3.2.1 Notes on the Hessian

The eigenvalues of the Hessian matrix can be used to derive information about the local geometry in the neighbourhood of a vessel. A second order Taylor expansion of an image I with spatial coordinate \mathbf{x} gives $I(\mathbf{x} + \Delta\mathbf{x}) = I(\mathbf{x}) + \nabla I(\mathbf{x})\Delta\mathbf{x} + \frac{1}{2}\Delta\mathbf{x}^T \mathbf{H}(I(\mathbf{x}))\Delta\mathbf{x} + O(\Delta\mathbf{x}^3)$, where $\nabla I(\mathbf{x})$ is the gradient of I at \mathbf{x} and $\mathbf{H}(I(x))$ is the Hessian of I at \mathbf{x} . $\mathbf{H} \in \mathbb{R}^{3 \times 3}$ is the matrix composed of the second-order derivatives of I . When $\nabla I(\mathbf{x})$ is small, as is the case on the centerline of a blood vessel, the Hessian fully describes, up to second order, the local variation of the intensity in the direction of the associated eigenvectors. More precisely, the gradient describes the contours of the vessels whereas the Hessian gives information about the centerline, the axis, and the cross-section of the vessels. The Hessian matrix \mathbf{H} can be computed by convolving the image with the second partial derivatives of a Gaussian function $G(\mathbf{x}, \sigma) = \exp\left(-\frac{|\mathbf{x}|^2}{2\sigma^2}\right)$, due to the associative properties of the convolution operator for differentiation. This decomposition has the advantage of reducing the amount of computation needed from $O(n^3)$ to $O(3n)$. Due to the equality of mixed partials theorem, $I_{xy} = I_{yx}$, $I_{xz} = I_{zx}$, $I_{yz} = I_{zy}$, and therefore \mathbf{H} is symmetric. This symmetry has important properties. The eigenvalues of \mathbf{H} are real and its eigenvectors form an orthogonal basis. \mathbf{H} being a second-order operator, its eigenvalues $\lambda_1, \lambda_2, \lambda_3$ and associated eigenvectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ directly relate to the principal curvatures at \mathbf{x} and can be used to locate and characterize the geometry of blood vessels in an intensity image, as described in Section 3.2.2.

3.2.2 A Generalized Measure of Vesselness (Sato et al. [19])

Ideal blood vessels are generally represented as contrasting tubes with an isotropic Gaussian cross-section and with no centerline curvature. The original work by Sato et al. [19] manipulates this idealized model with respect to the Hessian operator to derive a measure of vesselness and segment blood vessels at different scales. This work is one out of many early Hessian-based vessel segmentation approaches that blossomed at the end of the 1990's [10].

Sato et al. [19] construct a vesselness measure based on the Hessian eigenvectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ and respective eigenvalues $\lambda_1, \lambda_2, \lambda_3$. The image is convolved with the second derivative of an isotropic Gaussian,

$$G(\mathbf{x}, \sigma_f) = \exp\left(-\frac{|\mathbf{x}|^2}{2\sigma_f^2}\right), \quad (13)$$

such that filter responses can be tuned to a specific width of 3D line. The eigenvalues of the Hessian matrix are assumed to be sorted by increasing value such that $\lambda_1 \geq \lambda_2 \geq \lambda_3$ and where $\lambda_2, \lambda_3 \leq 0$.

Sato et al. [19] suggest heuristics based on the Hessian of the idealized vascular model. \mathbf{e}_1 is oriented along the vessel centerline and its corresponding eigenvalue is small ($\lambda_1 \approx 0$) and $\mathbf{e}_2, \mathbf{e}_3$ span the cross-sectional plane of the vessel. The vessel is bright and isotropic along its cross-section, so both λ_2 and λ_3 are close, negative and have a large magnitude ($\lambda_2 \approx \lambda_3 \ll 0$). Again based on idealized structural models, other local structures can be characterized by the Hessian eigenvalues and are summarized in Table 1. Following these observations, a measure λ_{123} is developed that can discriminate between

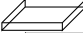





λ_1	λ_2	λ_3	Local Structure	2 nd Order Ellipsoid	Examples
$\pm\epsilon$	$\pm\epsilon$	$\pm\epsilon$	noisy	blob	noise
$\pm\epsilon$	$\pm\epsilon$	\blacktriangledown	plate (bright) 	line	cortex, skin
$\pm\epsilon$	$\pm\epsilon$	\blacktriangle	plate (dark) 		
$\pm\epsilon$	\blacktriangledown	\blacktriangledown	tube (bright) 	disk	vessel, bronchus
$\pm\epsilon$	\blacktriangle	\blacktriangle	tube (dark) 		
\blacktriangledown	\blacktriangledown	\blacktriangledown	blob (bright) 	blob	nodule
\blacktriangle	\blacktriangle	\blacktriangle	blob (dark) 		

Table 1: Hessian eigenvalues and their geometrical interpretation [19, 7, 12, 5]. The \blacktriangle symbol represents an eigenvalue much greater than 0, $\lambda_i \gg 0$; the \blacktriangledown symbol represents an eigenvalue much less than 0, $\lambda_i \ll 0$; ϵ represents an eigenvalue approximately zero, $\lambda_i \approx 0$.

line-like, plate-like, and vessel-like structures:

$$\lambda_{123} = \begin{cases} |\lambda_3| \left(\frac{\lambda_2}{\lambda_3}\right)^{\gamma_{23}} \left(1 + \frac{\lambda_1}{|\lambda_2|}\right)^{\gamma_{12}} & \text{for } \lambda_3 < \lambda_2 < \lambda_1 \leq 0 \\ |\lambda_3| \left(\frac{\lambda_2}{\lambda_3}\right)^{\gamma_{23}} \left(1 - \alpha \frac{\lambda_1}{|\lambda_2|}\right)^{\gamma_{12}} & \text{for } \lambda_3 < \lambda_2 < 0 < \lambda_1 < \frac{|\lambda_2|}{\alpha} \\ 0 & \text{otherwise,} \end{cases} \quad (14)$$

where γ_{23} and γ_{12} control the sharpness of the selectivity for the cross-section isotropy. A bright sheet-like structure will have $\lambda_3 \ll \lambda_2 \approx 0$, which yields $\lambda_{123} \approx 0$. A nonlinear structure will have $\lambda_1 \neq 0$, also implying a minimal vesselness measure for both $\lambda_1 < 0$ and $\lambda_1 > 0$. A bright line-like structure will have $\lambda_3 \approx \lambda_2 \ll 0$, thus λ_{123} is maximal with $\lambda_{123} \approx |\lambda_3|$.

The integration of a multiscale measure $M(\mathbf{x})$ is done by combining filter responses tuned to different line widths σ_f in Eq. 13. The line width σ_f maximizing the measure λ_{123} is selected. In practice, Sato et al. [19] use discretized values for σ_f on the basis of the width range of the anatomical structure of interest.

Sato et al. [19] place strict conditions on the Hessian eigenvalues. Only isotropic cross-sections are considered. Moreover, the Hessian eigenvalues are sorted by increasing value rather than absolute value, resulting in asymmetrical heuristics for defining bright and dark structures. A more important issue is that they fail to combine the vesselness measure in a neighborhood to achieve vessel completeness and connectivity.

3.2.3 Formalizing Scale and Hessian Eigenvalues (Krissian et al. [9])

Sato et al. [19] provide little mathematical formalism in the interpretation of the Hessian eigenvalues. Their vesselness measure is based on heuristics and on an experimental study of simple structure prototypes. Although Table 1 gives a feel for the relationship in between different eigenvalues for these structures, it is not clear how they transition from one to another.

Further work by Krissian et al. [9] addresses this issue by deriving an analytical expression for the eigenvalues of three different vessel models that together capture a large part of the structural variability found in blood vessels. 1) A cylindrical circular model describes the ideal blood vessel used by Sato et al. [19] ($\lambda_3 \approx \lambda_2 \ll \lambda_1 \approx 0$). This model is based on an isotropic 2D Gaussian intensity profile. 2) A toroidal

circular model extends the cylindrical circular model by allowing curvature in the centerline direction ($\lambda_3 \approx \lambda_2 \ll \lambda_1 \neq 0$). This model is based on the intensity profile

$$I_0(\mathbf{x}) = \exp\left(-\frac{(R - \sqrt{x^2 + y^2})^2 + z^2}{2\sigma_0^2}\right). \quad (15)$$

3) A cylindrical elliptical model extends the cylindrical circular model by providing anisotropy in the cross-section ($\lambda_3 \leq \lambda_2 \ll \lambda_1 \approx 0$). The model is based on the intensity profile

$$I_0(\mathbf{x}) = \exp\left[-\frac{1}{2}\left(\frac{x}{\sigma_x}\right)^2 + \left(\frac{y}{\sigma_y}\right)^2\right]. \quad (16)$$

This study shows that in general, eigenvalues are very sensitive to the vessel curvature in the centerline direction and to elliptical cross-sections, two variations that are ignored by Sato et al. [19]. Moreover, the estimation of the cross-sectional plane spanned by the two eigenvectors $\mathbf{e}_2, \mathbf{e}_3$ is shown to be stable to noise and different vessel geometries. This suggests the use of the Hessian eigenvalues for discriminating vessel-like structures from other structures, alike the work by Sato et al. [19], but with the addition of the Hessian eigenvectors and the image gradient to extract vessel centerlines. Based on these findings, an algorithm for vessel centerline extraction and vessel radius estimation is developed. The algorithm is divided into three principal steps: a) a vessel response function is evaluated at different scales, b) the local maxima of this response function is found, and c) vessels are reconstructed using centerline and size information.

The vesselness response function $R_t(\mathbf{x})$ at a point \mathbf{x} is a radial projection of the image gradient on a circle of radius $\tau\sqrt{t}$ in the cross-sectional plane spanned by the eigenvectors ($\mathbf{e}_2, \mathbf{e}_3$):

$$R_t(\mathbf{x}) = \frac{1}{2\pi} \int_0^{2\pi} -\nabla I(\mathbf{x} + \tau\sqrt{t}\mathbf{v}_\theta) \cdot \mathbf{v}_\theta d\theta, \quad (17)$$

where $\mathbf{v}_\theta = \cos(\theta)\mathbf{e}_2 + \sin(\theta)\mathbf{e}_3$. The purpose of the parameter τ is to align the circle boundary with the frontier of the vessel at a scale that yields the maximal response.

The computation of centerline extrema is done via the *marching lines* algorithm [22], an isocontour extraction method. The algorithm transforms the 3D volume into an implicit volume $\mathbb{R}^3 \rightarrow \mathbb{R}$ defined by the response of the kernel defined in Eq. 17 at each voxel \mathbf{x} . At \mathbf{x} , the algorithm can extract a line by interpreting the response function in a neighborhood. For instance, a small response on the left and right side of a location where a large response is found indicates that \mathbf{x} is a local extrema and potentially part of a vessel centerline.

Finally, a hysteresis thresholding is applied such that non-tubular structures and noise are removed. Short segments that do not change the topology of the image are then also removed in order to get rid of spurious details. The centerlines are then smoothed out to give a better visualization. More complete vessels are reconstructed by using the scale maxima σ found during the detection of centerline extrema.

Following this work, several points remain to be addressed. The circular cross-section is an assumption that has yet to be dropped. Moreover, the extrema detection method will poorly work in regions where non-vascular structures have a high intensity gradient, and near vessel edges. No subsequent vessel reconnection is carried out in order to reconstruct the vasculature network. This reconnection could be performed straightforwardly after the vessel extraction process but this would be costly since neighbouring vessel segments can be discovered independently and are disconnected from one another and possibly numerous.

3.2.4 Vessel Centerline as Intensity Ridges (Aylward and Bullitt [1])

Ridge-based methods embed an N -dimensional image in $(N + 1)$ -dimensional space where the extra dimension is the intensity. By doing this, a 2D vessel image becomes a ridge in a 3D image where intensity maps to height, as shown in Fig. 3 c) and d).




Figure 3: Vessels seen as ridges in an intensity map. (Adapted from Aylward and Bullitt [1])

Ridge traversal is employed by Aylward and Bullitt [1] for segmenting blood vessels. This work is motivated by the observation that identifying blood vessels as a single entity, a ridge, is less sensitive to image noise than identifying them as a combination of separate entities, two edges. Indeed, edge-detection is generally performed using smaller-scale measures and the requirement of identifying two boundaries can potentially reject many true positives. They also point out that the vessel centerline can be located as the high point in an intensity ridge. This high point can be used for subsequent processing to stabilize the detection of vessel boundaries and to resolve boundary ambiguities.

Analogously to Sato et al. [19] and Krissian et al. [9], Aylward and Bullitt [1] use Hessian matrix eigenvalues to help distinguish between local structures in the image, more precisely ridges, valleys, saddles, planes, and spheres (blobs). In the development of their methods, Aylward and Bullitt [1] assume that the blood vessel are brighter than the background. They define three conditions that must hold for a point \mathbf{x} to be located on the centerline of a vessel: a) \mathbf{x} is located on a ridge, i.e. $\lambda_3 \leq \lambda_2 < 0$, b) \mathbf{x} is an extrema, i.e. $\mathbf{e}_2 \cdot \nabla I = \mathbf{e}_3 \cdot \nabla I = \epsilon_2 \approx 0$, and c) the ridge is central to an object that has a nearly circular cross-section, i.e. $\frac{\lambda_2}{\lambda_3} \geq 1 - \epsilon_3 \approx 0$. The first condition implies that the eigenvectors $\mathbf{e}_2, \mathbf{e}_3$ are the orthogonal directions normal to the ridge and that \mathbf{e}_1 approximates a tangent to the ridge. The second condition implies that the image gradient ∇I points towards the intensity maxima which, in the intensity height map, roughly points in the direction of \mathbf{e}_1 which is orthogonal to $\mathbf{e}_2, \mathbf{e}_3$. The value of ϵ_2, ϵ_3 directly controls how close to an idealized tubular vessel a particular structure must be to be characterized as such.

The focus of the work of Aylward and Bullitt [1] is therefore to determine the remaining two ridge parameters that will satisfy the ridge conditions. That is, \mathbf{x} , a point located on the centerline of a vessel, i.e. on a ridge, and σ , the scale at which computations are made, and thus the width of the vessel at \mathbf{x} . In doing so, they develop two algorithms: one being the localization of \mathbf{x} given a fixed σ , the second being a dynamic optimization of σ while localizing \mathbf{x} . Initial ridge points \mathbf{x} are found by thresholding the image based on the assumption that the blood vessels are brighter than the background. Thresholding is done both globally (intensity-based) and locally (contrast-based). Global thresholding will generally not recover complete vessels but the authors assume that a sufficient number of points will be obtained and that gradient ascent will be able to reconstruct the vessel's centerline. Given an initial scale σ_0 , using gradient ascent an intensity maximum is reached by following the local gradient ∇I and the eigenvectors

$\mathbf{e}_i, i = 2, 3$ that define the cross-sectional plane, i.e.

$$\mathbf{x}_t = \mathbf{x}_{t-1} + \left(\underset{h}{\operatorname{argmax}} I(\mathbf{x}_{t-1} + h\nabla I) \right) \nabla I, \quad (18)$$

$$\mathbf{x}_t \rightarrow \mathbf{x}_t + \left(\underset{h}{\operatorname{argmax}} I(\mathbf{x}_t + h\mathbf{e}_i) \right) \mathbf{e}_i. \quad (19)$$

At each iteration t , \mathbf{x}_t is accepted if the ridge conditions are satisfied. The dynamic optimization of the ridge scale σ at \mathbf{x} is performed by convoluting a kernel composed of a ring of spherical boundary operators at a distance ρ . The ridge scale σ is selected as the ρ that maximizes the response given by the kernel. Finally, a complete vessel centerline is obtained by following the ridge at \mathbf{x}_n in the positive \mathbf{e}_1 and negative $-\mathbf{e}_1$ local tangent direction by an amount β . This approach is similar to the marching lines algorithm employed by Krissian et al. [9]. Here, β determines how far from the current point the shifted ridge normal plane should be extended. The new ridge point \mathbf{x}_{n+i} is kept if a line search based on Eq. 19 at \mathbf{x}_{n+i} yields a maxima that satisfies the ridge conditions.

3.2.5 Minimum Paths in 4D Scale Space (Li and Yezzi [11])

The methods [19, 9, 1] presented so far explore the effect of local structure on the eigenvectors and eigenvalues of the Hessian matrix. They present practical algorithms for extracting the centerline and the width of vessel-like objects from medical images. In these approaches, the choice of an appropriate scale for performing kernel computations is discussed profusely, particularly in Krissian et al. [9]; however, scale is always treated as a local quantity and no scale correspondence is made between adjacent voxels. Wink et al. [26] address this issue by designing a centerline extraction method based on a minimum cost path approach in a 4D image, where the fourth dimension is defined by scale.

The filter is based on the work by Sato et al. [19] for describing the Hessian eigenvalues λ_i . The vesselness response filter $R(\mathbf{x}, \sigma)$ at a voxel \mathbf{x} and at scale σ is defined by

$$R(\mathbf{x}, \sigma) = \begin{cases} 0 & \text{for } \lambda_2 > 0 \\ \exp\left(-\frac{\lambda_1^2}{2\beta^2}\right) \left(1 - \exp\left(\frac{\lambda_1^2 + \lambda_2^2}{-2c^2}\right)\right) & \text{otherwise} \end{cases} \quad (20)$$

where β and c are two parameters that control the sensitivity of the filter. The response is maximal when the scale σ approximates the vessel radius. At each voxel \mathbf{x} in the image I , the response $R(\mathbf{x}, \sigma)$ is evaluated for a discrete range of σ values. The response from all voxels is combined in a 3D cost image, where intensity is replaced by a cost function, $C(\mathbf{x}, \sigma)$, such that a high response yields a small cost. The minimum value of the response image, R_l , is used as an epsilon to prevent a division by zero when the response $R(\mathbf{x}, \sigma)$ is null:

$$C(\mathbf{x}, \sigma) = \begin{cases} \frac{1}{R_l} & \text{for } R(\mathbf{x}, \sigma) = 0 \\ \frac{1}{R(\sigma)} & \text{otherwise.} \end{cases} \quad (21)$$

The transition cost $a(\mathbf{n}, \mathbf{n}')$ for going from a node $\mathbf{n} = (\mathbf{x}, \sigma)$ to a node $\mathbf{n}' = (\mathbf{x}', \sigma')$ in the cost image is a weighted function of the cost at \mathbf{n}' , $a(\mathbf{n}, \mathbf{n}') = w(\sigma, \sigma')C(\mathbf{x}', \sigma')$, where w is a function of the scale in two adjacent nodes that is used to control the extent to which the scale is allowed to change in a given vessel. If vessels are known to have a distinct and constant width, a weight function that yields a large cost for different scales will be preferable. The minimum path is found using Dijkstra's algorithm.

3.2.6 Minimum Paths in 4D Potential Fields (Li and Yezzi [11])

Li and Yezzi [11] develop a semi-automatic vessel segmentation approach based on minimal paths in 4D space, similarly to Wink et al. [26]. The main differences are found in the formulation of the energy formulation that defines a minimum path and in the automation of the procedure. In Li and Yezzi [11], vessel centres are not detected automatically. Instead, the user selects two vessel endpoints in-between which a globally minimizing trajectory is found. This minimum trajectory is found as a path in a gradient potential. Two differential potential functions are used. The first potential function is defined by exploiting the fact that blood vessels have a nearly constant intensity, that is the image intensity gradient is negligible for voxels within a blood vessel. The mean $\mu(\mathcal{B})$ and standard deviation $\sigma(\mathcal{B})$ of voxel intensities in a ball $\mathcal{B}(\mathbf{x}, r)$ of volume $V(\mathcal{B})$ centered at \mathbf{x} and with radius r are defined as

$$\mu(\mathcal{B}) = \frac{1}{V} \int_{\mathcal{B}} I(\mathbf{p}) d\mathbf{p}, \quad \sigma^2(\mathcal{B}) = \frac{1}{V} \int_{\mathcal{B}} (I(\mathbf{p}) - \mu(\mathcal{B}))^2 d\mathbf{p}. \quad (22)$$

A potential \tilde{P}_1 is then defined that is targeted towards vessel structures with constant intensity:

$$\tilde{P}_1(\mathbf{x}) = \omega + \frac{\lambda_1}{r^2} [\mu(\mathcal{B}(\mathbf{x})) - \mu(\mathcal{B}_0)]^2 + \frac{\lambda_2}{r^2} [\sigma^2(\mathcal{B}(\mathbf{x})) - \sigma^2(\mathcal{B}_0)]^2, \quad (23)$$

where \mathcal{B}_0 is the sphere at the first endpoint of the path, x_0 , ω is a constant that controls the smoothness of the path, and λ_1, λ_2 are weights for the mean difference and variance difference between spheres along the path and the starting sphere. The potential \tilde{P}_1 is designed to increase if a sphere's radius differs from the width of the tubular structure. The potential is also weighted by the inverse of the sphere radius r such that larger spheres produce a smaller potential and are therefore favoured over smaller ones. This characteristic ensures that the vessel surface is maximally reconstructed.

The second potential \tilde{P}_2 relaxes the intensity constancy condition along the vessel centerline. Rather than deriving an expression based on the volume of the ball $\mathcal{B}(\mathbf{x}, r)$ of radius r centred at \mathbf{x} , its boundary $\partial\mathcal{B}$ is used instead and $V(\partial\mathcal{B}), \mu(\partial\mathcal{B}), \sigma^2(\partial\mathcal{B})$ are recomputed in Eq. 22. The mean difference $\delta_\mu(\partial\mathcal{B}) = |\mu(\partial\mathcal{B}) - \mu(\partial\mathcal{B}')|$ and the variance difference $\delta_{\sigma^2}(\partial\mathcal{B}) = |\sigma^2(\partial\mathcal{B}) - \sigma^2(\partial\mathcal{B}')|$ are also introduced, where $\partial\mathcal{B}' = \partial\mathcal{B}(\mathbf{x}, r - h)$ is a shell of radius less than $\partial\mathcal{B}$. The potential is given by

$$\tilde{P}_2(\mathbf{x}) = \omega + \frac{\lambda_1}{1 + \delta_\mu^2(\partial\mathcal{B}(\mathbf{x}))} + \frac{\lambda_2}{1 + \delta_{\sigma^2}^2(\partial\mathcal{B}(\mathbf{x}))}, \quad (24)$$

where $\omega, \lambda_1, \lambda_2$ are parameters introduced similarly as in Eq. 23. The resulting Eikonal equation $||\nabla U_{p0}|| = \tilde{P}$ with $U_{p0} = 0$ is solved for using the fast marching method, similar in design to the method employed in Wink et al. [26] for finding minimal paths. Once the minimal potential field is obtained, the minimal path can be traced back from the last endpoint to the first by following the gradient field.

The advantage of the work described by Li and Yezzi [11] over many segmentation methods is its ability to simultaneously extract vessel surfaces and centerlines. As the vessel centerline is reconstructed by solving the minimal path problem in a potential field, the vessel surface can be directly obtained as the envelope of the spheres \mathcal{B} that lie along the vessel centerline. This method also has several drawbacks. First, it is semi-automatic such that endpoints between vessels need to be specified manually. It can be argued that this method could be combined with other filter-based methods for detecting vessel centres from which minimal paths could be propagated, similarly to Wink et al. [26]. Another drawback is the extensive use of model parameters. The authors explicitly mention that the optimal set of parameters is very application dependent. By the number of required parameters, it is unclear how much work needs to be put in the optimization of these parameters for the method to work optimally.

3.2.7 Combining Filters and Flows (Descoteaux et al. [5])

So far, Section 3.2 has shown that pure filter-based approaches can be used successfully to segment neurovasculature. These approaches are fundamentally different from the active contours techniques described in Section 3.1, where computations are done in a dynamic setting rather than a static one. The method presented by Descoteaux et al. [5] combines the response given by filters with an active contour formulation. As before [26, 9, 19], a vesselness measure [7] is used based on the eigenvalues of the Hessian matrix and where vessels are assumed to be dark against a light background

$$V(\sigma) = \begin{cases} 0 & \text{for } \lambda_2 < 0 \text{ or } \lambda_3 < 0 \\ \left(1 - \exp\left(-\frac{R_A^2}{2\alpha^2}\right)\right) \exp\left(-\frac{R_B^2}{2\beta^2}\right) \left(1 - \exp\left(-\frac{S^2}{2c^2}\right)\right) & \text{otherwise} \end{cases} \quad (25)$$

where $\mathcal{R}_B = \frac{|\lambda_1|}{\sqrt{|\lambda_2\lambda_3|}}$, $\mathcal{R}_A = \left|\frac{\lambda_2}{\lambda_3}\right|$ and $\mathcal{S} = \|\mathbf{H}\|_F = \sqrt{\sum_{i=1}^3 \lambda_i^2}$. α, β, c are thresholds that control the sensitivity of the line filter to the measures $\mathcal{R}_A, \mathcal{R}_B, \mathcal{S}$. \mathcal{R}_A differentiates sheet-like from tube-like structures and \mathcal{R}_B differentiates tube-like from blob-like and noise-like structures. The vesselness measure is used to find potential vessel centerline candidates. A multiscale analysis is performed by considering multiple values of σ in Eq. 25 and selecting the scale σ_0 at which the response is maximal. This maximizing σ_0 approximates the radius of the underlying vessel.

The second step is to propagate the vesselness measure from the centerline of a vessel to its expected boundary by considering all voxels that lie within an ellipsoid aligned with the centerline and with a semi-minor axis length proportional to the estimated vessel radius. The contribution of the vesselness measure from a centerline maxima \mathbf{x} to a given voxel \mathbf{x}_e depends on its position in the elliptical coordinate system. The vesselness measure is scaled by the magnitude of the projection of the vector $\mathbf{x}_e - \mathbf{x}$ onto the cross-sectional plane formed by the eigenvectors $\mathbf{e}_2, \mathbf{e}_3$ of the Hessian matrix. The resulting accumulation of vesselness measures is defined as the ϕ distribution and is maximal close to vessel boundaries.

An extended vector field is then constructed from the image gradient ∇I and the distribution ϕ , $\mathbf{V} = \phi \frac{\nabla I}{|\nabla I|}$. This vector field is directed in the direction of the gradient of the image, which is orthogonal to the vessel boundary. The gradient is maximal on the vessel boundary and ϕ is large near vessel boundaries, making \mathbf{V} a good detector of vessel boundaries. A level set active contour formulation is then used to embed the gradient field and conclude the segmentation algorithm, as shown in Section 3.1. More precisely, \mathbf{V} is used in a flux maximizing geometric flow [24] that yields

$$\mathcal{S}_t = \left(\nabla \phi \cdot \frac{\nabla I}{|\nabla I|} + \phi \kappa_I \right) \mathbf{N}, \quad (26)$$

where κ_I is the mean curvature of the iso-intensity level set of the image. The first term in Eq. 26 pushes back the evolving surface if it overshoots the boundary since $\nabla \phi$ has a zero-crossing the vessel boundaries and ∇I does not change sign. The second term has anisotropic smoothing properties and thus stabilizes the evolution of the manifold. Outside vessels, both ϕ and $\nabla \phi$ are zero which prevents the flow from leaking. Again, vessel centerlines can be extracted as the zero level set of the manifold \mathcal{S} .

3.3 Model-based Segmentation (Tyrrell et al. [23])

Tyrrell et al. [23] propose a vessel segmentation method that is neither based on image filters nor active contours. Rather than construct a virtual ellipsoid based on the eigendecomposition of the Hessian

matrix, here blood vessels are modelled explicitly as physical piecewise *superellipsoids*. The superellipsoid approximates both the boundary and centerline, through its semi-major and semi-minor axes.

Cylindroidal forms of a superellipsoid $S(\mathbf{x}, \epsilon)$ are used which can be described as

$$S(\mathbf{x}, \epsilon) = \left(|x|^2 + |y|^2 \right)^{\frac{1}{\epsilon}} + |z|^{\frac{1}{\epsilon}} \quad (27)$$

where ϵ controls the shape of the superellipsoid, ranging from ellipsoidal to elliptical cylindrical forms. Eq. 27 describes a superellipsoid that has a fixed scale, is aligned with the canonical coordinate system, and is centred at the origin. A more general expression can be obtained by applying a coordinate transformation: $S(\mathbf{x}, \boldsymbol{\beta}) = S(\mathbf{D}(\boldsymbol{\sigma})^{-1} \mathbf{R}(\boldsymbol{\phi})^T (\mathbf{x} - \boldsymbol{\mu}), \epsilon)$, where $\mathbf{D}(\boldsymbol{\sigma}) = \text{diag}(\sigma_1, \sigma_2, \sigma_3)$ is an anisotropic scaling matrix, $\mathbf{R} \in \mathbb{R}^3$ is a rotation matrix described by $\boldsymbol{\phi}$ and $\boldsymbol{\mu}^T = (\mu_x, \mu_y, \mu_z)$ is a translation vector. These parameters are represented by the parameter vector $\boldsymbol{\beta} = (\boldsymbol{\mu}^T, \boldsymbol{\sigma}, \boldsymbol{\phi}, \epsilon)$. Tyrrell et al. [23] derive a procedure for estimating parameters in $\boldsymbol{\beta}$. A strong assumption is placed on the grayscale appearance of voxels in the image that are contained inside (foreground) and outside (background) of the superellipsoid. These two regions, $\mathcal{R}(\boldsymbol{\beta})$ and $\Omega \setminus \mathcal{R}(\boldsymbol{\beta})$, are respectively described as homogeneous with constant intensity values I_F and I_B . Vessels are assumed to be dark and contrast against a light background such that $I_B \gg I_F \approx 0$.

Based on the foreground and background intensity assumptions, a cost function $L(\boldsymbol{\beta}, I_F, I_B)$ is developed such that its maxima yields the optimal superellipsoid describing a vessel. The cost function is evaluated both at the interior $\mathcal{R}(\boldsymbol{\beta})$ and exterior $\Omega \setminus \mathcal{R}(\boldsymbol{\beta})$ points \mathbf{x} of the superellipsoid:

$$L(\boldsymbol{\beta}, I_F, I_B) = \int_{\mathcal{R}(\boldsymbol{\beta})} g(I(\mathbf{x}), I_F) d\mathbf{x} + \int_{\Omega \setminus \mathcal{R}(\boldsymbol{\beta})} g(I(\mathbf{x}), I_B) d\mathbf{x} \quad (28)$$

where $g(I_1, I_2) \equiv \log f(I_1 - I_2)$ describes noise in the image, i.e. given a noise model, $f(I_1 - I_2)$ is the probability that the intensity $I_1 - I_2$ corresponds to noise. Hence, if $I_1 - I_2$ is sufficiently small for all \mathbf{x} in Eq. 28, the superellipsoid fits well the data. Since I_F and I_B are not known in advance, Eq. 28 is maximized by alternating between intensity and parameter estimates. At each iteration t , the parameter vector $\boldsymbol{\beta}$ is updated using a gradient descent

$$\boldsymbol{\beta}_{t+\Delta t} = \boldsymbol{\beta}_t + \frac{d\boldsymbol{\beta}}{dt} \Delta t. \quad (29)$$

The foreground and background intensities are then updated by fixing the model parameters $\boldsymbol{\beta}$ and observing how the cost in Eq. 28 varies, based on image noise estimates.

Some issues arise in the fitting of the superellipsoid model to vessel segments. Adjacent vessels appear as background intensity outliers to one another such that intensity estimates may fail for densely intertwined vessel networks. Moreover, the superellipsoid model has an explicit geometry, making the connection between two superellipsoids piecewise continuous. An implicit description of the model could yield a smoother representation for subsequent visualization.

3.4 Histogram-based Segmentation (Wilson and Noble [25])

Histogram-based approaches use the histogram of voxel intensities to identify whether a vessel is detected at a given voxel given its intensity. A thresholding algorithm can then be used to reconstruct vessels from this voxel volume. These methods contrast heavily from the filter-based and active contour methods

presented so far, as they address the distribution of voxel intensities directly and ignore their spatial component. Early work is carried out by Wilson and Noble [25]. They observe that in an MRA image, the histogram shape is typically bimodal. The first, low-intensity, peak represents mainly background tissue (external to the head) and cerebrospinal fluid (CSF). The second, high-intensity, peak represents a combination of brain tissue, eyes, skin, and bone. Vessel voxels lie at the upper tail of the high intensity peak.

Wilson and Noble [25] propose to reconstruct the distribution of voxel intensities by a mixture of three distributions: one Gaussian $G(x, \sigma_1)$ for the low-intensity peak, one Gaussian $G(x, \sigma_2)$ for the high-intensity peak, and a uniform distribution $\omega_0 \frac{1}{I}$ located at the right-tail of the high-intensity Gaussian. The uniform distribution corresponds to the region where voxels start having a smaller probability of being a blood vessel. This is represented by the mixture distribution

$$p(x) = \omega_0 \frac{1}{I} + \sum_{k=1}^2 w_k \frac{1}{\sqrt{2\pi\sigma_k^2}} \exp \left[-\frac{1}{2} \left(\frac{x - \mu_k}{\sigma_k} \right)^2 \right]. \quad (30)$$

The parameters of the distributions are found in an expectation maximization (EM) algorithm based on the voxel histogram. The region where the second Gaussian meets the uniform distribution is marked as the high intensity threshold in a hysteresis thresholding approach. The lower threshold is determined more carefully since slow blood flow (in aneurysms, for instance) can lead to lower intensities in MRA. Following the hysteresis thresholding, the segmentation is refined by propagating the vessel classification using local connectivity. Sufficiently high intensity voxels that are close to segmented vessel voxels are given a higher probability of being vessel voxels. An histogram is constructed for these lower intensity voxels and the EM approach is used once more to fit the distribution, the minimum of which is retained for the lower bound of another hysteresis thresholding.

There are many issues related to the sole use of histograms to determine vessel voxels. The method presented in Wilson and Noble [25] is purely global since the whole spectrum of voxel intensities is fitted by a mixture of probability distributions. Other methods will try to fit the intensity histogram using more comprehensive models of the intensity distribution. Nevertheless, histogram-based methods lack structural accuracy in vessel extraction – the vessels are not modelled explicitly and no multiscale analysis is performed – and their applicability typically lies in a pre-filtering pass. More important issues arise from the nature of the imaging itself. Tissues with a short T1 relaxation time such as fat, methemoglobin, and other contrast enhancing structures can produce a very bright signal, undistinguishable from flowing blood [20].

4 Conclusion

This literature review presented a non-exhaustive selection of relevant papers in the field of brain vessel segmentation. Segmentation methods were introduced that make use of a combination of active contours, filters, vessel models, and image histograms. Filter-based approaches were shown to be robust and flexible to vessel scale, and to be useful in the initialization of other approaches. No assumptions are made on the topology of a vessel so branching is dealt with naturally. However, these methods are hard to assist manually, one major advantage of active contour methods, particularly in their explicit formulation. Model-based approaches such as superellipsoids show great potential. The degrees of

freedom of an explicit superellipsoid simultaneously captures centerlines and boundaries while being robust to noise. Finally, histogram-based approaches are useful for a fast thresholding and visualization of vessels, assuming they are well represented in the image.

The segmentation methods presented in this literature review are mostly concerned with the detection of vascular structures rather than in the interpretation of the resulting images with respect to the image modality. Neurovascular images are considered as volumes containing tubular structures for validating their algorithms. The focus of this review was put on these algorithms in order to make an abstraction of the vascular structures. A temporal quantification of vessel branch points, vessel spatial relationship, segment extraction, diameter measurements, complexity and complication assessment, intra-patient models, and inter-patient comparisons leave a lot of room for future research. More detailed reviews include Lesage et al. [10], and Kirbas and Quek [8].

References

- [1] S. Aylward and E. Bullitt. Initialization, noise, singularities, and scale in height ridge traversal for tubular object centerline extraction. *Medical Imaging, IEEE Transactions on*, 21(2):61–75, 2002.
- [2] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International journal of computer vision*, 22(1):61–79, 1997.
- [3] E. Cheney and D. Kincaid. *Numerical mathematics and computing*. Brooks/Cole Pub Co, 2007.
- [4] T. Deschamps. *Curve and Shape Extraction with Minimal Path and Level-Sets techniques - Applications to 3D Medical Imaging*. PhD thesis, Université Paris-IX Dauphine, Place du maréchal de Lattre de Tassigny, 75775 Paris Cedex, Dec. 2001.
- [5] M. Descoteaux, D. Collins, and K. Siddiqi. A geometric flow for segmenting vasculature in proton-density weighted mri. *Medical image analysis*, 12(4):497–513, 2008.
- [6] E. Dijkstra. A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1):269–271, 1959.
- [7] A. Frangi, W. Niessen, K. Vincken, and M. Viergever. Multiscale vessel enhancement filtering. *Medical Image Computing and Computer-Assisted Intervention—MICCAI’98*, pages 130–137, 1998.
- [8] C. Kirbas and F. Quek. A review of vessel extraction techniques and algorithms. *ACM computing surveys*, 36(2):81–121, 2004.
- [9] K. Krissian, G. Malandain, N. Ayache, R. Vaillant, and Y. Troussel. Model-based detection of tubular structures in 3d images. *Computer Vision and Image Understanding*, 80(2):130–171, 2000.
- [10] D. Lesage, E. Angelini, I. Bloch, and G. Funka-Lea. A review of 3d vessel lumen segmentation techniques: Models, features and extraction schemes. *Medical Image Analysis*, 13(6):819–845, 2009.
- [11] H. Li and A. Yezzi. Vessels as 4-d curves: Global minimal 4-d paths to extract 3-d tubular surfaces and centerlines. *Medical Imaging, IEEE Transactions on*, 26(9):1213–1223, 2007.

- [12] Q. Lin. *Enhancement, Detection, and Visualization of 3D Volume Data*. PhD thesis, Linköping University, SE-581 83 Linköping, Sweden, May 2003. Dissertations No. 824, ISBN 91-7373-657-0.
- [13] L. Lorigo, O. Faugeras, W. Grimson, R. Keriven, R. Kikinis, A. Nabavi, and C. Westin. Curves: Curve evolution for vessel segmentation. *Medical Image Analysis*, 5(3):195–206, 2001.
- [14] T. McInerney and D. Terzopoulos. T-snakes: Topology adaptive snakes. *Medical Image Analysis*, 4(2):73–91, 2000.
- [15] P. Morris. *Practical neuroangiography*. Lippincott Williams & Wilkins, 2007.
- [16] D. Nain, A. Yezzi, and G. Turk. Vessel segmentation using a shape driven flow. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2004*, pages 51–59, 2004.
- [17] A. Osborn. *Diagnostic cerebral angiography*. Lippincott Williams & Wilkins, 1999.
- [18] P. Reimer and P. Landwehr. Non-invasive vascular imaging of peripheral vessels. *European radiology*, 8(6):858–872, 1998.
- [19] Y. Sato, S. Nakajima, N. Shiraga, H. Atsumi, S. Yoshida, T. Koller, G. Gerig, and R. Kikinis. Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. *Medical Image Analysis*, 2(2):143–168, 1998.
- [20] G. Schneider, M. Prince, and J. Meaney. *Magnetic resonance angiography: techniques, indications, and practical applications*. Springer Verlag, 2005.
- [21] G. Sundaramoorthi, A. Yezzi, and A. Mennucci. Sobolev active contours. *International Journal of Computer Vision*, 73(3):345–366, 2007.
- [22] J. Thirion and A. Gourdon. The marching lines algorithm: New results and proofs. Technical Report RR-1881-1, INRIA, Sophia-Antipolis, 1993.
- [23] J. Tyrrell, E. di Tomaso, D. Fuja, R. Tong, K. Kozak, R. Jain, and B. Roysam. Robust 3-d modeling of vasculature imagery using superellipsoids. *Medical Imaging, IEEE Transactions on*, 26(2):223–237, 2007.
- [24] A. Vasilevskiy and K. Siddiqi. Flux maximizing geometric flows. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(12):1565–1578, 2002.
- [25] D. Wilson and J. Noble. Segmentation of cerebral vessels and aneurysms from mr angiography data. In *Information Processing in Medical Imaging*, pages 423–428. Springer, 1997.
- [26] O. Wink, W. Niessen, and M. Viergever. Multiscale vessel tracking. *Medical Imaging, IEEE Transactions on*, 23(1):130–133, 2004.
- [27] D. Yang, J. Zheng, A. Nofal, J. Deasy, and I. El Naqa. Techniques and software tool for 3d multimodality medical image segmentation. *Journal of Radiation Oncology Informatics*, 1(1):1–22, 2009.

Overall Performance: satisfactory_____, NOT satisfactory_____

Comments of the Evaluation Committee:

Chair of the Evaluation Committee:

Date: