Mueller, Daniel C. and Maeder, Anthony J. (2008) Robust Semi-automated Path Extraction for Visualising Stenosis of the Coronary Arteries. *Computerized Medical Imaging and Graphics*.

Robust Semi-automated Path Extraction for Visualising Stenosis of the Coronary Arteries

Daniel Mueller [a,*], Anthony Maeder [b]

[a] *Queensland University of Technology, Brisbane, Queensland, Australia*

[b] *School of Computing and Mathematics, University of Western Sydney, New South Wales, Australia*

Abstract

Computed tomography angiography (CTA) is useful for diagnosing and planning treatment of heart disease. However, contrast agent in surrounding structures (such as the aorta and left ventricle) makes 3-D visualisation of the coronary arteries difficult. This paper presents a composite method employing segmentation and volume rendering to overcome this issue. A key contribution is a novel Fast Marching minimal path cost function for vessel centreline extraction. The resultant centreline is used to compute a measure of vessel lumen, which indicates the degree of stenosis (narrowing of a vessel). Two volume visualisation techniques are presented which utilise the segmented arteries and lumen measure. The system is evaluated and demonstrated using synthetic and clinically obtained datasets.

Key words: vessel segmentation; volume visualisation; minimal path; Fast Marching; stenosis; cardiology

1. Introduction

Coronary heart disease is a major health concern. This disease stems from the underlying problem of *atherosclerosis*, which is a build up of plaque (consisting of substances including cholesterol, calcium, and others) on the interior surface of arteries supplying the heart. Coronary heart disease typically manifests in two forms: heart attack, and angina. A heart attack occurs when blood flow is completely blocked, typically from a

* Corresponding author: 2 George Street GPO Box 2434 Brisbane, Queensland, Australia, 4001.
Email addresses: d.mueller@qut.edu.au (Daniel Mueller), anthony@scm.uws.edu.au (Anthony Maeder).

dislodged portion of plaque. Angina — typically brought on by physical activity — is a chest pain or discomfort caused by an inadequate blood flow due to a narrowed artery.

Computed tomography angiography (CTA) provides high-resolution, high-contrast images of the thoracic cavity and as such is emerging as a useful imaging modality for diagnosis and treatment planning for coronary heart disease [1]. An intravenous contrast agent (such as an iodine-based compound) is injected into the patient causing the blood — and hence vessels — to exhibit high intensities in the resultant images [31]. In practice, motion artefacts due to the beating heart must be reduced using electrocardiographic (ECG) retrospective reconstruction (called ECG gating) [26].

From the acquired images, radiologists and cardiac surgeons require tools to easily identify *stenotic* (narrowed or constricted) arteries. A number of post-processing techniques are currently employed including: thin-slab maximum intensity projection (MIP) [12], curved planar reformatting (CPR) [17], and direct volume rendering (DVR) [35]. Each of these techniques exhibit varying strengths and weaknesses, depending on the task: MIP is useful for visualising calcified plaques, however 3-D information is discarded; CPR lays flat vessels of interest, but surrounding contextual structures appear deformed; DVR can depict the 3-D relationship between vascular and contextual structures, however specifying display parameters to clearly visualise the arteries is difficult and sometimes not possible. Hybrid rendering approaches [38] (which display both direct volume rendered and segmented polygonal structures) or tagged volume rendering [25] (which uses a number of *a priori* binary volumes to separate structures of interest) are other suitable techniques.

This paper proposes the use of segmentation methods to aid visualisation of stenotic vessels. The proposed technique is relevant for a range of vascular images and applications, however the focus is on the coronary arteries in CTA. The method consists of four stages: vessel enhancement, centreline extraction, stenosis measure computation, and volume visualisation. The centreline extraction stage extends an existing technique [29] by deriving a novel cost function using morphological top-hat by opening to enhance the vessels. The resultant continuous centreline is then used to compute a quantitative measure of stenosis based on the vessel radius. Two volume visualisation techniques using the resultant segmentation and measure — one based on MIP, one based on DVR — are discussed and demonstrated using clinically obtained datasets. The segmentation method is evaluated using various synthetic and clinical datasets using three criteria: efficiency, accuracy, and reproducibility.

2. Related Work

Vessel enhancement and segmentation is a broad area of research; a partial review of the field can be found in [19]. The existing work can be loosely organised into five categories: differential geometry, active contours, skeletonization, tracking, and minimal path extraction methods.

Differential geometry approaches utilise the differentiability of Euclidean space to derive measures which indicate the 'vesselness' of each pixel in an image. Sato et al. [28] and Frangi et al. [13] both proposed multi-scale line filters based on Gaussian differentials. Manniesing et al. [23] extended this to control anisotropic diffusion filtering, which smooths inside the vessels while maintaining sharp boundaries. Another approach pre-

sented in [20] computes the degree of belonging to the medial axis (centreline). This 'medialness function' is realised by convolving the whole image with a kernel; the kernel is typically a multi-scale construct utilising second-order derivatives (computed using the Hessian matrix). These methods tend to be computationally expensive because they process the whole dataset.

Active contour methods segment vascular structures by propagating a surface or front. The surface is typically embedded in a higher dimensional function such as the zero level-set. Lorigo et al. [21] was one of the first to extend the classical geodesic active contour model to 3-D images for use segmenting vascular structures. In this method a one-dimensional curve (corresponding to the tubular centreline) was evolved in 3-D space. Holtzman et al. [15] formulated a level-set cost function based on three terms: the zero crossings of the second-order derivative, a minimal variance term to penalise lack of homogeneity inside and outside the evolving surface, and a geodesic active surface term to regularise the functional. They showed this approach was suitable for detecting thin vascular structures with low contrast compared to their background. Yan et al. [42] also presented a level-set based approach suitable for segmenting thin vessels. Their method was founded on the principle of capillary action — the attraction of fluid to the walls of a bounding tube. This phenomena was used to derive an adhesion energy term for propagating an active contour. In [9], three separate approaches were brought together: firstly a multi-scale Hessian-based line filter was used for enhancement, then a level-set based approach (which implicitly handled change in topology) was used to provide an initial segmentation, and finally a geometric deformable model (triangulated mesh) driven by a gradient energy cost function was evolved to provide the final result.

Skeletonization converts a binary volume to a discrete centreline — or skeleton — using a method such as distance-ordered homotopic thinning [27]. Pruning and graph analysis techniques must then be applied to transform the unordered discrete set of points into an acyclic graph [14,10]. The challenge for skeletonization is to obtain a good initial segmentation. In [14] a region-growing technique was used to produce the initial binary volume. Such intensity-based techniques are susceptible to noise and varying intensity within the vessel, so therefore in [8] an additional gradient magnitude criteria was added to the traditional lower and upper thresholding strategy. Because these methods only consider pixels comprising the vessel, they are relatively fast; however, they operate in discrete pixel space which can cause the centreline to exhibit stair-case artefacts. In [22] a level-set based method was applied to produce the initial segmentation, followed by a graph analysis method to order and smooth the skeleton.

Similar to skeletonization, *tracking* methods only consider the pixels in close proximity to the vessel centreline, and therefore tend to be relatively fast. Such methods are typically iterative in nature; at each step an operator is applied to compute a continuous point on the centreline and a direction to step. Wink et al. [39] proposed a 'centrelikelihood' operator based on the termination of a number of radially projected lines. Aylward et al. [4] used a Hessian-based metric to compute both the centre point and step vector, as well as estimate the radius. Tracking methods tend to be highly susceptible to noise: once the computed centre point deviates from the actual, it is difficult for the algorithm to recover.

Minimal path techniques frame the centreline extraction problem in terms of cost function minimisation. In [37] a hybrid tracking–path technique was presented. An initial estimate of the centreline was found by tracking in a helical or 'corkscrew' motion. A

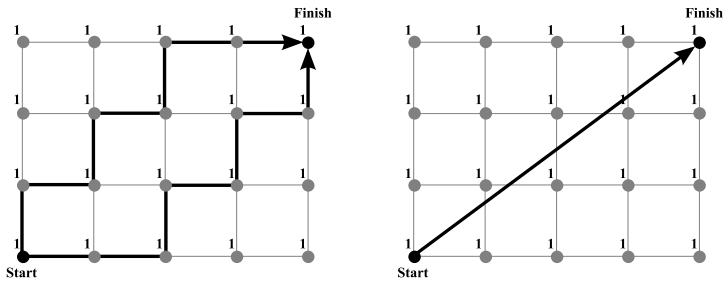

(a) Dijkstra's algorithm (b) Fast Marching method

Fig. 1. Dijkstra's algorithm is restricted to the grid points, whereas the Fast Marching method solves the underlying continuous solution.

cost function — based on a centrelikelihood measure similar to that discussed above in [39] — was then iteratively minimised. Unfortunately, the authors indicate this method was not robust in the presence of noise. Wink et al. [40] explored two best-first minimal path search algorithms: Dijkstra's algorithm and A*. Dijkstra's algorithm operates on a cost function and fans out from the start position, accumulating the cost of each possible discrete path until the end point is reached. The A* algorithm is a generalization which uses a heuristic function $h(\vec{x})$ to better steer the search process (a simple — but common — heuristic function is the Euclidean distance between the current node and the goal). If $h(\vec{x}) = 0$ then A* reduces to Dijkstra's method. The authors compared both unidirectional (expanding from start to end) and bidirectional (expanding from both start and end) searches. As discussed by Sethian [29], Dijkstra's method (and by generalization A*) is inconsistent with the underlying continuous problem: the resultant minimal path is bound to the discrete grid. In contrast, Fast Marching approximates the continuous solution to the underlying partial differential equation (see Fig. 1). For this reason, our proposed method builds upon the Fast Marching minimal path extraction framework set out in [29], which is discussed in greater detail in the following section.

3. Method

Our goal is to enable radiologists and cardiac surgeons to visualise the morphology and degree of stenosis of the coronary arteries from 3-D CTA images. Visualisation is important for diagnosis, treatment planning, training, and patient education. However, this task is not easy for a number of reasons:

(i) CTA images of the thoracic cavity are often large ($\approx 512^3 \times 12$–bits).
(ii) It is common for images to contain noise and artefacts from breathing and other motion, metal streaking, and calcified plaques.
(iii) Both desired structures (coronary arteries) and undesired structures (ventricles, atria, aorta, and pulmonary arteries) contain contrast agent and exhibit similar intensity distributions.
(iv) The coronary veins reside in close proximity to the arteries and in some cases can have higher intensity (particularly the great cardiac vein and coronary sinus).

(v) The contrast agent can vary along the length of the arteries; that is, distal segments have lower intensity than proximal segments.

To address these concerns, a composite method employing segmentation to aid visualisation is proposed. To ensure robustness, a semi-autonomous segmentation methodology — which recognises the user as an important participant of the system capable of correcting problematic regions — is adopted. The path extraction must be fast to allow the user to interact and quickly process large images. Therefore, the proposed method consists of four stages: (1) vessel enhancement for generating an efficient cost function, (2) vessel centreline extraction using Fast Marching minimal path extraction, (3) computation of a measure of stenosis from the resultant centreline, and (4) visualisation of the resultant segmentation and measure. The following sections details each of these steps.

3.1. *Vessel Enhancement*

As will become apparent, it is important to suppress non-vascular structures to ensure fast and robust centreline extraction. In this section two vessel enhancement strategies are considered: morphological top-hat by opening [30], and multi-scale line filtering [28,13]. Readers already familiar with these techniques are invited to skip to the next section.

The relevant morphological operations are defined following the treatise in [30]. The erosion of an image I is denoted by $\varepsilon_B(I)$ and is defined at a given pixel $\vec{x}$ as the minimum value of the image in the window defined by the structuring element B when its origin is at $\vec{x}$:

$$[\varepsilon_B(I)](\vec{x}) = \min_{b \in B} I(\vec{x} + b) \tag{1}$$

Similarly the dilation of an image is defined as the maximum value of the image in the window defined by the structuring element:

$$[\delta B(I)](\vec{x}) = \max_{b \in B} I(\vec{x} + b) \tag{2}$$

The opening of an image I by a structuring element B is denoted by $\gamma_B(I)$ and is defined as the erosion of I by B, followed by the dilation by B:

$$\gamma_B(I) = \delta_B[\varepsilon_B(I)] \tag{3}$$

Top-hats use knowledge about the shape characteristics that are *not* shared by the relevant image structures. Top-hat by opening (otherwise known as white top-hat or simply WTH) is the difference between the original image I and its opening γ:

$$\mathrm{WTH}_B(I) = I - \gamma_B(I) \tag{4}$$

The top-hat by opening with a ball structuring element suppresses structures which do not fully contain the element (ie. large non-tubular structures are removed).

The computation time for morphological filtering rises exponentially with increasing structuring element radius. Therefore, a number of methods have been proposed to accelerate morphological methods to overcome this hurdle. The approach presented by van Herk [34] in the early 1990's is still one of the most efficient. This approach considers each dimension separately and uses a recursive procedure to track the extrema, resulting

in an operation independent of the structuring element size, but only supporting box structuring elements. More recent techniques which support arbitrarily shaped elements include moving histogram [33] and anchor [32] methods.

A more sophisticated approach for enhancing tubular structures is to take advantage of second-order local structure. There exist a number of vesselness measures which utilise the eigenvalues (λ_1, λ_2, λ_3) of the Hessian matrix $H = \nabla^2 I$. Two noteworthy methods include the line filter proposed by Sato et al. [28], and then further refined by Frangi et al. [13]. Both of these methods firstly compute the Hessian matrix using convolution with a derivative of Gaussians. Given H, the eigenvalues are then ordered: Sato's approach sort based on value (ie. $\lambda_1 > \lambda_2 > \lambda_3$), whereas Frangi's method perform the ordering using *absolute* value (ie. $|\lambda_1| > |\lambda_2| > |\lambda_3|$). Different ratios of the eigenvalues are then used to construct the respective vesselness measures.

Sato's method analysed the ideal case and asserted that a bright line exhibits the following relations: $\lambda_1 \approx 0$, $\lambda_2 \approx \lambda_3 << 0$. With these in mind the line filter was suggested as follows:

$$V_S(I,\sigma) = \begin{cases} \exp\left(-\dfrac{\lambda_1^2}{2(\alpha_1\lambda_c)^2}\right); & \lambda_1 \geq 0, \lambda_c \neq 0 \\ \exp\left(-\dfrac{\lambda_1^2}{2(\alpha_2\lambda_c)^2}\right); & \lambda_1 > 0, \lambda_c \neq 0 \\ 0; & \lambda_c = 0 \end{cases} \tag{5}$$

where $\lambda_c = \min(-\lambda_2, -\lambda_3)$, and α_1 and α_2 are control parameters. Frangi's method asserted that the following relations hold for an ideal line: $\lambda_1 \approx 0$, $\lambda_1 \leq \lambda_2$, $\lambda_2 \approx \lambda_3$. The proposed line filter consequently used of all three of the eigenvalues:

$$V_F(I,\sigma) = \begin{cases} \left(1 - \exp\left(-\dfrac{R_A^2}{2\alpha^2}\right)\right) \exp\left(-\dfrac{R_B^2}{2\beta^2}\right) \left(1 - \exp\left(-\dfrac{|H|_F^2}{2\gamma^2}\right)\right) \\ 0; \qquad \lambda_2 > 0 \text{ or } \lambda_3 > 0 \end{cases} \tag{6}$$

where α, β, and γ are control parameters, $R_A = |\frac{\lambda_2}{\lambda_3}|$, $R_B = \frac{|\lambda_1|}{\sqrt{\lambda_2\lambda_3}}$, and $|H|_F$ is the Frobenius norm of the Hessian matrix. Both methods detect tubular structures of varying radius using a multi-scale approach:

$$V(I, \sigma_{min}, \sigma_{max}) = \max_{\sigma_{min} \leq \sigma \leq \sigma_{max}} V(I,\sigma) \tag{7}$$

The unavoidable computational burden with such line filtering methods resides in the construction of the Hessian matrix H for the entire image by convolution with potentially large kernels. Fig. 2 depicts results from each of the discussed measures. Notice that the line filtering methods suppress more of the non-tubular structures than morphological top-hat by opening.

3.2. *Centreline Extraction*

The ability to quickly and accurately extract vessel centrelines from large and noisy images is difficult. This section proposes techniques to achieve this task for contrast-enhanced CT images of the coronary arteries, however other modalities and vasculature

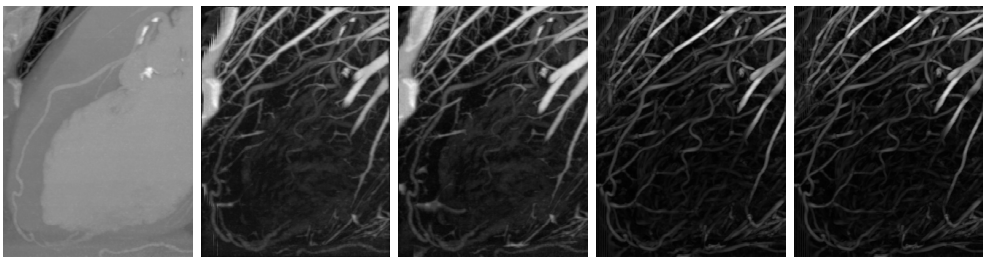

(a) Original (b) Ball WTH (c) Box WTH (d) Sato Line (e) Frangi Line

Fig. 2. A maximum intensity projection (MIP) of the dataset used in Section 4.1 for each of the different vessel enhancement measures: white top-hat (WTH) and multi-scale line filtering.

are inherently supported. Firstly an overview of conventional Fast Marching minimal path extraction techniques and limitations is given. Three improvements are then presented: (1) an efficient cost function is constructed using vessel enhancing filters, (2) a mechanism is presented for correctly handling calcification artefacts in CTA images, and (3) regularisation is introduced to control the step vector length for gradient descent back-propagation. Finally, a proposal is made for the incorporation of way-points into the Fast Marching minimal path framework, which provides the user with a means to guide the process.

3.2.1. *Conventional Method*

Fast Marching is typically known as an active contour method, however it is also suitable for a range of optimisation problems including minimal path extraction (see [29, pp. 284–312]). Active contour methods model the motion of a contour normal to itself with a known speed function F. A contour is a curve (2-D) or surface (3-D) separating two regions (inside, outside). The speed function F may depend on various factors, however this paper deals exclusively with fronts that propagate *outwards* (ie. $F \geq 0$). Sethian [29] describes one way to characterise such an expanding front: each grid position is assigned an arrival time T at which the front crosses it. Fig. 3 illustrates the arrival function for an arbitrary curve travelling with speed $F = 1$. The arrival function $T(\vec{x})$ can be computed for n-dimensional images using the fact that *distance* = *speed* × *time*:

$$1 = F \, \|\nabla T\| \quad \text{such that} \quad \Gamma = \{(\vec{x}) \mid T(\vec{x}) = 0\} \tag{8}$$

where ∇ is the grad function orthogonal to the level sets of T, and Γ is the initial location of the front. If the speed function F only depends on position, then the equation reduces to the boundary value problem known as the *Eikonal equation*: a non-linear partial differential equation.

In practice the Eikonal equation is solved using numerical methods. Sethain [29, pp. 68] reports a backward finite difference approach:

$$\begin{bmatrix} \max(D^{-x}T, -D^{+x}T, 0)^2 \\ + \max(D^{-y}T, -D^{+y}T, 0)^2 \\ + \max(D^{-z}T, -D^{+z}T, 0)^2 \end{bmatrix}^{1/2} = \frac{1}{F} \tag{9}$$

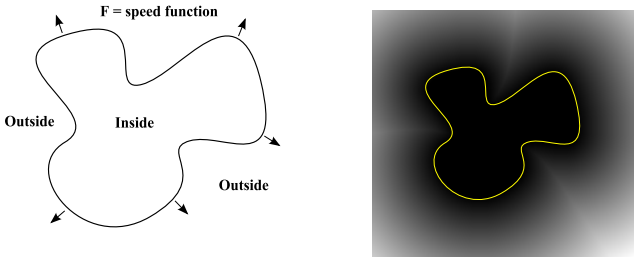

(a) Illustration (b) Arrival Function

Fig. 3. Active contour methods model curves which propagate with speed F normal to their surface. The arrival function on the right was computed with $F = 1$.

```
Definition:
  Alive is the set of grid points with known T values.
  Trial is the set of grid points with recently computed T values.
  Far is the set of remaining grid points with no computed T value.

Loop:
  Let A be the Trial point with the smallest T value.
  Add A to Alive; remove it from Trial.
  For all neighbours of A:
    If neighbour is Far, remove from Far and add to Trial.
  Recompute T for all Trial neighbours of A according to Eq. 9
End
```

Listing 1. Fast Marching pseudo-code

where D is the spatial derivative operator:

$$D^{+x}u \equiv \frac{u(x+h,t) - u(x,t)}{h} \tag{10}$$

$$D^{-x}u \equiv \frac{u(x,t) - u(x-h,t)}{h} \tag{11}$$

The problem now reduces to solving the quadratic equation given in Eq. 9 in an efficient manner. Fast Marching is the name given to one such solution which propagates information from smaller values of T to larger values. This is typically achieved using a heap sort data structure to iteratively compute T for the neighbours of the smallest known value. The pseudo-code for the method is listed in Listing 1.

In simple terms the Fast Marching method introduces an optimal ordering to all grid points. The order is propagated outward from the set of *Trial* points and the order can only grow (due to Eq. 9). When using a heap sort data structure for the set of *Trial* points, the method requires only a single pass over the image. The worst-case scenario for updating one point in the heap (ie. moving a point from the bottom to the top) results in $O(\log N)$ complexity. Therefore, for N grid points the Fast Marching method has $O(N \log N)$ complexity.

Given a cost function F and an initial front consisting of a single point p_0, the resultant arrival function $T(p)$ contains the arrival time t for each point. The minimal path P between p_0 and p_1 can be found by back-propagating from $T(p_1)$ to $T(p_0)$ (remember $T(p_0) = 0.0$). Because the front moves normal to the initial point, it can be shown that

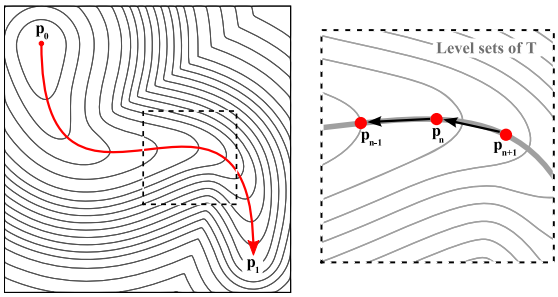

Fig. 4. The minimal path is extracted from the arrival function T using back-propagation. The extracted path is comprised of a set of line segments computed by stepping along the gradient, which is orthogonal to the level sets of T.

tracking orthogonal to the level sets of $T(p)$ from p_1 will always converge to p_0 along the minimal geodesic path (see [18]). In practice this continuous path is approximated using numerical methods. Sethian originally proposed a second-order scheme (Heun's method) to compute the step vector [29, pp. 285–286]. In [11] a simple first-order gradient descent method was used:

$$p_{n-1} = p_n + \alpha \, \nabla T(p_n) \tag{12}$$

where p_n is the current position, p_{n-1} the previous position, and α is the step length (see Fig. 4).

Deschamps [11] also proposed a cost function for Fast Marching minimal path extraction based directly on the input image intensity:

$$F = \text{Rescale}\,(I) \tag{13}$$

where I is the input image, Rescale is a function which rescales the intensity to the range $[0, 1]$, and F is the resultant cost function. This cost function allows the front to expand to unwanted portions of the image, which is computationally expensive. Deschamps proposed two possible mechanisms for dealing with this issue: (1) alter the Fast Marching algorithm to freeze *Trial* points which have reached the boundary of the object, and/or (2) stop the front propagating once the end point of the path has been reached (which he termed *partial propagation*).

3.2.2. *Improved Method*

The conventional method described above suffers a number of drawbacks: (1) the expanding front visits unnecessary pixels which increases computation time and potentially creates undesirable 'short-cuts', (2) the path is incorrectly attracted to high-intensity calcification artefacts, and (3) gradient descent back-propagation is sensitive to oscillation. This section proposes various improvements to address each of these concerns.

The value of the cost function outside of vascular structures (denoted δ) influences both the correctness and efficiency of the minimal path extraction. Large δ values cause the front to visit unnecessary pixels which is computationally expensive. Also, if δ is too large, the back-propagation may find an undesirable 'short-cut' between p_1 and p_0. Fig. 5 illustrates the effect of different values of δ. Ideally the value of the cost function outside of vascular structures should be as close to zero as possible. With this observation in mind,

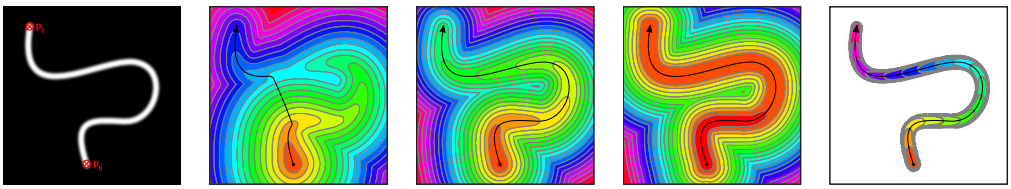

(a) Input (b) $\delta = 0.25$ (c) $\delta = 0.1$ (d) $\delta = 0.001$ (e) $\delta = 0.0$

Fig. 5. The value of the cost function outside of the vessel (δ) plays a significant role in both the correctness and efficiency of Fast Marching minimal path extraction. Cost functions in which δ is small ensure correct results and are computationally efficient (only visiting pixels near the desired centreline).

it is possible to construct a simple cost function which prevents the front propagating to unnecessary pixels:

$$\widetilde{I} = \begin{cases} \delta; & \text{Enhance}(I) \leq t_1 \\ \text{Enhance}(I); & \text{otherwise} \end{cases} \tag{14}$$

$$F = \text{Rescale}\Big(\text{Smooth}\big(\widetilde{I}\big)\Big) \tag{15}$$

where I is the input image, Enhance is a vessel enhancement strategy, Smooth is a suitable smoothing function, δ is a small value which impedes the motion of the front outside of vessels (typically $\delta = 0.0$), and t_1 is the lower threshold value (typically $t_1 \approx 100$). As shown later, the white top-hat (WTH_B) with box structuring element of length $2B$ is a suitably fast vessel enhancement strategy. The smoothing function is required to suppress discontinuities in the cost function due to the thresholding. Such discontinuities can skew or prematurely halt the gradient descent back-propagation. Furthermore, the smoothing helps exaggerate the Gaussian profile of the vessels, which is important for extracting a centred path. The curvature flow denoising method described in [29, pp. 200–213] is an appropriate — yet relatively fast — smoothing technique, controlled using two parameters: time step (t_{cvf}) and number of iterations (n_{cvf}).

In some cases high-intensity calcified plaque accumulates on the interior surface of arteries, resulting in narrowing. Plaque blocks the passage of blood, and as such, should be excluded when computing the vessel centreline. As depicted in Fig. 6(a), the currently proposed cost function causes the path to be *attracted* to the high-intensity calcifications, rather than repelled. Another threshold must be added to the cost function to suppress this unwanted attraction:

$$\widetilde{I} = \begin{cases} \delta; & \text{Enhance}(I) \leq t_1 \\ \delta; & \text{Enhance}(I) \geq t_2 \\ \text{Enhance}(I); & \text{otherwise} \end{cases} \tag{16}$$

$$F = \text{Rescale}\Big(\text{Smooth}\big(\widetilde{I}\big)\Big) \tag{17}$$

where t_2 is the additional calcification threshold value (typically $t_2 \approx 700$). Fig. 6(b) depicts the corrected path.

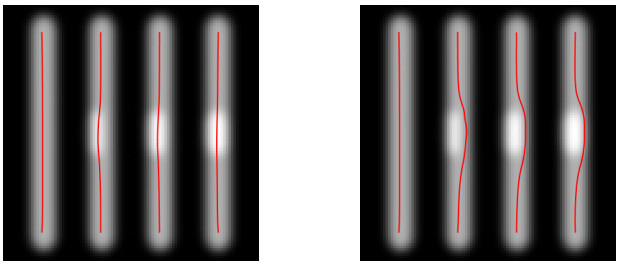

(a) Uncorrected (b) Corrected

Fig. 6. An upper threshold (t_2) must be added to the cost function to ensure the path correctly negotiates high-intensity calcifications. From left to right, the synthetic vessels represent the following degrees of stenosis (S_{deg}): 0%, 25%, 50%, and 75%. The radius of the calcified plaque (r_c) was computed for fixed vessel radius ($r_v = 4.75$mm) using the following relation: $r_c = r_v\sqrt{S_{deg}}$.

Minimal path back-propagation using fixed length gradient descent is susceptible to unwanted oscillation (see [11, pp. 19]). To prevent this the step length can be regularised as follows:

$$\alpha = \frac{\alpha_0}{\|\nabla T(p_n)\|} \tag{18}$$

where α_0 is the maximum desired step length. Using this approach, the step length will vary based on the degree of change in the arrival function (areas of large change will result in smaller steps, more closely approximating the underlying continuous function). The value of α_0 allows the user make a trade-off between efficiency and accuracy. Empirical evidence suggests that $\alpha_0 = \min(\text{spacing})$ is a suitable choice (where spacing is the set of physical distances between pixels for each dimension).

3.2.3. *Way-point Guidance*

Despite the addition of vessel enhancement to the cost function, it is still possible for the back-propagation to find undesirable 'short-cuts' (see the first column of Fig. 7 for one such example). To address this issue, Wink et al. [41] proposed an automated method which searches for the optimal path through a multi-scale vesselness cost function. This technique requires the pre-computation of the vesselness measure proposed by Frangi et al. [13] which, as already discussed, is computationally expensive for large 3-D images. Instead a semi-automated approach is adopted, similar to live-wire [6], in which the user can specify a number of way-points which the path must pass near.

The proposed way-point guidance method is a generalisation of the partial-propagation method proposed by Deschamps [11]. The user must provide a cost function and list of paths; each path consists of start and end points, and any number of way-points (if no way-points are specified the method reduces to the original). Starting with the first way-point (or end point if there are no way-points) a front is propagated in the typical Fast Marching manner and terminates when both the previous and next points in the path have been reached. A path segment is extracted by back-propagation from the previous point to the current *Trial* point. The process is repeated with each remaining way-point and end point acting as the initial front. The resultant path is computed by concatenating

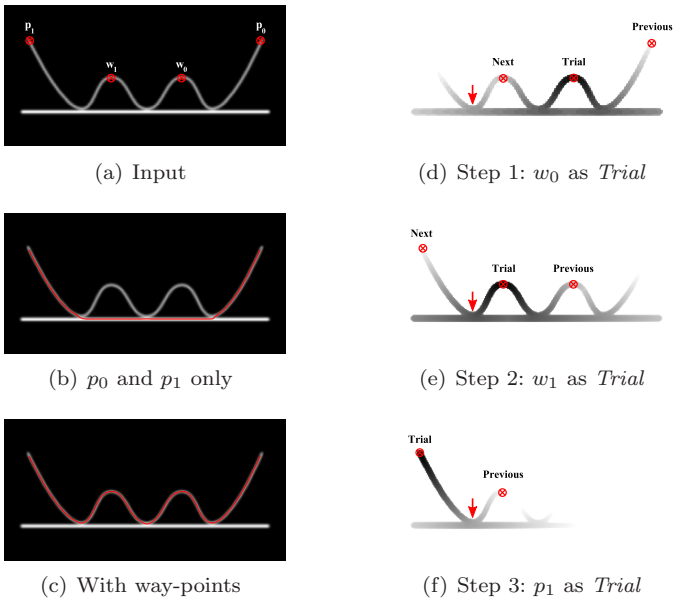

(a) Input

(b) p_0 and p_1 only

(c) With way-points

(d) Step 1: w_0 as *Trial*

(e) Step 2: w_1 as *Trial*

(f) Step 3: p_1 as *Trial*

Fig. 7. This figure depicts the proposed way-point guidance method. The first column (a–c) demonstrates the minimal path is not always the desired path. The second column (d–f) displays the steps for extracting the path. For each arrival function the path segments are extracted and combined. The arrow indicates the unfortunate redundancy associated with the method.

each path segment. The second column in Fig. 7 depicts the arrival function for each step of the process applied to a 2-D synthetic image.

Unfortunately, as indicated by the arrow in Fig. 7, some grid positions are visited by the front on multiple occasions, which is computational expensive. The worst case complexity is $O(M \times N \log N)$ (where M is the number of way-points and N the number of pixels in the image). However, as will be shown in a succeeding section, this worst case scenario is rarely encountered; in fact the way-points can actually help *reduce* the computation time. Adding a heuristic function (such as those in [2] or [40]) could also further improve the efficiency.

3.3. *Measuring Vessel Lumen*

It is important for clinicians to have tools to visually quantify the prevalence of disease. For coronary heart disease this is achieved through a measure of stenosis, typically derived using an estimate of the vessel radius or cross-sectional area. Fig. 8 shows that these measures are sensitive at bifurcations and tend to inflate the measure (see [5]). This section proposes a simple measure based on the smallest Euclidean distance from the centreline to the vessel wall; the *minimum* distance ensures the radius estimate is not over-inflated near branch points (notice in Fig. 8 that $r_2 < r_1$).

One of the strengths of the Fast Marching minimal path method is that the centreline is computed *without* requiring the vessel surface (compare with skeletonization methods which use a binary segmented volume to compute the centreline). However, given the centreline, it is straight-forward to compute the binary volume representing the vascu-

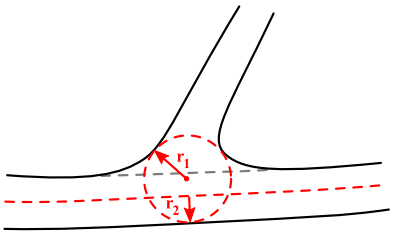

Fig. 8. The estimate of the vessel radius at bifurcations is subject to inflation (r_1). Choosing the minimum distance to the vessel wall on the centreline avoids this over estimation (r_2).

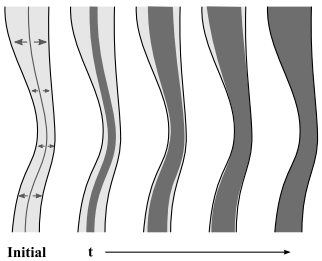

Fig. 9. The vessel surface can be segmented using Fast Marching. The cost function given in Section 3.2.2 is used as the speed image and the discretized centreline as *Trial* positions. The front will expand perpendicular to the centreline and slow down as it reaches the vessel boundary.

lature boundary. The discretized centreline acts as an initial front which is propagated outward using Fast Marching (Fig. 9 depicts the process at increasing time steps). Traditional Fast Marching — with a user initialised front — is not suitable for segmenting vessels because the front inevitably leaks (remember there are no curvature or homogeneity regularisation terms). However, because the centreline is a good initial estimate of the vessel surface, the marching can be terminated after a small number of iterations (typically $t \leq 10$), which not only prevents leakage but is also very fast. The cost function given in Section 3.2.2 (without the smoothing term) is a suitable speed image.

The minimum Euclidean distance to the surface can be efficiently computed using the Maurer distance transform [24]. This value can then be interpolated at each point along the centreline and stored for later retrieval. Because vascular structures represent only a small portion of the whole image, it is advantageous to compute a low-resolution mask (on a subsampled image) to indicate which pixels are near the centreline. Each pixel under the mask can then perform a brute-force search to find the closest centreline point and stored radius estimate. This masked brute-force method may not be as elegant as a space partition approach (such as kd-trees), however it is adequate for our situation.

3.4. *Visualisation*

The extracted vessel volume and measure of vessel lumen create a number of opportunities for improved visualisation. In this section two visualisation techniques are proposed which address limitations with existing methods. The first method extends maximum intensity projection (MIP) to allow the coronary arteries to always remain visible. The second method uses direct volume rendering (DVR) to pseudo-colour the arteries with the measure of vessel lumen.

3.4.1. *Weighted Maximum Intensity Projection*

Maximum intensity projection (MIP) is a favoured technique amongst radiologists because it allows for three dimensional visualisation of the whole dataset, without the need for adjusting display parameters (which can lead to large inter-observer variability). This advantage comes at the cost of discarded 3-D information. A common issue for

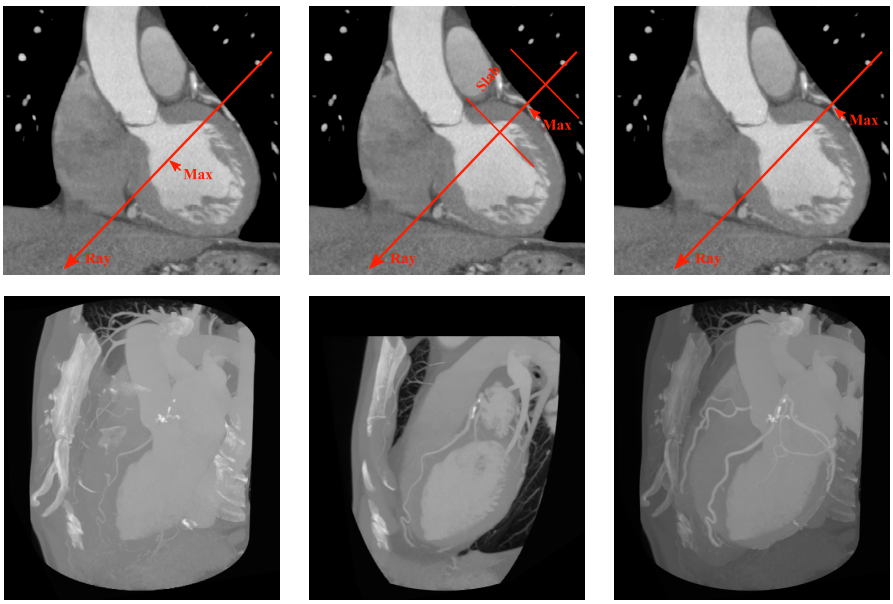

(a) MIP (b) Thin-slab MIP (c) Weighted MIP

Fig. 10. The coronary arteries are obscured when using conventional maximum intensity projection (MIP) to visualise cardiac angiography data. Thin-slab MIP alleviates this issue at the cost of decreased usability and removal of contextual cues. Weighted MIP enables the coronary arteries to be propagated to the resultant rendering while still maintaining contextual cues ($w_{out} = 0.8, w_{in} = 1.0$).

cardiac angiography is that the left ventricle, left atrium, and aorta (all of which contain high concentrations of contrast) often obscure the vasculature in normal MIP images. Thin-slab MIP addresses this concern by projecting the maximum value along rays which are restricted to a thin slice [12]. Despite the visualisation improvement, the usability of thin-slab MIP is tedious, the user is required to scroll through the data similar to many 2-D visualisation methods. Furthermore, the resultant rendering only shows pixels within the slab, removing important contextual cues.

Given the segmented volume representing the coronary arteries, it is possible to weight the vessels to prevent them being obscured by other higher intensity structures. Values outside of the arteries are assigned weight w_{out} and values inside the arteries weight w_{in}, where $w_{out} < w_{in}$. If w_{out} is sufficiently less than w_{in}, then the arteries will be propagated as the maximum value, while still retaining contextual structures (with reduced intensity). We call this method weighted maximum intensity projection. Fig. 10 depicts the differences between traditional MIP, thin-slab MIP, and weighted MIP (for this example, it was empirically found that $w_{out} = 0.8$ and $w_{in} = 1.0$ produced suitable results). Weighted MIP has the advantage of displaying the entire coronary network, with contextual cues, using relatively simple display parameters. However, the method still suffers from loss of dimensionality, the fundamental drawback associated with all MIP techniques.

3.4.2. *Pseudo-colouring the Vessel Lumen*

In contrast to MIP, direct volume rendering uses a number of perceptual cues to represent three dimensional structures [12]. The user is required to specify various display parameters (opacity and colour), typically in an iterative and interactive fashion via a 'transfer function'. In some cases it may not be possible to specify parameters which clearly depict the structures of interest. In these cases tagged volume rendering [25] can be employed. Tagged volume rendering uses a combination of segmentation before and during rendering to separate structures of interest with similar intensity characteristics.

A similar mechanism can be employed which, in addition to delineating structures of interest, applies pseudo-colouring to the arteries to quantitatively depict the vessel lumen. The method can be described as follows: cast a number of rays into the volume and sample values at equidistant positions along each ray; query the label of each sample to determine if it lies outside the heart, inside the heart, or inside an artery; discard samples outside of the heart; apply typical volume rendering for samples inside the heart; use a lookup table indexed using the lumen measure to pseudo-colour samples inside the arteries; combine the samples along each ray to form the final output image. Diseased coronary artery segments can be visually identified in the resultant renderings as bands, and the colour map can be used to quantify the degree of stenosis. Fig. 14 depicts visualisations generated using this method.

4. Results

Because the proposed method comprises both segmentation and visualisation, validation is a difficult task. The efficiency, reproducibility, and accuracy of the vessel segmentation is firstly considered using a mixture of both real and synthetic datasets. The two proposed visualisation methods were then applied to a small cohort of clinically obtained datasets, the results of which are depicted in Fig. 14. The following tests were executed on a high performance desktop computer (Intel Xeon, 4×2.66 GHz processors, 8 GB RAM, Windows XP 64-bit) with a recent graphics card (NVIDIA Quadro 4600, 768 MB VRAM). The algorithms were implemented in C++ using the Insight Toolkit (ITK) [1] . The volume visualisation methods were realised using a custom made OpenGL application with fragment shaders written in OpenGL Shading Language (GLSL).

4.1. *Efficiency*

In terms of efficiency, there are three interesting questions requiring investigation: (1) Which vessel enhancement method performs the fastest? (2) Does adding way-points increase the computation time? and (3) How long does the proposed method take?

Firstly, the question of vessel enhancement speed was considered. All of the methods discussed in Section 3.1 are suitable for the task (the line filtering methods have particularly desirable characteristics), however practicality is an important consideration. A simple experiment was devised to compare the computation times for five different vessel enhancement methods: (1) naive top-hat by opening with ball structuring element (Ball WTH), (2) histogram white top-hat with ball structuring element (Histogram WTH), (3)

[1] Insight Toolkit(ITK): `http://www.itk.org`

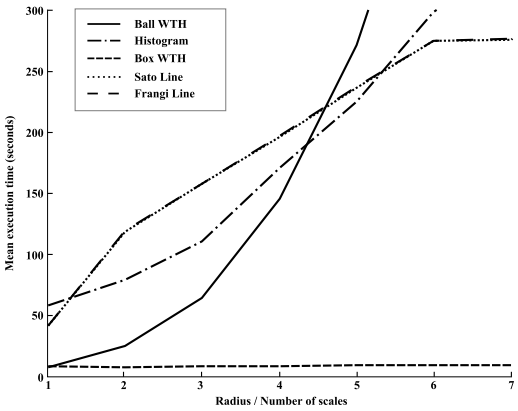

Fig. 11. Experimental results comparing computation time of five different vessel enhancement strategies: three morphological white top-hat (WTH) methods and two Hessian-based line filters.

van Herk white top-hat (Box WTH), (4) multi-scale line filtering using Sato's method (Sato Line) with $\alpha_1 = 0.5$, $\alpha_2 = 0.5$, $\sigma = [0.1, 1.0]$, and (5) multi-scale line filtering using Frangi's method (Frangi Line) with $\alpha = 0.5$, $\beta = 0.5$, $\gamma = 5.0$, $\sigma = [0.1, 1.0]$. The implementations detailed in [7] and [3] were used for this experiment. Each method was executed in a sequential fashion on a single thread a total of five times, the mean time was recorded. The dataset was a clinically obtained CTA dataset of the thoracic cavity, cropped to include only the left anterior descending (LAD) (size [220, 260, 300] pixels, spacing [0.35, 0.35, 0.4] mm). As can be seen in Fig. 11, the van Herk morphological top-hat by opening (Box WTH) performed orders of magnitude faster than the other methods (dashed line). Furthermore, the results confirm that the computation time is independent of the structuring element size.

Next, the efficiency of the way-point guidance method was examined (note that the reproducibility is addressed separately in the next section). The centreline of the LAD from the same dataset as above was extracted using the proposed Fast Marching minimal path method with varying number of way-points. The following parameters were employed: $t_1 = 200$, $t_2 = 800$, $\delta = 0.0$, $B = 8$, $t_{cvf} = 0.15$, $n_{cvf} = 4$, $\alpha_0 = 0.35$. The synthetic example given in Section 3.2.3 indicated that adding way-points could potentially degrade the path extraction performance. However, the results listed in Table 1 indicate that for real datasets the reverse is true: additional way-points actually help focus the expanding front, reducing computation time.

Finally, the overall time taken for each step of the proposed method was evaluated using a small cohort of clinically obtained datasets. Each dataset was acquired in a hospital setting using contrast-enhanced, ECG-gated, 40-slice spiral computed tomography angiography (Brilliance 40, Philips Medical Systems, Best, Netherlands), with approximate size [512, 512, 512] pixels and spacing [0.4, 0.4, 0.4] mm. The results listed in Table 2 show that — once the path information has been specified — the centreline, surface, and measure of vessel lumen for the entire coronary network can be automatically computed in approximately 6 minutes. The longest phase was for the user to specify the start, way, and end points for the path extraction; future work could focus on methods to help automate this stage.

Way-points	Time (s)	DSC
1	3.55	0.946
2	3.93	0.898
3	2.12	0.930
4	1.87	0.923
5	1.44	0.896

Table 1
Execution time and Dice similarity coefficients (DSC) for varying number of way-points. The values indicate that increasing the number of way-points does not increase the execution time and produces similar resultant paths.

Dataset	Size	Paths	Cost	Centreline	Surface	Measure
A	[512, 512, 428]	10m 30s	50.4s	71.5s	53.0s	171.9s
B	[512, 512, 404]	10m 40s	40.0s	54.7s	65.2s	158.5s
C	[512, 512, 334]	9m 45s	58.2s	56.4s	41.7s	129.0s
D	[512, 512, 377]	14m 0s	64.7s	92.9s	53.3s	196.1s
E	[512, 512, 400]	12m 40s	66.6s	86.8s	69.4s	158.9s

Table 2
The processing time for each stage of the proposed method on a small cohort of clinically obtained datasets. Each dataset had approximate size [512, 512, 400] pixels and spacing [0.4, 0.4, 0.4] mm.

4.2. *Reproducibility*

It is desirable for the method to be invariant to user specified parameters. Dice similarity coefficients (DSC) have been shown to be a useful validation metric to evaluate reproducibility [43]. The DSC metric is defined as follows:

$$DSC = \frac{2 \times (A \cap B)}{(A + B)} \tag{19}$$

where A and B are target regions, $A \cap B$ is the intersection of the regions, and $A + B$ the sum of the regions. To assess the reproducibility, dataset A had a start point (p_0) placed in the distal portion of the LAD and an end point (p_1) at the origin of artery (near the aorta). A number of way-points were spaced evenly along the artery ($n = [1, 2, 3, 4, 5]$) and the DSC computed between the resultant path and the path without way-points. Table 1 lists the resultant values. Given that a good overlap is considered to occur when $DSC > 0.7$, these results indicate the placement of superfluous way-points does not significantly alter the extracted centreline.

4.3. *Accuracy*

To access the accuracy of the centreline extraction method the Simultaneous Truth And Performance Level Estimation (STAPLE) method proposed by Warfield et al. [36] was employed. The STAPLE algorithm takes a set of binary segmentations and generates an assessment of the accuracy of each input using expectation-maximisation to characterise the performance in terms of sensitivity (true positive fraction) and specificity (true

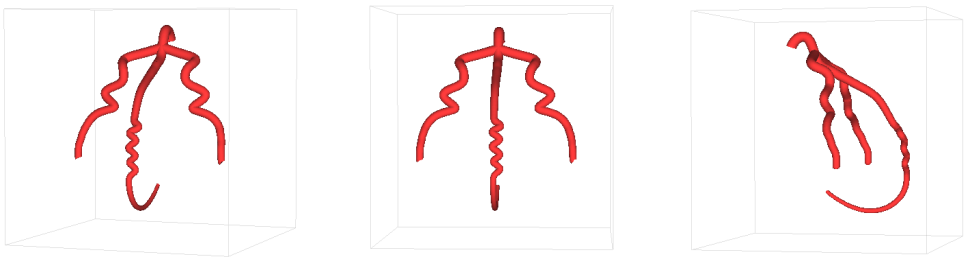

Fig. 12. A surface rendering of the synthetic dataset using the OpenGL Tubing and Extrusion API (GLE). The tubular object is tortuous, contains a branch, and exhibits varying radius.

negative fraction). Because the algorithm expects the set of binary segmentations to overlap, an iterative scheme must be used which morphologically dilates the centreline with increasing radius [16]. This allows the given centreline segmentations to be compared for each dilation factor in terms of sensitivity and specificity.

To test the algorithm in the presence of noise, a synthetic dataset with known characteristics was created. The dataset contained a tortuous, branching, tubular object, with size [256, 256, 256] pixels and spacing [1.0, 1.0, 1.0] mm, and varying radius in the range 3.0 mm to 1.0 mm. The background intensity was set to 0 and the object intensity to 500 (typical CTA values). The dataset was smoothed using a Gaussian kernel to represent the parabolic cross-sectional intensity of real vessels ($\sigma_k = 1.5$). Fig. 12 depicts a surface rendering of this dataset. Five centrelines extracted from the synthetic dataset were fed into the STAPLE algorithm: the discretized centreline used to generate the synthetic dataset (Known), the skeleton obtained using distance-ordered homotopic thinning [27] (Skeleton), and three segmentations using the Fast Marching minimal path technique with additive Gaussian noise of varying standard deviation $\sigma_n = [0, 10, 20]$ (Path). To emphasise the robustness of the method, the Smooth function was omitted from the cost function generation. The algorithm was configured with the following parameters: $t_1 = 100$, $t_2 = \text{N/A}$, $\delta = 0.05$, and $\alpha_0 = 0.75$. Fig. 13 shows that the Fast Marching path extraction method performs well in comparison and the sensitivity is relatively constant for varying levels of noise. It should be noted that no useful comparison could be made using the specificity, and as such these values were omitted from the analysis.

5. Summary

A semi-autonomous composite segmentation and visualisation system for the coronary arteries in 3-D CTA images was presented. The proposed method consisted of four stages: (1) vessel enhancement to construct an efficient cost function, (2) coronary artery centreline extraction using Fast Marching minimal paths, (3) computation of a measure of vessel lumen (which directly relates to the degree of stenosis) using the minimum Euclidean distance from the centreline to the vessel surface, and (4) visualisation employing the resultant vessel segmentation and measure.

It was noted that the conventional Fast Marching minimal path method has the tendency to visit unnecessary pixels (increasing computation time and the risk of undesirable 'short-cuts'). A novel cost function which utilised a vessel enhancement strategy

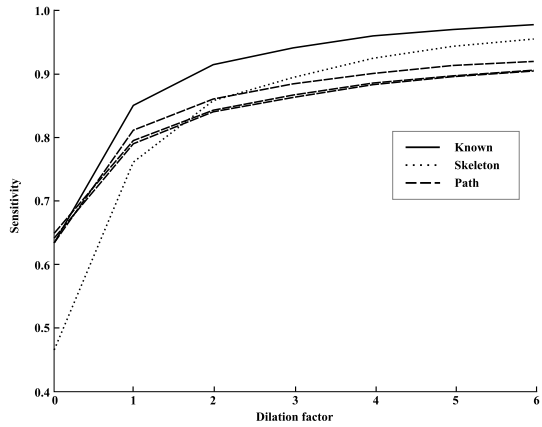

Fig. 13. This figure depicts the sensitivity comparison for the known, skeleton, and path techniques. The Fast Marching path extraction method performs well in comparison and is relatively constant for varying levels of noise ($\sigma_n = [0, 10, 20]$).

was proposed to address this issue. An analysis of existing vessel enhancement methods showed that Hessian-based line filters, while having desirable properties, are computationally expensive. The morphological top-hat by opening was in this case a suitable replacement, and — when using an efficient implementation — can be very fast and independent of structuring element size. Furthermore, a mechanism was added to the cost function to correctly handle calcification artefacts. The step length for gradient descent back-propagation was also regularised to help avoid oscillations in the extracted path.

Despite the improvements offered by the proposed cost function, it was still possible for the back-propagation to find undesirable 'short-cuts'. A semi-automated approach was adopted in which the user could specify a number of way-points which the path must pass near. It was hypothesised that additional way-points could degrade the path extraction performance, however results indicated that for real datasets the reverse was true: additional way-points actually helped focus the expanding front, reducing computation time. Furthermore, Dice similarity coefficients (DSC) indicated that additional way-points did not significantly alter the reproducibility of the extracted centreline.

To provide clinicians with a tool to visually quantify the prevalence of disease, the lumen of the arteries (which indicates the degree of narrowing) was measured. The measure was then mapped to volume renderings using pseudo-colour. Severely diseased coronary artery segments could be visually identified in the resultant renderings. The vessel surface was also used to weight maximum intensity projection (MIP) images, allowing the arteries to remain visible, even when in alignment with higher intensity structures.

Using a proven methodology (STAPLE), the centreline extraction method was compared against known synthetic results, demonstrating its robustness in the presence of noise. Furthermore, the overall timings showed that — once the path information has been specified — the centreline, surface, and measure of vessel lumen for the entire coronary network can be automatically computed in approximately 6 minutes for large ($\approx 512^3 \times 12$–bits) CTA datasets.

The proposed method was applied to the problem of accessing the degree of stenosis of the coronary arteries, which is vital for successful diagnosis and treatment planning

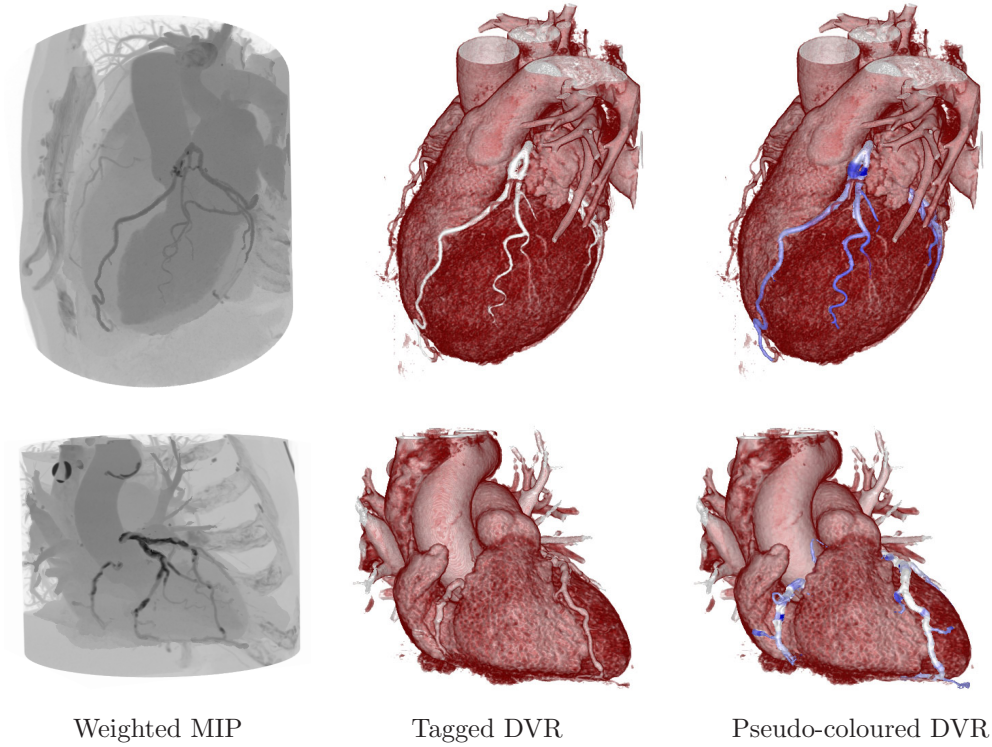

Fig. 14. Resultant direct volume renderings (DVR) with pseudo-coloured coronary arteries; white indicates largest radius, blue smallest. Arterial segments exhibiting stenosis can be easily identified as bands in the pseudo-coloured renderings (last column).

of heart disease. Aspects of the presented system are also applicable to other problem domains, including segmentation and visualisation of other modalities and vascular structures, optimal path planning for image guided surgery, and robot navigation.

Acknowledgements

Thanks to Dr. Richard Slaughter (The Prince Charles Hospital, Brisbane, Australia) for supplying the CTA datasets and offering clinical advice. Thanks also to Prof. Peter O'Shea (Queensland University of Technology, Brisbane, Australia) for commenting on an early draft of the manuscript.

References

[1] S. Achenbach. Cardiac CT: State of the art for the detection of coronary arterial stenosis. *Journal of Cardiovascular Computed Tomography*, 1(1):3–20, 2007.

[2] J. Andrews and J. Sethian. Fast marching methods for the continuous traveling salesman problem. *Proceedings of the National Academy of Sciences (PNAS)*, 104(4):1118–1123, 2007.

[3] L. Antiga. Generalizing vesselness with respect to dimensionality and shape. *Insight Journal*, Aug:1–14, 2007.

[4] S. Aylward and E. Bullitt. Initialization, noise, singularities, and scale in height ridge traversal for tubular object centerline extraction. *IEEE Transactions on Medical Imaging*, 21(2):61–75, 2002.

[5] S. Aylward, E. Bullitt, S. Pizer, and D. Eberly. Intensity ridge and widths for tubular object segmentation and description. In *Proceedings of the Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 131–138. IEEE, 1996.

[6] W. Barrett, , and E. Mortensen. Interactive live-wire boundary extraction. *Medical Image Analysis*, 1(4):331–341, 1997.

[7] R. Beare and G. Lehman. Efficient implementation of kernel filtering. In *Proceedings of ISC/NA-MIC Workshop on Open Science at MICCAI 2007*, pages 1–18. Insight Software Consortium, 2007.

[8] T. Boskamp, D. Rinck, F. Link, B. Kummerlen, G. Stamm, and P. Mildenberger. New vessel analysis tool for morphometric quantification and visualization of vessels in CT and MR imaging data sets. *Radiographics*, 24(1):287–297, 2004.

[9] J. Chen and A. Amini. Quantifying 3-d vascular structures in mra images using hybrid pde and geometric deformable models. *IEEE Transactions on Medical Imaging*, 23(10):1251–1262, 2004.

[10] Z. Chen and S. Molloi. Automatic 3D vascular tree construction in CT angiography. *Computerized Medical Imaging and Graphics*, 27(6):469–479, 2003.

[11] T. Deschamps. *Curve and Shape Extraction with Minimal Path and Level-Sets techniques: Applications to 3D Medical Imaging*. PhD dissertation, University of Paris Dauphine, 2001.

[12] E. Fishman, D. Ney, D. Heath, F. Corl, K. Horton, and P. Johnson. Volume rendering versus maximum intensity projection in CT angiography: What works best, when, and why. *Radiographics*, 26(3):905–922, 2006.

[13] A. Frangi, W. Niessen, K. Vincken, and M. Viergever. Multiscale vessel enhancement filtering. In *Proceedings of Medical Image Computing and Computer Assisted Intervention (MICCAI 1998)*, volume 1496, pages 130–137. LNCS, Springer, 1998.

[14] H. Hahn, B. Preim, D. Selle, and H.-O. Peitgen. Visualization and interaction techniques for the exploration of vascular structures. In *Proceedings of IEEE Visualization*, pages 395–578. IEEE, 2001.

[15] M. Holtzman-Gazit, R. Kimmel, N. Peled, and D. Goldsher. Segmentation of thin structures in volumetric medical images. *IEEE Transactions on Image Processing*, 15(2):354–363, 2006.

[16] J. Jomier, V. LeDigarcher, and S. Aylward. Comparison of vessel segmentations using STAPLE. In *Proceedings of Medical Image Computing and Computer Assisted Intervention (MICCAI 2005)*, volume 3749, pages 523–530. LNCS, Springer, 2005.

[17] A. Kanitsar, R. Wegenkittl, D. Fleischmann, and M. Gröller. Advanced curved planar reformation: flattening of vascular structures. In *Proceedings of IEEE Visualization*, pages 43–50. IEEE, 2003.

[18] R. Kimmel, A. Amir, and A. Bruckstein. Finding shortest paths on surfaces using level sets propagation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(6):635–640, 1995.

[19] C. Kirbas and F. Quek. A review of vessel extraction techniques and algorithms. *ACM Computing Surveys*, 36(2):81–121, 2004.

[20] K. Krissian, G. Malandain, N. Ayache, R. Vaillant, and Y. Trousset. Model-based detection of tubular structures in 3D images. *Computer Vision and Image Understanding*, 80(2):130–171, 2000.

[21] L. Lorigo, W. Faugeras, O.and Grimson, R. Keriven, R. Kikinis, A. Nabavi, and C.-F. Westin. CURVES: Curve evolution for vessel segmentation. *Medical Image Analysis*, 5(3):195–206, 2001.

[22] R. Manniesing, B. Velthuis, M. van Leeuwen, I. van der Schaaf, P. van Laar, and W. Niessen. Level set based cerebral vasculature segmentation and diameter quantification in CT angiography. *Medical Image Analysis*, 10(2):200–214, 2006.

[23] R. Manniesing, M. Viergever, and W. Niessen. Vessel enhancing diffusion: A scale space representation of vessel structures. *Medical Image Analysis*, 10(6):815–825, 2006.

[24] C. Maurer, R. Qi, and V. Raghavan. A linear time algorithm for computing exact euclidean distance transforms of binary images in arbitrary dimensions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(2):265–270, 2003.

[25] D. Mueller, A. Maeder, and P. O'Shea. Tagged volume rendering of the heart. In *Proceedings of Medical Image Computing and Computer Assisted Intervention (MICCAI 2007)*, volume 4791, pages 194–201. LNCS, Springer, 2007.

[26] H. Pannu, T. Flohr, F. Corl, and E. Fishman. Current concepts in multidetector row CT evaluation of the coronary arteries: Principles, techniques, and anatomy. *Radiographics*, 23:S111–S125, 2003.

[27] C. Pudney. Distance-ordered homotopic thinning: A skeletonization algorithm for 3D digital images. *Computer Vision and Image Understanding*, 72(3):404–413, 1998.

[28] Y. Sato, S. Nakajima, H. Atsumi, T. Koller, G. Gerig, S. Yoshida, and R. Ron Kikinis. 3D multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. In *Proceedings of CVRMed-MRCAS'97*, volume 1205, pages 213–222. LNCS, Springer, 1997.

[29] J. Sethian. *Level Set Methods and Fast Marching Methods*. Cambridge Press, 2nd edition, 1999.

[30] P. Soille. *Morphological image analysis : principles and applications*. Springer, New York, 2nd edition, 2003.

[31] M. Sonka, A. Stolpen, W. Liang, and R. Stefancik. Vascular imaging and analysis. In M. Sonka and J. Fitzpatrick, editors, *Handbook of Medical Imaging: Medical Image Processing and Analysis*, volume 2, pages 809–914. SPIE, 2000.

[32] M. van Droogenbroeck and M. Buckley. Morphological erosions and openings: Fast algorithms based on anchors. *Journal of Mathematical Imaging and Vision*, 22(2-3):121–142, 2005.

[33] M. van Droogenbroeck and H. Talbot. Fast computation of morphological operations with arbitrary structuring elements. *Pattern Recognition Letters*, 17(14):1451–1460, 1996.

[34] M. van Herk. A fast algorithm for local minimum and maximum filters on rectangular and octagonal kernels. *Pattern Recognition Letters*, 13(7):517–521, 1992.

[35] P. van Ooijen, K. Ho, J. Dorgelo, and M. Oudkerk. Coronary artery imaging with multidetector CT: visualization issues. *Radiographics*, 23(6):e16, 2003.

[36] S. Warfield, K. Zou, and W. Wells. Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. *IEEE Transactions on Medical Imaging*, 23(7):903–921, 2004.

[37] S. Wesarg and E. Firle. Segmentation of vessels: the corkscrew algorithm. In *Proceedings of Medical Imaging: Image Processing*, volume 5370, pages 1609–1620. SPIE, 2004.

[38] S. Wesarg, M. Khan, and E. Firle. Localizing calcifications in cardiac CT data sets using a new vessel segmentation approach. *Journal of Digital Imaging*, 19(3):249–257, 2006.

[39] O. Wink, W. Niessen, and M. Viergever. Fast delineation and visualization of vessels in 3-D angiographic images. *IEEE Transactions on Medical Imaging*, 19(4):337–346, 2000.

[40] O. Wink, W. Niessen, and M. Viergever. Minimum cost path determination using a simple heuristic function. In *Pattern Recognition*, volume 3, pages 998–1001. IEEE, 2000.

[41] O. Wink, W. Niessen, and M. Viergever. Multiscale vessel tracking. *IEEE Transactions on Medical Imaging*, 23(1):130–133, 2004.

[42] P. Yan and A. Kassim. Segmentation of volumetric MRA images by using capillary active contour. *Medical Image Analysis*, 10(3):317–329, 2006.

[43] K. Zou, S. Warfield, A. Bharatha, C. Tempany, M. Kaus, S. Haker, W. Wells, F. Jolesz, and R. Kikinis. Statistical validation of image segmentation quality based on a spatial overlap index. *Academic Radiology*, 11(2):178–189, 2004.