Shape and Appearance Models for Automatic Coronary Artery Tracking

Release 0.22

Sebastian Zambal[1], Jiří Hladůvka[1], Armin Kanitsar[2] and Katja Bühler[1]

August 11, 2008

[1]VRVis Research Center, Donau-City-Strasse 1, 1220 Vienna, Austria
[2]Agfa Healthcare GesmbH, Diefenbachgasse 35, 1150 Vienna, Austria

Abstract

Automatic tracking of coronary arteries in Computed Tomography Angiography (CTA) is a challenging task. To accomplish it we propose a method consisting of two main steps: (1) A 3D model of the heart is matched for detecting the approximate position of the heart. Based on this information candidates for coronary artery seeds are calculated. (2) Fitting of cylindrical sampling patterns is performed for extracting the vessel tree of coronary arteries. Branching and termination are handled by depth-first search and noise level estimation respectively. The average scoring is 39.4 for accuracy measures and 51.5 for overlap measures.

Latest version available at the Insight Journal [http://hdl.handle.net/1926/1926/1420]

Contents

1 Introduction

Automatic detection of coronary arteries in Computed Tomography Angiography (CTA) is a clinically very important but challenging task. Efforts are currently being undertaken to automate coronary artery tracking in order to improve the medical workflow and objectively quantify clinically important properties of coronary arteries.

There is a large amount of literature about vessel extraction in 2D and 3D medical images [6]. Many of the approaches rely on features such as the response of the Hessian matrix filter [8]. However, it is in general difficult to derive the full connectivity of a vessel tree from this information only. A method is required which takes into account the spatial relationships of features. For example, a shortest path algorithm with costs corresponding to feature responses may be applied for extracting a vessel [5].

A very promising direction in medical image analysis is the application of (statistical) models which represent shape (geometry) and appearance (texture) of the investigated anatomical structures. Active Appearance Models (AAMs) [2], for example, extract knowledge about shape and appearance from a set of training examples by applying Principal Component Analysis (PCA). Another approach uses Markov Random Fields (MRFs) in combination with symmetry features [4]. By matching a model of shape and appearance to unknown image data, high-level knowledge about spatial and gray value relationships of structures is exploited.

In this paper we refer to a model of shape and appearance as a set of spatial configurations of texture sample positions (shape model) and well defined relationships between gray values at these sample positions (texture model). When matching a model of shape and appearance to unknown image data, two critical issues arise: First, a reasonable objective function has to be used which sufficiently well reflects matching accuracy of the model. Second, an effective optimization technique has to be applied which robustly optimizes the objective function to achieve a match.

The method proposed in this paper employs two distinct shape and appearance models. First, an anatomical model of the complete heart is used (followed by symmetry feature extraction) to locate seed points of coronary arteries (section 2). Second, vessel tracking is accomplished by matching a cylinder-like model in combination with depth-first search (section 3). For detailed information on the data sets used for validation please refer to [7].

2 Detection of coronary artery seeds

Two steps are carried out to automatically detect potential seeds within coronary arteries close to the coronary ostia: First, a 3D model of the heart is matched to the volume data. Second, local maxima of a symmetry feature volume are calculated and identified as potential seeds.

2.1 3D model of the heart

A 3D model of the heart is built based on 2245 manually placed landmark points defined on data set 04. This data set is chosen since it exhibits low noise level according to the noise estimation algorithm outlined in section 3.3. Landmark points are defined on the left atrium, the left ventricle, the origin of the aorta, the pericardium, and outside around the heart. Figure 1(a) shows the model as wireframe and figure 1(b) as partially transparent solid. Figure 1(c) shows the model matched to data set 03.

A Delaunay tetrahedralization of the landmark points is performed to obtain a volumetric tetrahedron mesh. A grid of texture samples (with a cell size of 7mm edge length) is layed over the tetrahedron mesh. For each sample position the barycentric coordinates within the surrounding tetrahedron are stored. A deformation describing the displacements of landmark points can easily be passed on (via the precalculated barycentric coordinates) to the volumetric grid of texture samples.

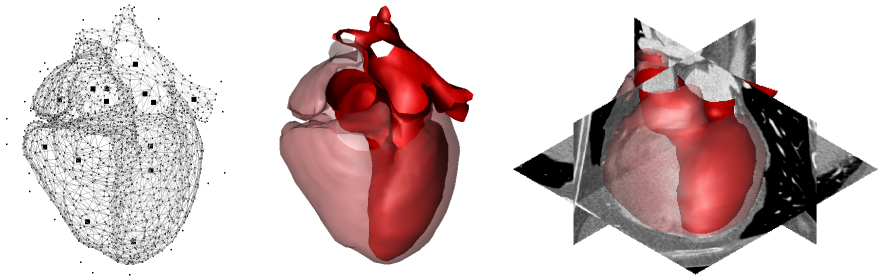

(a) (b) (c)

Figure 1: The 3D heart model (a) as wireframe (small black squares represent control points for diffeomorphic warps), (b) as partially transparent solid and (c) matched to data set 03.

To allow for smooth deformations of the landmark points during matching we apply a combination of diffeomorphic warps similar to those used, e.g., for diffeomorphic statistical shape models [1]. However, we do not use a regular grid for control points, but manually distribute them sparsely within the heart. Small black squares in figure 1(a) represent the 14 control points. Spatial transformations (scaling and translation) at control points are passed on to landmarks in the neighborhood. Deformation displacements are weighted with a function k of the distance between landmark and control point. As k we choose a scaled Gaussian function ($\sigma = 30mm$, $\mu = 0$). A scaling factor is used such that $k(0) = 1$.

The appearance part of the model assigns one out of three classes to each texture sample. A class is assigned to each voxel by investigation of the original gray values as observed in data set 04. Histogram equalization is performed for normalization first. A quantization of gray values to three quantization levels each representing one class is performed. As expected, the three classes correspond approximately to the regions of lung, myocardium and contrast enhanced blood.

Model matching accuracy (the objective function for matching) is evaluated by first sampling gray values according to the spatial configuration of the model in the current iteration. Next, gray value classification is estimated via maximum likelihood of the observed gray values. Finally, the *observed* classes are compared with the *expected* classes for each texture sample and a likelihood is calculated as in [10]. A probability of 0.1 is assigned to misclassified samples and a probability of 0.9 is assigned to correctly classified samples. We use a shape particle filtering approach [3] for calculating the optimal shape (with respect to the texture model) in the shape space spanned by the parameters of rigid transformation and diffeomorphic warps.

2.2 Local symmetry

Two of the 3D model's landmark points, **L** and **R**, are located directly on the left and right coronary arteries close to the aorta (figure 2(a)). After model matching these two landmarks define the locations where the ostia of the two main coronary arteries are approximately located. Due to anatomical variations, however, these points will probably not exactly reside on the coronary arteries of the unknown data set after model matching. This is why we look for coronary artery seeds within a region of radius r_{init} around **L** and **R**. r_{init} is chosen 2.5 times the distance of two landmark points $\mathbf{A_1}$ and $\mathbf{A_2}$ inside the aorta as illustrated in figure 2(a).

At this stage of the algorithm we are interested in *potential* coronary artery seeds. The idea is to gain robustness by calculating multiple candidates and to later select the best ones after more detailed evaluation. In fact we decide which of the potential seeds are used after few tracking iterations using the tracking algorithm outlined in section 3. We calculate a simple symmetry feature to identify locally promising seed

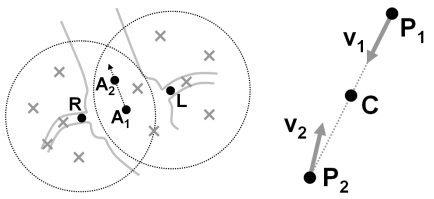

(a) (b)

Figure 2: (a) Four landmark points $\mathbf{A_1}$, $\mathbf{A_2}$, $\mathbf{L}$, and $\mathbf{R}$ are used to define regions that are searched for coronary artery seeds. Local maxima of symmetry features (potential seeds) are illustrated as bright crosses. (b) Calculation of symmetry feature contribution along scan lines.

candidates. A symmetry feature volume for regions around $\mathbf{L}$ and $\mathbf{R}$ is initialized with value 0 for every voxel and 3D image gradients are precalculated. A scan line of length $10mm$ (which is assumed the maximum diameter of a coronary artery) along the direction of the gradient vector $\mathbf{v_1}$ is traversed. Figure 2(b) illustrates this for a voxel at position $\mathbf{P_1}$. At every position along the scanline such as $\mathbf{P_2}$, a second gradient vector $\mathbf{v_2}$ is investigated. The squared length of $\mathbf{v_1} - \mathbf{v_2}$ is calculated and added to the voxel within the feature volume at position $\mathbf{C}$ (which lies in the middle between $\mathbf{P_1}$ and $\mathbf{P_2}$). This procedure is carried out for every voxel. Finally, the feature volume contains high values at voxels which locally exhibit high symmetry. After the symmetry volume has been generated, local maxima are identified as potential coronary artery seeds.

Each potential seed is tracked 3 iterations wide using vesselness function v defined below. For these three iterations the average vesselness is calculated and the seed with the best evaluation is used as the first coronary artery to be further tracked. Vector $\mathbf{a}$ from $\mathbf{A_1}$ to the first symmetry maximum is calculated. The next best candidate with a vector $\mathbf{b}$ from $\mathbf{A_1}$ to this candidate and an angle of more than 50° between $\mathbf{a}$ and $\mathbf{b}$ is selected as the second coronary artery seed. For these two finally selected seeds, tracking backward into the aorta is performed by using the tracking method outlined below.

3 Vessel tracking

We accomplish vessel tracking by incremental matching of a shape and appearance model. In contrast to the 3D model for the heart it is difficult to formulate a generic shape model for the vessel tree of coronary arteries. Shape and topology vary substantially between patients. We incrementally grow the vessel tree starting from the two seeds of left and right coronary arteries. In principle we proceed in a similar way as a recently proposed method that uses Bayesian tracking [9]. However, our method differs in how successive candidates are selected. Furthermore our method handles branching and termination.

3.1 Estimating vesselness

We track a vessel by iteratively matching a rigid model with simple symmetric shape. For a given position $\mathbf{x}$, a vector $\mathbf{d}$ indicating the tangential direction of the vessel, and a radius r, we define a function $v(\mathbf{x}, \mathbf{d}, r)$ which reflects the "vesselness" of the given configuration for a given volume data set. Texture samples are created by using a shape pattern of two concentric circles which are translated and scaled along direction $\mathbf{d}$. A pure translation of two concentric (an interior and an exterior) circles leads to a perfectly cylindrical pattern. However, if – as illustrated in figure 3(a) – the artery exhibits strong bending at the investigated position, a straight cylindrical model does not fit well. Along the central axis we scale the circles which carry texture samples by using a scaling function $s(x)$ where x is the distance to the center of the pattern divided by the radius r.

In fact we use two patterns for two different levels of detail. In the first step a pattern with larger margins

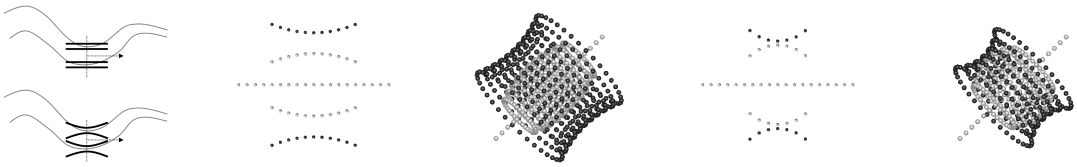

(a) (b) (c) (d) (e)

Figure 3: (a) Straight cylindrical pattern versus bent pattern for estimation of vesselness. (b) cross section and (c) 3D view of sampling pattern 1. (d) cross section and (e) 3D view of sampling pattern 2.

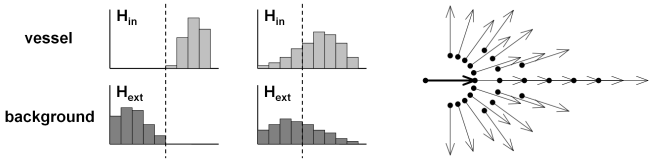

(a) (b)

Figure 4: (a) The histograms generated from locally sampled texture. (b) Candidates evaluated for finding the successive segment. Dots with arrows indicate positions and orientations of the segements. The bold arrow represents the current segment.

but smaller scaling is used for evaluating the approximate direction of the vessel. This pattern is depicted in figure 3(b) (cross section) and 3(c) (3D). Quadratic scaling functions $s_{1,in}(x) = 0.75 - x^2 * 0.2$ and $s_{1,ext}(x) = 1.25 + x^2 * 0.2$ of interior and exterior radii are applied along the axis. In the second step the estimated position of the vessel is refined by using a second pattern depicted in 3(d) (cross section) and 3(e) (3D). Circle-shaped scaling functions $s_{2,in}(x) = -0.05 + \sin(\arccos(x))$ and $s_{2,ext}(x) = 2.05 - \sin(\arccos(x))$ of interior and exterior radii are applied along the axis. For both patterns additional samples lying at the central axis along the direction of the vessel are sampled and marked as interior samples. In figures 3(b), 3(c), 3(d), and 3(e) the interior sample positions are colored gray and exterior samples are colored black.

For a given shape configuration $(\mathbf{x}, \mathbf{d}, r)$, the texture values of the sampling pattern are accordingly extracted from the volume data using nearest neighbor interpolation. These values are used to create two separate histograms H_{in} and H_{out} for interior and exterior sample values respectively. The basic assumption is made that in case of a perfect alignment with a vessel, there must be a threshold that optimally separates both histograms. This is illustrated in figure 4(a) where the dashed line represents the threshold and on the left side perfectly separates samples from inside and outside the vessel. On the right side the separation is not so obvious. The number of values which are (as expected) greater than the threshold for the interior histogram are divided by the total number of interior samples. The resulting fraction f_{in} is the fraction of gray values which have a value as expected by the model. Similarly, for the exterior histogram the fraction f_{ext} of values greater than the threshold (as expected) is calculated. The vesselness function v is finally defined as the mean of the above fractions:

$$v(\mathbf{x}, \mathbf{d}, r) = \frac{f_{in} + f_{ext}}{2} \tag{1}$$

where the threshold is chosen such that v is maximized.

3.2 Tracking of vessels

The proposed vessel tracking algorithm proceeds in a depth-first fashion by adding new vessel segments to the vessel tree that has been extracted so far. A vessel segment is defined again by a shape configuration $(\mathbf{x}, \mathbf{d}, r)$ as outlined in section 3.1. By convention $\mathbf{d}$ points from proximal to distal. The idea is to compare

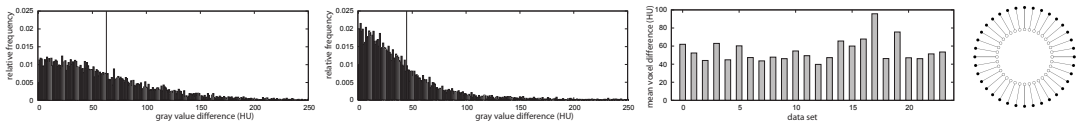

(a) (b) (c) (d)

Figure 5: The relative frequencies of local gray value differences in data sets 03 (a) and 04 (b) respectively (vertical lines indicate mean absolute differences $\overline{\delta}_g$). (c) The mean of estimated noise distributions for workshop data sets 00-23. (d) Pairs of samples compared with mean voxel difference for termination criterion.

possible candidates for the successive vessel segment distal to the current one. Vessel tracking is continued with the segment that exhibits the highest value of the vesselness function v.

Let $\mathbf{s}_c = (\mathbf{x}_c, \mathbf{d}_c, r_c)$ be the current vessel segment. A set of candidates for successive segments is calculated in the following way. Direction vectors $\mathbf{d}_{n,i}$ with angles α between $-90°$ to $90°$ relative to $\mathbf{d}_c$ in steps of $18°$ are considered. Radii for the new candidates range from $0.5r_c$ to $1.5r_0$ in scaling increments of 1.1. Positions for candidates are generated as $\mathbf{x}_{n,i} = \mathbf{p}_c + \Delta(\mathbf{d}_c + \mathbf{d}_{n,i})/2$ where Δ is the step size. We use a variable number of step sizes n_{steps} depending on α. The values 6, 3, 2, 1 are assigned to n_{steps} for $\alpha < 5°$, $30°$, $60°$, $90°$ respectively. The distance to the successive candidates is calculated as $(n_{steps} \cdot r_c)/2 + 0.5$. This scheme for the generation of successive candidate positions $\mathbf{x}$ and directions $\mathbf{d}$ is illustrated in figure 4(b).

From all generated candidates the one with the most significant (largest) vesselness $v(\mathbf{s}_{n,i})$ is selected as the successive candidate. In order to refine the candidate, a set of refined (altered position, scaling, radius) candidates close to the selected one is generated by using sampling pattern 2. Again the one with the most significant vesselness survives.

The second best candidate for which the direction vector differs more than $30°$ to the direction vector of the best successive candidate is stored as a possible branch. The vessel is tracked until the termination criterion which is discussed in section 3.3 is fulfilled. Then all possible branches of the terminated vessel are investigated. Each branch is tracked five segments wide. Depth-first search is continued with the branch that exhibits the largest average vesselness of the first five segments. The algorithm proceeds until a maximum number of vessels (in the reported experiments we used 30) is found. An additional restriction is that sub-branching is stopped for vessels where a branching depth of two is reached.

3.3 Termination criterion

We estimate the level of noise in a CTA volume before starting vessel tracking. Later this estimate is used in a noise-adaptive criterion for vessel termination. The vessel termination criterion is not related with the vesselness function v which is only used for selection of the optimal successive segment. We apply a simple noise estimation scheme where we consider gray value differences, δ_g, of neighboring voxel pairs. A neighboring voxel pair is defined as a pair of voxels with a distance of three voxel sizes. 100,000 voxel pairs are randomly chosen and absolute gray value differences are measured. In order to keep out regions within the lung, only voxel pairs are considered for which both gray values are greater than -500 Hounsfield units (HU). Figures 5(a) and 5(b) show the resulting relative frequencies of gray value differences for data sets 03 and 04. The mean values of gray value differences $\overline{\delta}_{g,03}$ and $\overline{\delta}_{g,04}$ for data sets 03 and 04 are 63 and 45 Hounsfield units respectively. The mean values for all data sets are shown in figure (c).

We consider a gray value difference, δ_g, less than or equal the mean, $\overline{\delta}_g$, most likely to originate from noise. A gray value difference greater than $\overline{\delta}_g$ rather indicates a meaningful transition between background and object at the corresponding voxel pair. During vessel tracking, pairs of gray values are considered: each

interior texture sample is compared to its radially exterior neighbor sample (figure 5(d); sample pairs are connected by lines). If more than half of the measured texture differences are less than the average gray value difference, the vessel is considered not distinguishable from the background. If two successive vessel segments fulfill this criterion the vessel is terminated.

In fact we use some additional vessel termination criteria. If the vessel intersects with itself (a segment overlaps with one of its preceding segments) the vessel is terminated. Further, if the mean gray value of exterior samples falls below a Hounsfield value of -500 (lung vessel), tracking is stopped.

4 Results

For selecting the right vessels for evaluation from the automatically extracted vessel tree we apply the following procedure. For a given point **A** the closest vessel with a distance of less than 5mm is selected. If no vessel is closer than 5mm to point **A**, point **B** is used. If there is also no vessel closer than 5mm to point **B** the vessel is not detected. For 8 out of the 64 reference vessels of data sets 08-23 point **B** is used for selection. More precisely point **B** is used for the following vessels: data set 08, vessels 1 and 2; data set 11, vessel 3; data set 14, vessel 3; data set 19, vessel 1; data set 22, vessel 1; data set 23, vessel 2. All remaining 56 vessels are identified using point **A**.

Although in general our method handles bifurcations well, errors occur in very complex cases. For example, the correct branch is not found for vessel 3 in data set 11 and the vessel is only tracked to point **B**.

Depending on the branching depth of the vessel tree the run times of the algorithm range from 4 to 8 minutes per data set (measured on a PC with Pentium 4 CPU at 3.2GHz and 2GB of RAM).

5 Conclusion

We have outlined a method for automatic centerline extraction of the coronary artery tree in CTA. The three most critical issues in vessel tracking probably are vesselness estimation, termination criterion, and branching. In the presented approach we handle these issues by putting multiple texture samples into relation for vesselness estimation. For this we avoid the use of a sample mean of gray values because of its sensitivity to outliers. Instead we apply a method which evaluates the separability of histograms. Noise estimation is performed to achieve an effective termination criterion. Branching is handled by applying a depth-first search where new branches are selected after few tracking iterations have been performed.

Table 1: Average overlap per dataset

Dataset	**OV**			**OF**			**OT**			**Avg.**
nr.	%	score	rank	%	score	rank	%	score	rank	**rank**
8	73.4	39.8	–	62.7	45.3	–	76.0	38.1	–	–
9	91.1	55.3	–	68.8	56.7	–	93.6	71.8	–	–
10	95.4	49.1	–	76.6	64.3	–	95.8	72.9	–	–
11	71.9	36.7	–	52.8	52.1	–	72.1	49.2	–	–
12	88.8	45.8	–	25.8	14.5	–	92.3	46.5	–	–
13	91.9	58.9	–	70.7	59.2	–	92.6	61.0	–	–
14	83.6	42.2	–	46.8	39.0	–	86.1	55.5	–	–
15	97.6	58.9	–	89.2	70.7	–	98.8	74.4	–	–
16	89.2	48.8	–	72.6	48.1	–	92.4	58.7	–	–
17	72.5	42.1	–	13.1	7.3	–	73.2	38.2	–	–
18	90.6	46.1	–	72.6	61.6	–	92.2	71.1	–	–
19	84.3	56.1	–	79.4	66.0	–	84.9	55.0	–	–
20	96.3	66.5	–	37.7	21.8	–	98.6	49.5	–	–
21	90.4	51.4	–	45.6	29.4	–	96.1	68.3	–	–
22	77.9	39.3	–	77.3	51.2	–	79.8	64.9	–	–
23	82.2	46.1	–	71.0	60.8	–	83.7	66.8	–	–
Avg.	**86.1**	**48.9**	**–**	**60.2**	**46.8**	**–**	**88.0**	**58.9**	**–**	**–**

Table 2: Average accuracy per dataset

Dataset	AD			AI			AT			Avg.
nr.	mm	score	rank	mm	score	rank	mm	score	rank	rank
8	3.93	36.8	–	0.32	48.3	–	3.75	38.0	–	–
9	0.95	40.0	–	0.20	43.9	–	0.80	41.1	–	–
10	0.47	36.0	–	0.28	37.7	–	0.46	35.4	–	–
11	7.03	31.1	–	0.33	44.0	–	6.96	31.2	–	–
12	1.22	37.0	–	0.26	41.2	–	0.53	38.3	–	–
13	1.15	40.7	–	0.27	44.1	–	1.13	41.3	–	–
14	2.03	37.9	–	0.31	44.2	–	1.67	38.7	–	–
15	0.34	41.8	–	0.23	42.7	–	0.29	42.6	–	–
16	1.44	37.5	–	0.26	41.9	–	1.29	37.8	–	–
17	4.26	39.2	–	0.34	47.3	–	4.27	39.6	–	–
18	1.37	37.9	–	0.24	41.6	–	1.29	38.4	–	–
19	4.26	38.0	–	0.30	44.7	–	4.21	38.2	–	–
20	0.69	41.5	–	0.37	42.9	–	0.39	42.5	–	–
21	1.45	37.2	–	0.21	40.8	–	0.58	39.1	–	–
22	5.06	29.7	–	0.31	37.3	–	4.86	30.2	–	–
23	5.75	35.9	–	0.25	41.9	–	4.76	36.5	–	–
Avg.	**2.59**	**37.4**	**–**	**0.28**	**42.8**	**–**	**2.33**	**38.1**	**–**	**–**

Table 3: Summary

Measure	% / mm			score			rank		
	min.	max.	avg.	min.	max.	avg.	min.	max.	avg.
OV	16.1%	100.0%	86.1%	8.3	100.0	48.9	–	–	–
OF	0.0%	100.0%	60.2%	0.0	100.0	46.8	–	–	–
OT	16.2%	100.0%	88.0%	8.5	100.0	58.9	–	–	–
AD	0.20 mm	25.30 mm	2.59 mm	9.9	51.3	37.4	–	–	–
AI	0.17 mm	0.57 mm	0.28 mm	30.2	59.0	42.8	–	–	–
AT	0.16 mm	25.06 mm	2.33 mm	9.9	52.8	38.1	–	–	–
Total							**–**	**–**	**–**

References

[1] T.F. Cootes, C.J. Twining, and C.J. Taylor. Diffeomorphic statistical shape models. In *British Machine Vision Conference*, volume 1, pages 447–456, 2004. 2.1

[2] Timothy F. Cootes, Gareth J. Edwards, and Christopher J. Taylor. Active Appearance Models. In *European Conference on Computer Vision*, volume 2, pages 484–498, 1998. 1

[3] M. de Bruijne and M. Nielsen. Shape particle filtering for image segmentation. In *Medical Image Computing and Computer–Assisted Intervention*, volume 1, pages 168–175, 2004. 2.1

[4] R. Donner, B. Micušík, G. Langs, and H. Bischof. Sparse MRF appearance models for fast anatomical structure localisation. In *British Machine Vision Conference*, 2007. 1

[5] Armin Kanitsar. *Curved Planar Reformation for Vessel Visualization*. PhD thesis, Institute of Computer Graphics and Algorithms, Vienna University of Technology, Favoritenstrasse 9-11/186, A-1040 Vienna, Austria, 2004. 1

[6] C. Kirbas and F.K.H. Quek. A review of vessel extraction techniques and algorithms. *ACM Computing Surveys*, 36(2):81–121, 2004. 1

[7] C. Metz, M. Schaap, T. van Walsum, A. van der Giessen, A. Weustink, N. Mollet, G. Krestin, and W. Niessen. 3D segmentation in the clinic: A grand challenge II - coronary artery tracking. *Insight Journal*, 2008. 1

[8] Y. Sato, S. Nakajima, H. Atsumi, T. Koller, G. Gerig, S. Yoshida, and R. Kikinis. 3D multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. In *CVRMed*, pages 213–222, 1997. 1

[9] M. Schaap, I. Smal, C.T. Metz, T. van Walsum, and W.J. Niessen. Bayesian tracking of elongated structures in 3D images. In *Information Processing in Medical Imaging*, 2007. 3

[10] S. Zambal, K. Bühler, and J. Hladůvka. Entropy-optimized texture models. In *Medical Image Computing and Computer–Assisted Intervention*, 2008. 2.1