Edge Based Tube Detection for Coronary Artery Centerline Extraction

Release 0.00

Christian Bauer[1] and Horst Bischof[1]

July 4, 2008

[1]Institute for Computer Graphics and Vision, Graz University of Technology, Austria

Abstract

The extraction of the coronary artery central lumen lines from CTA datasets is a necessary prerequisite for the computerized assessment of heart related disease. In this work, we present an automatic approach for this task that consists of generic methods for detection of tubular objects, extraction of their centerlines, grouping of the single centerlines into complete tree structures, and some application specific adaptions for the identification of the coronary arteries. The tube detection approach is based on the Gradient Vector Flow and an analysis of the resulting vector field. Contrary to conventional tube detection filters this approach avoids multi-scale analysis with related scale space problems and is able to identify tubular objects surrounded by different tissues such as blood vessels in proximity of calcifications. After identification of the tubular structures their centerlines are extracted and grouped into complete tree structures. Based on gray value information and centerline length tubular structures not belonging to the coronary arteries are removed. The approach has been evaluated on 16 clinical datasets showing a high overlap of 94% with undisclosed reference centerlines and an average distance of 0.58mm.

Contents

Coronary artery disease is the first leading cause of death in the USA. For the assessment of coronary artery disease like plaques, aneurysms or stenoses computer tomography angiography (CTA) has become an established method as with current multi-slice CT scanners the interesting structures can be imaged with sufficient detail. For diagnosis, visual inspection of 2D slices of coronary CTA data is still common. However, to facilitate radiologists in this process and to raise the quality of the assessment from a pure qualitative level to a quantitative level, computer aided tools for this task are desired. An integral part for these tools is the identification of the coronary arteries and an extraction of their central lumen line.

For coronary artery central lumen line extraction methods the main difficulties are closely adjacent vessels or calcifications, the proximate heart chambers, changing gray value along the vessels, and regions where

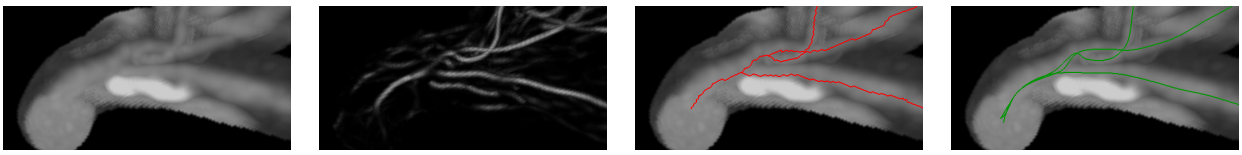

Figure 1: Centerline extraction at the proximal end of the conoary artery containing a calcification. From left to right: original dataset, tube detection response T, extracted centerlines, provided reference centerlines.

locally the vessels are partly indistinguishable from the background; e.g. due to occlusions or imaging artifacts. In the literature several techniques for the segmentation/centerline extraction of 3D vascular structures in general (see e.g [5]) and the coronary arteries in particular (see e.g [6]) have been presented ranging from highly interactive to fully automatic approaches. Methods based on tube detection filters (aka. vessel detection or lineness filters) have been applied for automatic detection of the blood vessels with a robustness against several before mentioned issues. However, tube detection filters typically assume that the blood vessels form ridges in Gaussian scale space, whereby this assumption is problematic in case of proximate image structures within the same gray value range (e.g. other vessels or the heart chamber) and the assumption does not hold in case of adjacent structures that appear brighter (e.g. calcifications in the blood vessels).

In previous works [2, 1] we presented a tube detection approach that is capable of successfully dealing with the before mentioned situations of proximate structures with the same gray value and different background conditions (see Figs. 1 and 3). Thus, it is well suited for the task of coronary artery detection. In this work, we present an automatic approach for the coronary artery centerline extraction based on this tube detection filter [2] and the centerline traversal method we used in [1].

1 Methodology

Our approach consists of three main parts: 1) Bottom up identification of all tubular objects in the whole dataset using a tube detection filter (TDF) and extraction of their centerlines using a ridge traversal procedure. 2) Grouping of these centerlines into complete tree structures. During this step also smaller gaps in the centerline representation are closed that may occur in proximity of junctions or in regions where locally the vessels are partly indistinguishable from the background; e.g. due to occlusions or imaging artifacts. 3) Discarding centerlines not belonging to the coronary arteries based on application specific prior knowledge.

Tube Detection Filtering: Initially, a TDF is used that identifies structures with a circular – or approximately circular – cross section profile that may be surrounded by arbitrary tissues [2, 1]. The approach consists of on an edge-preserving diffusion of directional edge information (gradient vectors) using the Gradient Vector Flow (GVF) [8] and an analysis of the second order structure of the resulting vector field (see Fig. 2).

The approach requires an appropriate noise-reduction in the given image I; a median filter of size $5\times5\times5$ voxels is used in this work. Depending on the application domain, the TDF can be adapted to the expected edge type: $F=\nabla I$ for tubes surrounded by darker tissue, $F=-\nabla I$ for tubes surrounded by brighter tissue, or $F=-\nabla|\nabla I|$ for tubes surrounded by arbitrary step edges such as blood vessels adjacent to calcifications. After normalization of the vector magnitudes to account for varying contrast situations: $F_n(\mathbf{x})=F(\mathbf{x})/|F(\mathbf{x})|*(min(|F(\mathbf{x})|,F_{max})/F_{max})$ where $\mathbf{x}=(x,y,z)$, the GVF as presented by Xu and Prince [8] is performed. The GVF is defined as the vector field $V(\mathbf{x})$ that minimizes:

$$E(V)=\iiint_{\Omega}\mu|\nabla V(\mathbf{x})|^2+|F_n(\mathbf{x})|^2|V(\mathbf{x})-F_n(\mathbf{x})|^2d\mathbf{x} \tag{1}$$

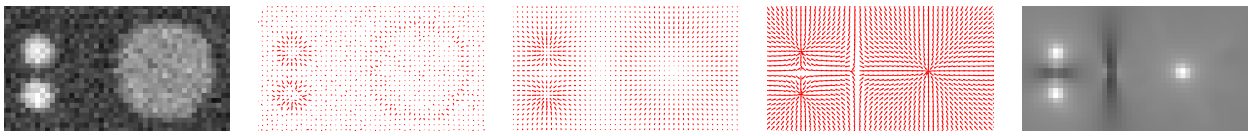

Figure 2: Processing steps of the GVF based tube detection. From left to right: Original dataset I, initial vector field F_n, GVF result V, normalized GVF result V_n, and the final tube likeliness T.

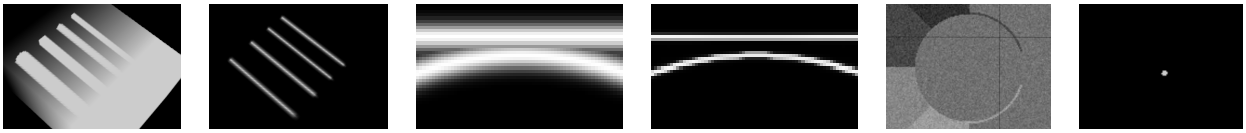

Figure 3: Properties of the GVF based tube detection approach. From left to right: Varying tube diameter and contrast (MIP), closely tangenting tubes (MIP), varying edge-types (2D cross section of 3D dataset)).

where μ originally has to be adapted to the noise level; as the initial image and the vector magnitudes are normalized (including a noise reduction), a fixed value of $\mu = 5$ can be used. The GVF performs an edge-preserving diffusion of the directional edge information, whereby for tubular objects the resulting vector field shows the same characteristic properties at the tubes centerline that are used by other tube detection filters for classification: All vectors point towards the centerline of tubular objects and the vector field shows a large variance in two directions and a low variance in the third direction (see Fig. 2).

To derive a measure of tube likeliness from the resulting vector field $V_n(\mathbf{x}) = V(\mathbf{x})/|V(\mathbf{x})|$ (the magnitude of the vectors is not of importance, just their direction) the Hessian matrix $H(\mathbf{x}) = \nabla V_n(\mathbf{x})$ with its eigenvalues $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$ can be used. Currently, we are using Frangi's vesselness measure [4] with default parameters ($\alpha = 0.5$, $\beta = 0.5$, and $\gamma = 100$) to derive the final tube-likeliness from these eigenvalues since this formulation is well known:

$$T(\mathbf{x}) = \begin{cases} 0 & \text{if } \lambda_2 > 0 \text{ or } \lambda_3 > 0 \\ \left(1 - exp\left(-\frac{R_A^2}{2\alpha^2}\right)\right) exp\left(-\frac{R_B^2}{2\beta^2}\right) \left(1 - exp\left(-\frac{S^2}{2c^2}\right)\right) & otherwise \end{cases} \tag{2}$$

with $R_A = |\lambda_1|/\sqrt{|\lambda_2||\lambda_3|}$, $R_B = |\lambda_2|/|\lambda_3|$, and $S = \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}$.

This combination of the normalized GVF field and Frangi's vesselness measure allows detection of tubular objects (at their centerlines). The response is independent of the tubes size and contrast, the approach is robust against leakage and allows handling of varying edge types, and – as an edge-preserving diffusion is performed – it allows separation of closely adjacent image structures with a similar gray value where most tube detection approaches (in particular those based on scale-space analysis) have problems (see Fig. 3).

Applying this approach for detection of coronary arteries on heart CTA datasest results in the response shown in Fig. 4(e), whereas applying Frangi's vesselness measure to the original vector field F_n results in the response shown in Fig. 4(d). The GVF based approach performs well in identifying major parts of the coronary arteries independent of their size and produces clean responses at their centerlines, but for the very thin low-contrast vessels at the distal ends of the arteries the edge-information is too low and the response falls off completely. Compared to that, applying Frangi's vesselness measure to the original vector field F_n the response for the thin low-contrast vessels just decreases, but it also produces responses of similar strength due to some imaging artifacts/noise that locally may look similar to such thin vessels. Resolving this ambiguity in the tube detection filtering without a more sophisticated pre-processing for noise-suppression

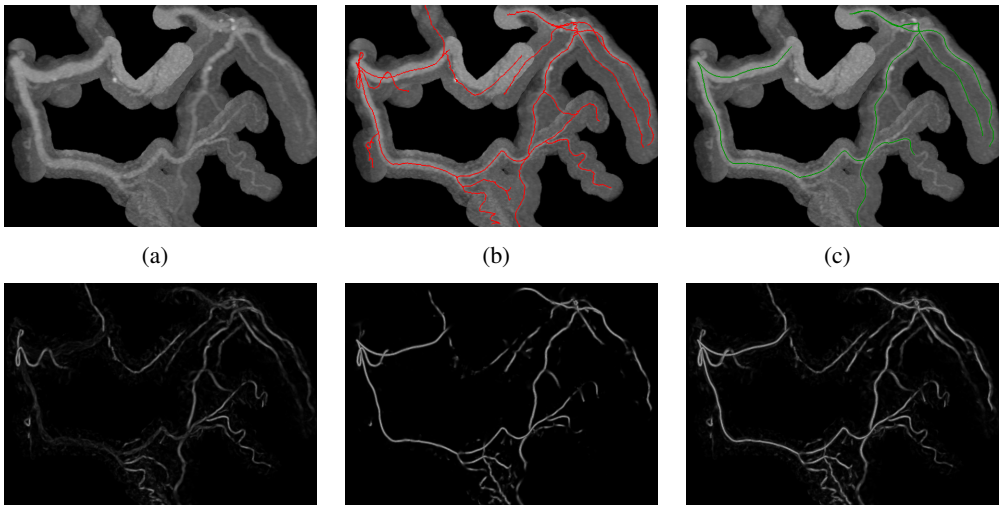

Figure 4: Tube detection step. For visualizations just the image regions in proximity of the coronary artery trees are shown. (a) MIP of the original dataset. (b) Extracted coronary artery trees. (c) Provided reference centerlines. (d) Tube detection result on the initial vector field F_n. (e) Tube detection result on the normalized GVF field V_n. (f) Combined tube detection result T.

seems unrealistic. However, the usage of a hysteresis thresholding (see next paragraph) enables resolving this ambiguity. Of course, applying Frangi's vesselness measure to F_n, does not allow detection of larger image structures. To enable an extraction of larger vessels and thin-low contrast vessels, the maximum of both filter responses – Frangi's vesselness measure applied to F_n and V_n – is used which is in the remainder of this article considered as the TDF response T (see Figs. 1, 4(f), 5(b) and 5(c)).

Centerline Extraction: After applying the TDF, it's response T is analyzed and centerlines of elongated structures are extracted using a ridge traversal procedure with a hysteresis thresholding. This discards short spurious responses and accounts for the fact that for thin low-contrast vessels or in proximity of junctions the TDF response also decreases. Further, it is computationally efficient and immediately allows an extraction of a higher level description of the centerlines.

The procedure requires a given starting point $\mathbf{x}_0$ for every ridge and an estimation of its tangent direction $\mathbf{t}(\mathbf{x})$; $\mathbf{t}(\mathbf{x})$ is estimated using the Hessian matrix. The complete medial curve is extracted by traversing the ridge in both directions $\mathbf{t}(\mathbf{x}_0)$ and $-\mathbf{t}(\mathbf{x}_0)$. Given a current point on the ridge $\mathbf{x}_i$ the next point $\mathbf{x}_{i+1}$ in the traversal direction $\mathbf{t}_i$ is chosen as the local neighbour $\mathbf{x}_i^n$ with the highest value that satisfies $\overrightarrow{\mathbf{x}_i\mathbf{x}_i^n} \cdot \mathbf{t}_i > 0$. To maintain the correct direction during traversal $\mathbf{t}(\mathbf{x}_{i+1})$ is set to $\mathbf{t}(\mathbf{x}_{i+1}) = \text{sign}(\overrightarrow{\mathbf{x}_i\mathbf{x}_{i+1}} \cdot \mathbf{t}(\mathbf{x}_{i+1}))\mathbf{t}(\mathbf{x}_{i+1})$ and the procedure is repeated until a stopping criterion is met: an already traversed point is found (a junction of two centerlines) or the value falls below a given threshold (an endpoint of the centerline).

A starting point for every ridge can be found easily: As we know that the TDF response increases for tubular structures and falls off in proximity of junctions, we can conclude that every tubular object contains at least one local maximum in the TDF response image. Thus, all local maxima are considered as potential starting points for the ridge traversal (they are processed in descending order of their TDF response). Starting from these points connected regions with a minimum value $T(\mathbf{x}) > t_{high} * \max_{y\in\Omega} T(\mathbf{y})$ and a minimum length ($length > l_{min}$) are identified as centerlines of tubular objects and further traversed while the TDF response stays above $T(\mathbf{x}) > t_{low} * \max_{y\in\Omega} T(\mathbf{y})$. This way very short spurious responses are discarded.

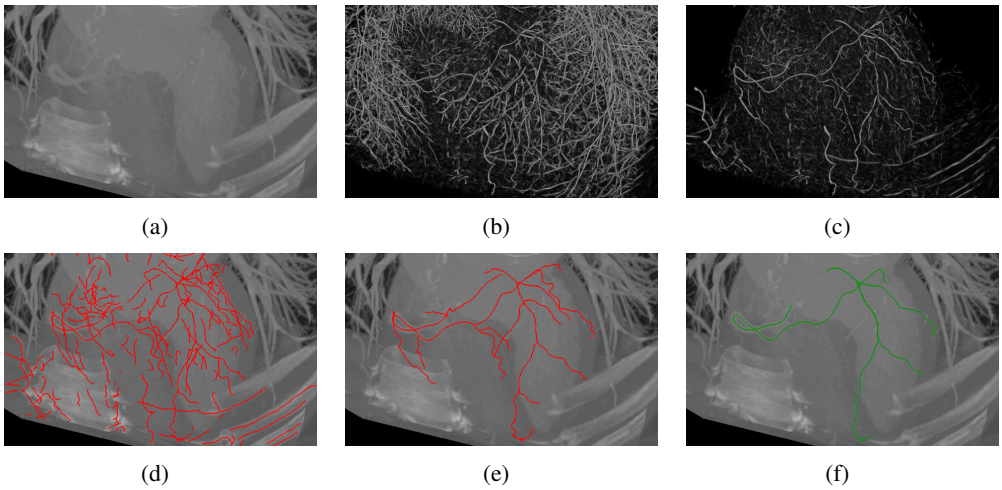

Figure 5: Centerline extraction and vascular tree reconstruction. (a) MIP of the original dataset (b) Tube detection result. (c) Tube detection result after suppression of responses inside the lung. (d) Result of the centerline extraction. (e) Identified coronary artery trees after tree reconstruction. (f) Provided reference.

In Fig. 5(d) the extracted centerlines from the TDF response shown in Fig. 5(c) are shown. The outlined ridge traversal performs well in extracting the centerlines of the tubular objects also at the thin low-contrast ends where the TDF response falls off considerably while still discarding short spurious responses. This way major parts of the coronary artery centerline are already extracted, but in junction areas or regions where the vessels are partly indistinguishable from the background, the centerlines may break up.

Grouping into Tree Structures: To account for these situations and to reconstruct the complete vascular trees the single centerlines are grouped and linked together. Therefore possible connections (and connection costs) between the centerlines are identified based on distance, angle, and gray value difference and a minimum spanning forest is constructed (see [3], Chapter 23), that contains the coronary artery trees. The connection costs are only computed between complete centerlines and not all centerline points. The endpoints of single centerlines are considered and the approach searches in the outward pointing tangent-direction of the centerline for potential continuations. Based on the centerline points coordinate $\mathbf{x}_S$ and its tangent-direction $\mathbf{t}_S$ connection costs to other centerline points $\mathbf{x}_E$ are computed and used for the minimum spanning forest construction; possible connections with a too large gray value difference $|I(\mathbf{x}_S) - I(\mathbf{x}_E)| > d_{max}$ are immediately rejected. The connection costs $C(\mathbf{x}_S, \mathbf{t}_S, \mathbf{x}_E)$ represent a trade-off between distance and angle:

$$C(\mathbf{x}_S, \mathbf{t}_S, \mathbf{x}_E) = ||\mathbf{x}_S - \mathbf{x}_E|| / exp(-\angle(\overrightarrow{\mathbf{x}_S\mathbf{x}_E}, \mathbf{t}_S)/(2\rho^2)); \tag{3}$$

and can also can be considered as a kind of coned region (with an opening angle specified by ρ) pointing away from the centerline where the algorithm searches for continuations; potential connections with too large costs $C > c_{max}$ are discarded. The final reconstructed trees (after removal of the non-coronary artery trees; see next paragraph) are shown in Figs. 4(b) and 5(e).

Adaptions and Parameters for Coronary Artery Centerline Extraction: The overall procedure as described above, the tube detection, centerline extraction, and tree reconstruction of course also extracts the centerlines of other tubular objects (see Fig. 5). Thus, not just the coronary arteries are identified, but also other structures such as blood vessels in the lung and chest area, the aorta and also some of the bones. These other

Table 1: Summary

Measure	% / mm			score			rank		
	min.	max.	avg.	min.	max.	avg.	min.	max.	avg.
OV	67.2%	100.0%	94.0%	40.4	100.0	56.0	–	–	–
OF	0.0%	100.0%	71.7%	0.0	100.0	56.8	–	–	–
OT	68.8%	100.0%	96.7%	40.2	100.0	67.6	–	–	–
AD	0.28 mm	1.92 mm	0.58 mm	20.4	60.9	30.2	–	–	–
AI	0.25 mm	0.57 mm	0.37 mm	20.9	69.4	31.5	–	–	–
AT	0.26 mm	1.81 mm	0.45 mm	20.7	62.8	30.9	–	–	–
Total							**–**	**–**	**–**

structures have to be removed. The lung tissue is identified and removed based on thresholding ($I(\mathbf{x}) < -700$ HU) and morpological closing using a spherical structuring element with a radius of 10 voxels; see Fig. 5(c) for the TDF response after removal of the lung tissue. Other tree structures not belonging to the coronary artery trees are typically isolated and relatively small; thus, the coronary artery trees can be identified as the largest connected components (concerning their centerline length). Figs. 4(b) and 5(e) show an example of the final coronary artery trees that are extracted as a whole with our approach while the reference only contains four of these vessels as shown in Figs. 4(c) and 5(f).

At the proximal ends of the coronary arteries – the transition to the aorta – the centerlines extracted with our approach are sometimes slightly shorter than the reference centerlines (see Fig. 1). This is due to properties of the tube detection filter as these regions are not really tubular any more. To account for this issue (see the "overlap until first error" measure in Section 2), the proximal ends of the centerlines are extrapolated along a straight line, solving this issue in most datasets, but not all of them.

Most of the above outlined methods are of a generic nature and may also be used in other application domains [2, 1]; for this specific task, following set of parameters was used: $F_{max} = 100$, $t_{high} = 0.5$, $t_{low} = 0.1$, $l_{min} = 10$ voxels, $\rho = 0.7$, $d_{max} = 200$ HU, $c_{max} = 20$. For GVF computation 500 iterations were used [8].

2 Evaluation and Results

Our approach was evaluated on 16 coronary CTA datasets with undisclosed reference centerlines of four coronary arteries per dataset. These datasets were provided by the organizers of the "Coronary Artery Tracking Challenge"[1] [7]. Based on provided reference points the centerlines of the associated arteries were selected from the coronary artery trees – which are extracted as a whole with our approach – and sent to the organizers, which in return provided evaluation results.

The performance measures are grouped into overlap measures and accuracy measures. For the exact description of the performance evalutation and the used measures we refer to [7]. The provided results on the testing datasets are summarized in Tables 1, 2, and 3.

Overlap: The overlap measures are used to assess the ability of the approach to identify the structures of interest (see Tables 1 and 2). Measures are the overlap OV, the overlap until first error OF, and the overlap with clinically relevant parts of the vessels OT. Summarizing, our approach shows a high overlap with the provided reference centerlines: $OV = 94.0\%$, $OF = 71.7\%$, and $OT = 96.7\%$ on average; better than the performance of a human user. The OF of 71.7% on average seems quite low; for some vessels the overlap until first error is 0.0%, while the overlap is still very high ($OV > 99\%$), meaning that just the first few centerline points at the proximal end are assumed to be incorrect. This behaviour is related to properties

[1] see `http://cat08.bigr.nl`

Table 2: Average overlap per dataset

Dataset	OV			OF			OT			Avg.
nr.	%	score	rank	%	score	rank	%	score	rank	rank
8	86.0	51.3	–	57.1	41.4	–	88.6	44.4	–	–
9	96.3	50.7	–	87.4	58.9	–	98.9	61.9	–	–
10	93.1	47.8	–	61.4	31.5	–	94.8	59.9	–	–
11	93.2	47.7	–	54.0	53.5	–	93.8	60.5	–	–
12	92.8	52.0	–	27.1	15.2	–	98.0	61.5	–	–
13	97.2	62.5	–	92.0	71.6	–	98.1	86.6	–	–
14	96.9	56.3	–	67.3	61.8	–	99.8	74.9	–	–
15	97.1	64.6	–	47.3	36.3	–	99.2	62.1	–	–
16	93.8	62.8	–	83.2	65.1	–	98.6	74.3	–	–
17	84.9	50.7	–	34.3	30.2	–	88.9	58.9	–	–
18	92.6	47.1	–	75.9	60.4	–	95.5	64.4	–	–
19	95.4	55.4	–	99.6	92.8	–	99.9	87.4	–	–
20	96.4	82.2	–	90.9	72.5	–	97.3	73.7	–	–
21	95.5	54.2	–	89.8	77.0	–	98.5	62.1	–	–
22	96.6	48.7	–	94.9	72.5	–	99.5	74.8	–	–
23	97.0	61.5	–	84.3	67.6	–	97.1	73.5	–	–
Avg.	**94.0**	**56.0**	**–**	**71.7**	**56.8**	**–**	**96.7**	**67.6**	**–**	**–**

of our tube detection filter as discussed in Section 1. The strategy for performance evaluation just tolerates too long centerlines at the proximal end leading into the aorta (see "Clipping the proximal part" in [7]), but it does not tolerate slightly too short centerlines as delivered by our approach. However, we think that in practice the behaviour of our approach is acceptable.

Accuracy: The used accuracy measures are the average distance AD, the average distance inside the vessel AI, and the average distance for the clinical relevant part of the vessel AT (see Tables 1 and 3). The average accuracy measures achieved with our approach are: $AD = 0.58mm$, $AI = 0.37mm$, and $AT = 0.45mm$, a deviation of typically more than 1.5 voxels. This relatively large error has two main reasons: first, the accuracy measures also incorporate the distances for unidentified centerline points (parts of the vessel that where not found). Thus, the accuracy measures are strongly correlated with the overlap measures; a low overlap automatically implies a large accuracy error (AD and AT). For the AI this correlation is minimal, while this error is still about one voxel ($AI = 0.37mm$). Second, the ridge traversal used for the centerline extraction as described in Section 1 is at most voxel-accurate; also the gap closing that is performed in regions where our tube detection filter does not respond (e.g. at junctions) does not guarantee centered paths. A refinement of the centerline point positions as a post-processing could be used to improve the centerline accuracy of our approach, but we have not consider this possibility in more detail yet.

3 Conclusion

We presented an approach for the extraction of the centerlines of the coronary artery trees from CTA datasets. The approach builds on generic methods for the detection of tubular objects, extraction of their centerlines and grouping of these centerlines into tree structures. The use of the presented tube detection filter is beneficial in proximity of calcifications or other image structures with a similar gray value (see Fig. 1). In combination with the presented ridge traversal that incorporates a hysteresis thresholding, centerlines of the coronary arteries can be extracted also in proximity of calcifications and in case of thin low-contrast vessels. The grouping process also accounts for situations where locally parts of the vessels are indistinguishable from the background; e.g. due to occlusions or imaging artifacts. Thus, our approach enables a reliable extraction of the coronary artery trees centerlines without the need for any user specified seed points. The

Table 3: Average accuracy per dataset

Dataset	AD			AI			AT			Avg.
nr.	mm	score	rank	mm	score	rank	mm	score	rank	rank
8	0.78	35.2	–	0.41	37.8	–	0.59	36.0	–	–
9	0.53	28.4	–	0.31	29.3	–	0.39	29.1	–	–
10	0.45	27.2	–	0.38	28.5	–	0.43	27.8	–	–
11	0.52	32.2	–	0.40	33.6	–	0.51	32.4	–	–
12	0.59	26.1	–	0.37	27.8	–	0.39	27.3	–	–
13	0.42	28.1	–	0.37	28.5	–	0.42	28.6	–	–
14	0.48	34.7	–	0.36	35.5	–	0.36	35.4	–	–
15	0.46	26.8	–	0.39	27.4	–	0.40	27.2	–	–
16	0.66	25.9	–	0.35	27.0	–	0.45	26.8	–	–
17	1.27	42.4	–	0.37	46.8	–	0.89	44.0	–	–
18	0.59	26.9	–	0.34	28.3	–	0.40	27.6	–	–
19	0.63	35.0	–	0.38	36.3	–	0.37	36.3	–	–
20	0.54	31.4	–	0.42	32.3	–	0.50	31.6	–	–
21	0.46	24.8	–	0.33	25.6	–	0.36	25.2	–	–
22	0.48	29.5	–	0.34	30.4	–	0.34	30.5	–	–
23	0.47	28.2	–	0.38	28.8	–	0.47	28.2	–	–
Avg.	**0.58**	**30.2**	**–**	**0.37**	**31.5**	**–**	**0.45**	**30.9**	**–**	**–**

achieved overlap with reference centerlines is high ($OV = 94.0\%$), thus being better than the performance of a human user. The centerlines for the identified parts show an average distance to the references of about one voxel ($AI = 0.37mm$), and may require some further refinement of the centerline point positions as a post-processing step if a higher accuracy is required. Still, as our approach performs well in identification of the vessels (a high overlap), we think that the results are a good starting point for such a refinement and also for other clinically relevant measurements such as tube diameter measurement for assessment of aneurysms/stenoses or severity of adjacent calcifications.

Acknoledgment: This work was supported by the Austrian Science Fund (FWF) under the doctoral program Confluence of Vision and Graphics W1209.

References

[1] C. Bauer and H. Bischof. Extracting curve skeletons from gray value images for virtual endoscopy. In *Proc. of MIAR*, pages 393–402, 2008. (document), 1, 1

[2] C. Bauer and H. Bischof. A novel approach for detection of tubular objects and its application to medical image analysis. In *Proc. of DAGM*, pages 163–172, 2008. (document), 1, 1

[3] T. H. Cormen, C.E. Leiserson, and R. L. Rivest. *Introduction to Algorithms (MIT Electrical Engineering and Computer Science)*. The MIT Press, June 1990. 1

[4] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever. Multiscale vessel enhancement filtering. In *Proc. of MICCAI*, pages 130–137, 1998. 1

[5] C. Kirbas and F. Quek. A review of vessel extraction techniques and algorithms. *ACM Comput. Surv.*, 36(2):81–121, 2004. (document)

[6] G. Lavi, J. Lessick, P.C. Johnson, and D. Khullar. Single-seeded coronary artery tracking in CT angiography. *Nuclear Science Symposium Conference Record, 2004 IEEE*, 5:3308–3311, 2004. (document)

[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*, -(-):–, - -. 2, 2

[8] C. Xu and J. L. Prince. Snakes, shapes, and gradient vector flow. *IEEE Transactions on Image Processing*, 7(3):359–369, Mar 1998. 1, 1