Coronary Centerline Extraction Using Multiple Hypothesis Tracking and Minimal Paths

Release 1.00

Ola Friman, Caroline Kühnel and Heinz-Otto Peitgen

July 8, 2008

Abstract

This paper describes an interactive approach to the identification of coronary arteries in 3D angiography images. The approach is based on a novel multiple hypothesis tracking methodology which is complemented with a standard minimal path search, and it allows for a complete segmentation with little manual labor. When evaluated using the 3D CT angiography data supplied with the MICCAI'08 workshop *3D Segmentation in the Clinic: A Grand Challenge II*, 98% of the target coronary arteries could be segmented in about 5 minutes per data set with the same spatial accuracy achieved in manual segmentations by human experts.

Latest version available at the Insight Journal [http://hdl.handle.net/1926/114]

Contents

1 Introduction

The methodology described in this paper was developed for the MICCAI'08 workshop *3D Segmentation in the Clinic: A Grand Challenge II - Coronary Artery Tracking* [4]. This workshop has the form of a competition where the aim is to locate the centerlines of the coronary arteries as accurately as possible in 3D CT angiography image volumes. Our goal is to provide an implementation with which it is possible to segment the complete coronary artery vessel system, regardless of data quality or artifacts. To this end, an accurate segmentation approach over which the human operator has good interactive control is required. Three different segmentation methods are combined in our implementation:

- Multiple hypothesis vessel tracking
- Minimal paths
- Manual setting of points

The multiple hypothesis vessel tracking [3, 2] is the working horse that identifies the major part (about 90%) of the coronary centerlines with high spatial precision. This recent tracking method is described in Section 3.2. Where the multiple hypothesis tracking terminates prematurely, the user can complete the coronary centerlines by connecting points using a standard minimal paths method based on Fast Marching [1], as described in Section 3.3. As a final resort, should the first two methods fail, can the user manually place points along the vessel centerline. A feature of our implementation is that the coronary centerlines at all times are represented by points in a world coordinate system. That is, at no point do we use a voxelized segmentation mask, as this means a quantization and loss of spatial accuracy, as well as a larger memory footprint.

2 Data

The competition data consist of 8 CT angiography data sets for which manually drawn centerlines are provided and 16 data sets for which only the start (S) and end (E) points of 4 coronary arteries per data set are provided. See [4] for more information on the data. The 8 training data sets were used to tune the parameters of the segmentation algorithms described below. All parameters were then kept fixed when segmenting the 16 competition data sets.

3 Methods

3.1 Image preprocessing

The original 3D CT volumes have an in-plane size of 512×512 voxels. However, an inspection of the image spectra reveals that there is no power in the frequencies above $\pi/2$, meaning that the intrinsic resolution of the images is only 256×256, see Fig. 1a. Therefore, as a first step, the input 3D image volume is downsampled by a factor 2, resulting in faster processing and less memory usage without a corresponding loss in accuracy.

In a second preprocessing step, the images are prepared for segmentation by setting the voxel intensity for lung tissue and vessel calcifications equal to the intensity myocardial tissue t_{myo}. Expressed mathematically,

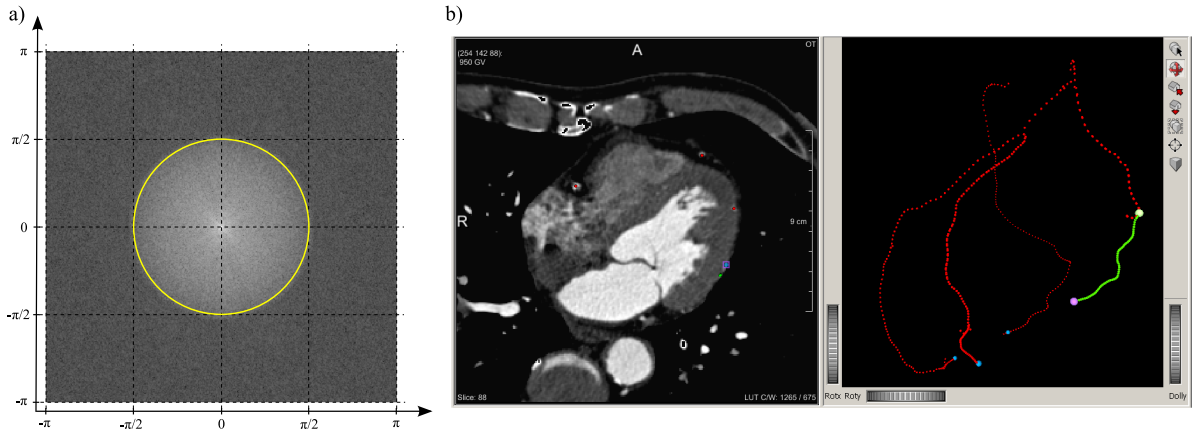

Figure 1: **a)** Log-spectrum of a 512×512 image slice. The spectral power can be found in frequencies less than $\pi/2$, as indicated by the yellow circle. Hence, one can downsample by a factor 2 without loosing information. **b)** The two data viewers used in the segmentation. The left viewer shows axial slices. In this viewer, seed points for the multiple hypothesis tracking are set. The right viewer shows the segmented centerline points. In this viewer, points to be connected with a minimal path algorithm are chosen. The green points illustrates such a path. The data set shown is the competition data set 12.

a preprocessed image $\tilde{I}(\mathbf{x})$ is produced from the original image $I(\mathbf{x})$, where $\mathbf{x} \in \mathbb{R}^3$ is a spatial position, as follows:

$$\tilde{I}(\mathbf{x}) = \begin{cases} t_{myo} & \text{if } I(\mathbf{x}) < t_{myo} \text{ (raise lung tissue),} \\ I(\mathbf{x}) & \text{if } t_{myo} \leq I(\mathbf{x}) \leq t_{calc}, \\ t_{myo} & \text{if } I(\mathbf{x}) > t_{calc} \text{ (remove calcifications).} \end{cases} \quad (1)$$

The rationale for raising the intensity of the lung tissue to the level of the myocardial tissue is to adapt the image to the vessel template model employed for tracking vessels (see Section 3.2) and thereby improve the segmentation of the coronary arteries running along the heart-lung interface. Vessel calcifications are not part of the vessel lumen and they are for this reason also eliminated. The thresholds were fixed to $t_{myo} = 950$ and $t_{calc} = 1700$, expressed in the units of the raw data.

3.2 Multiple hypothesis tracking

A tracking approach to vessel segmentation iteratively places model segments in front of each other so as to form a chain of segments that represents the vessel. In general, a tracking algorithm proceeds by first predicting a new vessel segment from the current position and then updating the model parameters (position, radius, orientation, etc.) based on the observed image data at the new position. Multiple hypothesis tracking (MHT) for vessel segmentation was first presented in [3] and a more detailed description can be found in [2]. The MHT provides a computationally efficient alternative to particle filters for evaluating multiple hypothetical vessel trajectories. From a given point on the vessel centerline, a search tree is built by recursively evaluating several possible vessel continuations, see Fig. 2a for an illustration. Similar strategies are used in game theory and dynamic programming. For the coronary artery segmentation, we search 4 steps forward and the search tree will therefore have a depth of 4. Each leaf in the search tree represents a trajectory consisting of 4 steps, where the length of each step is set to 1.5 times the local vessel radius (in millimeters). Once the search tree has been built, each hypothetical path is evaluated by assessing how well each model segment along the path fits the image data. A step along the most promising path, i.e., the path

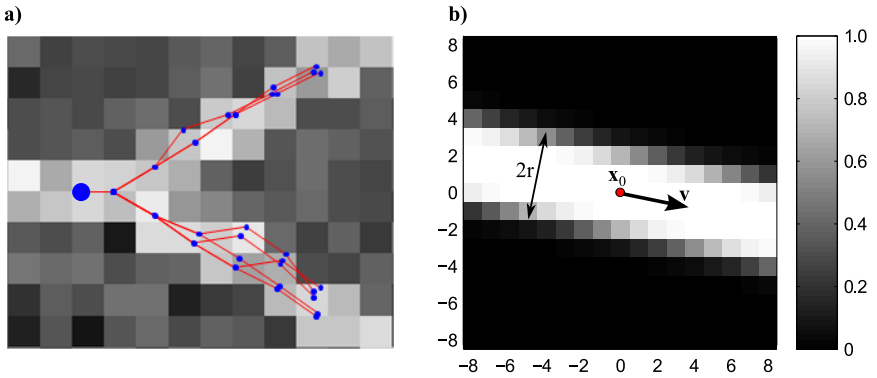

Figure 2: **a)** Illustration of the multiple hypothesis tracking principle. From the start point (large blue point), several hypothetical paths forward are evaluated before deciding on the next step. The search depth is here 6. **b)** A slice through the 3D vessel template function $T(\mathbf{x};\mathbf{x}_0,r,\hat{\mathbf{v}})$ with $\mathbf{x}_0 = \mathbf{0}$, $r = 2.2$ and $\hat{\mathbf{v}} = [0.98,-0.20,0.0]^T$.

with the highest score, is then taken and the search tree is rebuilt from the new position. The advantage of this approach is that the decision of where to take the next tracking step is based on the goodness-of-fit of many steps forward as opposed to just a single step ahead as in a conventional tracking method. This enables the MHT to traverse difficult passages where one or two model segments may fit the data poorly due to low contrast or artifacts.

The MHT approach can be used together with many different vessel models. In this work we use a tube segment model in form of a template $T(\mathbf{x};\mathbf{x}_0,r,\hat{\mathbf{v}}) : \mathbb{R}^3 \rightarrow [0,1]$ that models a small ideal image neighborhood containing a vessel. Here, $\mathbf{x} \in \mathbb{R}^3$ is a spatial coordinate, $\mathbf{x}_0 \in \mathbb{R}^3$ the center point of the template lying on the vessel centerline, $r \in \mathbb{R}$ the local vessel radius and $\hat{\mathbf{v}} \in \mathbb{R}^3$ a normalized vector pointing out the vessel direction. A template example is shown in Fig. 2b and more information on its construction can be found in [3, 2]. Important is that $T(\mathbf{x};\mathbf{x}_0,r,\hat{\mathbf{v}})$ has a closed-form expression and that the partial derivatives with respect to the parameters can be derived analytically. This means that a non-linear least squares approach can be efficiently used to fit the template to the image data, whereby accurate estimates of the local radius, the local vessel direction and a point on the vessel centerline are obtained. The Levenberg-Marquardt algorithm is used for this purpose to adapt the template radius, direction and center point to the image data in a least square error sense. Finally, to reconnect with the preprocessing step, the template model assumes a uniform image intensity around a brighter vessel. Therefore, to conform with this assumption at the heart-lung interface, the lung tissue image intensity is raised to the level of the myocardial tissue in the image preprocessing step.

To start the tracking, initial values of the position ($\mathbf{x}_0$), radius (r) and direction ($\hat{\mathbf{v}}$) parameters are required. In our implementation, the user provides the start point $\mathbf{x}_0$. The Hessian matrix is then analyzed at this point to find the orientation $\hat{\mathbf{v}}$ of the local image neighborhood. The tracking is performed bidirectionally along $\hat{\mathbf{v}}$ and -$\hat{\mathbf{v}}$. The initial radius is set to 0.75 mm. A first fitting step using the Levenberg-Marquardt optimization will refine these initial guesses and the tracking will then commence. Hence, all the user has to do is to give a spatial start point, the tracking then finishes in a few (1-3) seconds. Finally, while the MHT algorithm has the ability to detect vessel bifurcations, this option was turned off as no significant reduction in segmentation time was achieved.

3.3 Minimal paths

A minimal path approach is employed to complement the segmentation provided by the MHT algorithm described above. Typically, a minimal path is here used to bridge gaps the MHT algorithm is unable pass due to low vessel contrast or interfering neighboring structures. A minimal path is the shortest path between two points according to a given metric. The key problem is to specify a metric so that the minimal path coincides with the vessel centerline. Once the metric has been specified, the standard methods for finding the minimal path are to apply either Dijkstra's shortest path search, where the metric is given as weights between the graph vertices, or to apply a Fast Marching originating from the start point and then perform a back-tracking from the end point to find the optimal path [1]. The metric is in this latter case given as a speed image. We adopt the Fast Marching approach and use the implementation in the Insight Segmentation and Registration Toolkit (ITK). The speed image $S(I) \rightarrow [0,1]$ is constructed as a combination of the original image I (the spatial index $\mathbf{x} \in \mathbb{R}^3$ is not written out here) and the Sato et al. [5] vesselness measure $V(I)$ as implemented in ITK. To speed up the processing, the vesselness measure is only calculated for a single scale and for a box around the user-given start and end points. The correct scale of the vesselness measure can be inferred as the MHT algorithm provides us with the radii of the points we want to connect. The speed image is then calculated as $S(I) = 0.25\sigma_1(I) + 0.75\sigma_2(V(I))$, where $\sigma_i(\cdot) : \mathbb{R} \rightarrow [0,1]$ denotes the itkSigmoidImageFilter sigmoid function. The two parameters in $\sigma_1(\cdot)$ and $\sigma_2(\cdot)$ can also be inferred automatically as we have samples lying on the vessel centerline in the vicinity of the start and end points of the minimal path. Hence, the user needs only to mark two points, the remaining parameters are then set up automatically and the minimal path is generally found within 1-2 seconds. An example path is shown in Fig. 1b.

4 Implementation

4.1 GUI & Interaction

All methods were implemented in C++ and integrated as modules in the free software package MeVisLab (http://www.mevislab.de). MeVisLab is a graphical programming environment for prototyping biomedical imaging applications. It embodies, for example, the ITK and the Visualization Toolkit (VTK), as well as many other modules for medical image analysis, interaction and visualization. The modules required for the coronary centerline identification were gathered under a graphical user interface which facilitates the required user interactivity. Two data viewers are used for the segmentation, one showing axial image slices and one 3D viewer displaying segmented points on the artery centerlines, see Fig. 1b. In the image viewer, the user can add and delete points, for example to be used as seeds for the MHT. In the 3D viewer, all points are shown, including the pre-defined end (E) points delivered with the competition data. In this viewer, the user can choose points to connect with the minimal path algorithm.

4.2 Segmentation approach

The segmentation of the coronary arteries was in general carried out as follows. First, the MHT was executed with one seed point in each of the two main coronary arteries exiting from the aorta, i.e., in the right coronary artery (RCA) and in the left main stem (LM) artery. This typically results in a nearly complete segmentation of the RCA and of one of the main left coronary arteries (no attempt to detect bifurcations were made). Next, tracking seed points were placed in the remaining target arteries. If the tracking terminates prematurely, it can be reinitialized by placing further seed points. When the tracking stage is completed, about 90% of the coronary arteries are generally segmented, though there might be gaps in the segmentation where the

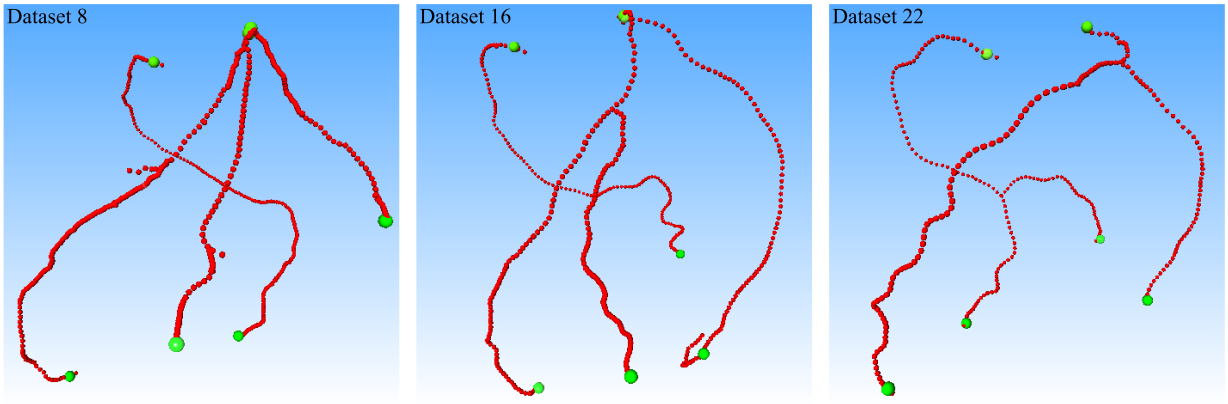

Figure 3: Segmented centerline points for the competition data sets 8, 16 and 22.

vessel contrast is low. To finalize the segmentation, the user selects points in the 3D viewer in Fig. 1b to be connected with the minimal paths algorithm. As our goal is a complete segmentation, a minimal path search was frequently necessary to find the last distal vessel parts to the pre-defined end points. Should the minimal path be incorrect, i.e., not run along the vessel centerline but along some neighboring structure, a last resort is to manually introduce intermediate points on the vessel centerline between which minimal paths are sought.

The final segmentation result is a collection of points, most of which lie on the centerlines of the target coronary arteries. For the competition purpose, the coordinates along the centerline of each target coronary artery must be extracted. To this end, the segmented points are converted into an undirected graph where each point becomes a vertex and where the vertexes are connected by edges with weights equal to the squared Euclidian distance between the corresponding points. The centerlines are found by applying a Dijkstra shortest path search between the pre-defined start (S) and end (E) points in this graph.

5 Results

The 16 competition data sets were segmented using the procedure described above. Results for three data sets are shown in Fig. 3. The number of user-set seed points for the MHT, the number of minimal path connections and the approximate time required for a complete segmentation were recorded for each data set, see Table 1(b). The segmentation time includes the entire process, i.e., from loading data to saving the competition results. An average data set is segmented in 5-6 minutes and requires 6 MHT seed points and 4 minimal path connections. The computational time for the MHT and the minimal paths constitute only a small part of the total segmentation time; most of the time is spent on visual screening and interaction to get the centerline identification correct at one or a few low contrast passages. For example, in data set 22 the MHT finds all target coronary arteries directly and the entire segmentation can be completed in less than 2 minutes. By contrast, data set 8 contains a few difficult passages which required more interaction. The problematic passages were almost exclusively areas of low vessel contrast; calcifications, plaques or stents did in general not cause any problems for the MHT. Again, it should be stressed that our primary goal when segmenting the competition data sets was to maximize the competition scores and not to optimize interaction or speed; a segmentation of clinically relevant parts can be produced with less interactivity and in shorter time.

Accuracy and overlap scores for the segmented vessels were calculated as described in [4]. There are 3

Table 1: User interaction & Average overlap per data set

(a) Average overlap per data set

Dataset	OV			OF			OT			Avg.
nr.	%	score	rank	%	score	rank	%	score	rank	**rank**
8	92.5	74.6	–	78.2	71.5	–	92.7	71.3	–	–
9	100.0	100.0	–	100.0	100.0	–	100.0	100.0	–	–
10	99.4	96.6	–	96.5	86.8	–	99.4	87.2	–	–
11	96.0	63.9	–	53.2	52.4	–	96.0	64.5	–	–
12	99.2	71.6	–	68.4	48.0	–	99.5	62.3	–	–
13	98.6	69.3	–	64.7	47.5	–	98.7	67.9	–	–
14	99.9	95.7	–	83.9	82.6	–	99.9	87.5	–	–
15	99.0	87.0	–	95.1	85.1	–	99.0	87.0	–	–
16	99.3	83.3	–	84.0	78.6	–	99.3	87.2	–	–
17	89.8	82.9	–	64.5	57.8	–	89.7	71.6	–	–
18	98.8	76.0	–	79.6	65.2	–	98.8	74.4	–	–
19	100.0	100.0	–	100.0	100.0	–	100.0	100.0	–	–
20	99.4	92.6	–	93.6	73.8	–	99.5	85.5	–	–
21	100.0	97.8	–	99.9	99.5	–	100.0	100.0	–	–
22	99.9	95.0	–	99.8	87.4	–	100.0	100.0	–	–
23	100.0	100.0	–	100.0	100.0	–	100.0	100.0	–	–
Avg.	**98.2**	**86.6**	**–**	**85.1**	**77.3**	**–**	**98.3**	**84.1**	**–**	**–**

(b) User interaction

# MHT	# Minimal	Time
seeds	**paths**	(min)
7	10	10
6	7	5
8	3	6
7	4	8
4	3	4
9	4	5
6	2	4
5	2	4
5	3	4
8	4	6
6	4	8
6	1	6
5	5	6
8	4	6
4	0	2
5	5	6
6.2	**3.8**	**5:38**

Table 2: Average accuracy per data set

Dataset	AD			AI			AT			Avg.
nr.	mm	score	rank	mm	score	rank	mm	score	rank	**rank**
8	0.39	46.9	–	0.31	47.7	–	0.39	47.9	–	–
9	0.16	51.6	–	0.16	51.6	–	0.16	52.2	–	–
10	0.23	43.5	–	0.23	43.7	–	0.23	43.5	–	–
11	0.33	45.2	–	0.28	45.8	–	0.33	45.2	–	–
12	0.24	47.4	–	0.24	47.7	–	0.24	48.5	–	–
13	0.23	46.4	–	0.22	47.0	–	0.23	47.2	–	–
14	0.24	50.6	–	0.24	50.7	–	0.24	50.4	–	–
15	0.19	51.3	–	0.18	51.8	–	0.19	52.1	–	–
16	0.23	46.7	–	0.22	46.9	–	0.25	45.3	–	–
17	0.75	52.6	–	0.30	53.0	–	0.76	52.6	–	–
18	0.20	51.5	–	0.18	52.0	–	0.20	51.5	–	–
19	0.25	50.8	–	0.25	50.8	–	0.25	50.8	–	–
20	0.30	47.7	–	0.29	47.8	–	0.30	47.7	–	–
21	0.17	49.1	–	0.17	49.1	–	0.16	49.3	–	–
22	0.21	47.1	–	0.21	47.2	–	0.21	47.3	–	–
23	0.23	45.5	–	0.23	45.5	–	0.23	45.5	–	–
Avg.	**0.27**	**48.4**	**–**	**0.23**	**48.6**	**–**	**0.27**	**48.6**	**–**	**–**

Table 3: Summary

Measure	% / mm			score			rank		
	min.	max.	avg.	min.	max.	avg.	min.	max.	avg.
OV	61.3%	100.0%	98.2%	47.4	100.0	86.6	–	–	–
OF	7.7%	100.0%	85.1%	10.4	100.0	77.3	–	–	–
OT	61.0%	100.0%	98.3%	35.9	100.0	84.1	–	–	–
AD	0.10 mm	2.12 mm	0.27 mm	37.7	61.9	48.4	–	–	–
AI	0.10 mm	0.53 mm	0.23 mm	38.1	62.9	48.6	–	–	–
AT	0.10 mm	2.13 mm	0.27 mm	32.7	62.0	48.6	–	–	–
Total							**–**	**–**	**–**

overlap scores: *Overlap* (OV), *Overlap until first error* (OF) and *Overlap with* $>$ 1.5 *mm vessel* (OT). These scores measure the overlap between the segmented centerlines and a ground truth centerline derived from manual segmentations by human experts. The scores are scaled so that 0 indicates complete failure, 50 corresponds to a result within the human inter-observer variability and 100 is a perfect result. The overlap scores for our algorithm and for each of the 16 data sets are presented in Table 1(a). A first observation is the nearly complete segmentation of the target vessels, more than 98% are on average segmented. This is also reflected in the high overlap scores, with an average of 86.6 for the OV score, i.e., significantly better than the human inter-observer variability.

The accuracy scores evaluate the distance to the ground truth centerline. The scores are: *Average distance* (AD), *Average distance inside vessel* (AI) and *Average distance to the clinical relevant part of a vessel* (AT). The distances and the scores are shown in Table 2. The average distance to the ground truth centerline is 0.27 mm which is about half of the (intrinsic) voxel size. This is also the accuracy the human experts achieve (average AD score of about 48.5 compared to the human score of 50). A summary of all scores is presented in Table 3.

6 Conclusions

Our goal in this work is to produce a complete segmentation of the target coronary vessels regardless of data quality or artifacts, and thereby achieve a good competition score. This ambition is reflected in the high overlap scores, but also in the interaction (clicks and connections) and time that were spent to get accurate segmentations also at the very distal parts of the vessels. The distal parts are of lesser importance from a clinical perspective and it is possible to achieve clinically relevant results with less interaction and in less time than in the current contest setting. In terms of centerline accuracy, human expert performance is achieved with our approach. This can largely be attributed to the precision of the multiple hypothesis tracking algorithm which segments the major part of the vessels. One can also remark that the high spatial accuracy was obtained although the data have been subsampled by a factor 2, confirming the observation that the intrinsic resolution is not as high as indicated by the original 512^3 image volumes.

References

[1] T. Deschamps and L. Cohen. Fast extraction of minimal paths in 3D images and application to virtual endoscopy. *Medical Image Analysis*, 5(4):281–299, 2001. 1, 3.3

[2] O. Friman, M. Hindennach, and H.-O. Peitgen. Multiple hypothesis template tracking of 3D vessel structures. *IEEE Transactions on Medical Imaging*, 2008. In review. 1, 3.2, 3.2

[3] O. Friman, M. Hindennach, and H.-O. Peitgen. Template-based multiple hypotheses tracking of small vessels. In *Proceedings of the 5th IEEE International Symposium on Biomedical Imaging (ISBI'08)*, 2008. 1, 3.2, 3.2

[4] 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. In *Medical Image Computing and Computer-Assisted Intervention - MICCAI'08*, 2008. 1, 2, 5

[5] 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 *Proceedings of the First Joint Conference on Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery*, 1997. 3.3