Semi-Automatic Shape Extraction from Tube-like Geometry

Jan Bruijns

Introduction to 3-dimensional x-ray systems

By Rolf Suurmond

Although the principle of using x-rays for diagnostic imaging has been known for a long time, there is still a lot of research and improvements going on in the field. One recent addition to the neuro-vascular imaging system—used for visualizing blood vessels in the brain and guidance of interventions—is 3D rotational angiography (3D-RA). Using 3D-RA, it is possible to view the vessel tree in 3D. Traditionally, 3D information could only be obtained with dedicated tube-shaped MR and CT scanners that are not available in the intervention room. With 3D-RA, the 3D images can be generated in the intervention room, without moving the patient to and from the CT scanner. Also, the spatial resolution is better than that of MR and CT scanners.

Figure 1: Philips Integris BV 5000 neuro-vascular x-ray system.

A picture of a neuro-vascular imaging system is shown in Figure 1. It is a system with two x-ray sources and two x-ray detectors that can be moved around the patient to obtain good viewing directions. With each source-detector pair real-time sequences of 2D projection images of the vessels can be made after injection of a contrast medium. For 3D-RA, only one source-detector pair is used. It is rotated around the patient while making projection images (see Figure 2). After that, the images are processed on a computer workstation to obtain a volumetric dataset.

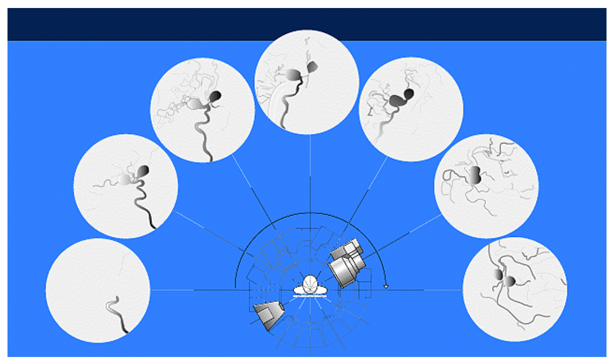

Figure 2: The x-ray source and detector rotate around the patient (center of figure) while making x-ray images.

During an intervention certain malformations can be repaired by means of a catheter. A catheter is a thin plastic hollow tube that can be inserted into the arterial system via a small incision, mostly at the groin. Examples of such malformations are aneurysms and stenoses. An aneurysm is a sac formed by the dilatation of the wall of an artery (see Figure 3 (left)). Viewing the aneurysm from multiple directions on the 3D-RA workstation helps the neuroradiologist to determine the location and the shape of such a malformation. Apart from viewing, it is also possible to measure the size of the aneurysm. The height and width of an aneurysm can be determined by drawing a line in 3D through the aneurysm and displaying the length of the line. A stenosis is a narrowing of the vessel lumen (see Figure 3 (right)). In making correct decisions about the treatment, it is important to know the diameter of the stenotic part of the vessel in relation to the non-stenotic part. Also the length of the stenotic part is relevant to know. Again, these measurements can be done by drawing lines manually, but this is a somewhat cumbersome process. Therefore, we are developing algorithms to automate the process of stenosis analysis. The following article by Jan Bruijns presents the results of this research.

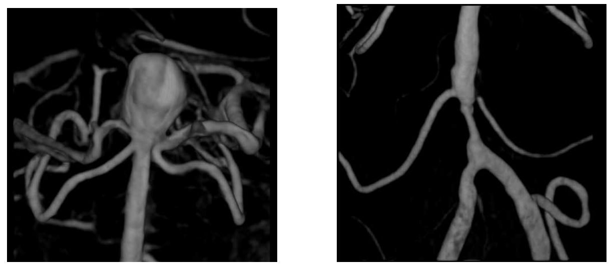

Figure 3: 3D-RA reconstruction of an aneurism (left). 3D-RA reconstruction of a stenosis (right).

Semi-Automatic Shape Extraction from Tube-like Geometry

Jan Bruijns

Introduction

These days, it is possible to acquire volume representations of the brain which show a clear distinction in gray values between tissue and vessel voxels (see [3] for an example). Shape extraction, such as measuring a vessel's diameter, is done by interactively positioning and orienting a plane. The intersection of this plane with the volume gives a 2D image of gray values in which the vessel pixels have a different gray value than the tissue pixels. The diameter of the selected vessel can be computed by selecting the appropriate area, for example, using a rectangle.

The plane should be oriented so that it is orthogonal to the vessel whose shape has to be measured. An oblique plane would give the wrong diameter. Unfortunately, interactively orienting the plane orthogonally to the vessel is a time-consuming and error-prone task. In addition, the plane frequently intersects more than one vessel. This means that the correct vessel in the 2D image also has to be selected interactively. If we consider diagnosing a stenosis, for example, it is therefore impractical to measure a set of diameters along a vessel.

We solved the problem of interactively orienting a plane orthogonally to the vessel by means of a self-adjusting probe (which, to the best of our knowledge, is a new tool). A probe is a combination of a sphere and a plane through the center of the sphere (including storage for some other vessel properties such as the radius of the vessel). After the probe has interactively been placed on a vessel in the neighborhood of the desired position (see figure 8; only the probe's sphere is outlined), the probe automatically adjusts itself so that its plane is orthogonal to the vessel and the center of its sphere is on the central axis of the vessel (see figure 9). The radius of the vessel is also estimated and stored in the probe. When the probe is aligned, it can be moved along the vessel in the direction of the plane normal. The probe aligns itself again after each step. It is also possible to let the probe follow the vessel automatically until the probe detects for example the end of a vessel or the beginning of an aneurysm. An endoview can be displayed while tracing a vessel (see figure 10).

First, we introduce the self-adjusting probe. Next, we describe its application for shape extraction. Finally, we present our results and we conclude the paper.

This paper is an extract of [6]. The interested reader is referred to this last paper for more details about the algorithms involved.

The Self-Adjusting Probe

As mentioned in the introduction, we solved the problem of orienting a plane orthogonally to a vessel by implementing a self-adjusting probe. A probe consists of a sphere:

$$(\boldsymbol{p} - \boldsymbol{p}_0)^2 = r^2 \quad (1)$$

and a plane through the center of the sphere:

$$\boldsymbol{n}_p \cdot (\boldsymbol{p} - \boldsymbol{p}_0) = 0 \quad (2)$$

with $\boldsymbol{p}$ the position of a point, $\boldsymbol{p}_0$ the center and r the radius of the sphere and $\boldsymbol{n}_p$ the normal of the plane.

For reasons that will be explained later, the radius of the sphere should be slightly greater than the radius of the vessel. Interactive selection of the correct radius is easy with a proper graphical user interface, when the 3D sphere is rendered together with a 3D triangle surface representation of the boundary of the vessels, created for example by a marching cubes algorithm [2] as in figure 7.

The sphere should be positioned so that the vessel intersects the sphere, at least partially. Interactive positioning is easy by creating three orthographic views of the volume of the vessel surface (see figure 8) and then placing the cursor in two different views in the neighborhood of the vessel to be measured.

Next, the orientation of the probe's plane, the local co-ordinate system of this plane and the position of the probe's sphere is adjusted. After alignment, the vessel radius is estimated. Finally, we explain how to deal with an initial orientation which deviates substantially from the orthogonal orientation.

Probe's Plane Orientation

In order to adjust the orientation of the plane, we use the normals of the surface vertices inside the sphere, at a small distance from the plane (see figure 4). The threshold for this distance to the plane is taken from the lengths of the triangle edges.

The plane normal should be as orthogonal to the selected surface normals as possible. Therefore, the sum of the squares of the inner products between the plane normal and the n_s selected surface normals:

$$error = \sum_{i=1}^{n\_s} (\boldsymbol{n}_p \cdot \boldsymbol{n}_i) \quad (3)$$

is minimized with the constraint that the plane normal has a length of 1:

$$||\boldsymbol{n}_p|| = 1 \quad (4)$$

Note that there is no need for connectivity information of the 3D triangles. Triangles may even intersect or be disconnected. All that is needed is a set of 3D triangles which together form a reasonable approximation of the boundary of the vessels.

Differentiating equation (3) to the plane normal $\boldsymbol{n}_p$ gives the following homogeneous least square problem:

$$N\boldsymbol{n}_p = \begin{bmatrix} 0\,0\,0 \end{bmatrix}^T \quad (5)$$

with

$$N = \begin{bmatrix} \sum_i \boldsymbol{n}_{i,x}^2 & \sum_i \boldsymbol{n}_{i,x}\boldsymbol{n}_{i,y} & \sum_i \boldsymbol{n}_{i,x}\boldsymbol{n}_{i,z} \\ \sum_i \boldsymbol{n}_{i,x}\boldsymbol{n}_{i,y} & \sum_i \boldsymbol{n}_{i,y}^2 & \sum_i \boldsymbol{n}_{i,y}\boldsymbol{n}_{i,z} \\ \sum_i \boldsymbol{n}_{i,x}\boldsymbol{n}_{i,z} & \sum_i \boldsymbol{n}_{i,y}\boldsymbol{n}_{i,z} & \sum_i \boldsymbol{n}_{i,z}^2 \end{bmatrix} \quad (6)$$

The solution of this homogeneous least square problem, together with the constraint in equation (4), is given by the null space of the matrix N. This null space is computed by means of singular value decomposition (see for example section 2.6 of [4]). If this null space has a dimension of 1, the plane normal is equal to the basis vector of this null space. Else, the selected surface normals do not span a plane (for example, if they have all the same direction).

If the selected surface normals do not define a plane direction, the plane cannot be adjusted so that it is orthogonally to the vessel: the probe cannot align itself to the vessel. Probably the probe does not intersect the vessel enough.

Local Co-ordinate System of the Plane

If the orientation of the plane is adjusted to the vessel (the probe is aligned), the direction of the normal is reversed if the new plane normal is opposite to the old plane normal.

Next, the u and v axis of the plane are chosen so that they are orthogonal to the new plane normal and orthogonal to each other and are as close as possible to the old u and v axes (if any). Consistent u and v axes are necessary for a smooth endoview of the vessel during vessel tracing (given further detail in the section "Shape Extraction").

These u and v axes are also used for the adjustment of the sphere center, as described in the next section.

Probe Position and Vessel Radii

The center of the sphere is adjusted so that it coincides with the central axis of the vessel.

The projection on the aligned plane of the average position of the surface vertices is used as an initial estimate for the new sphere center.

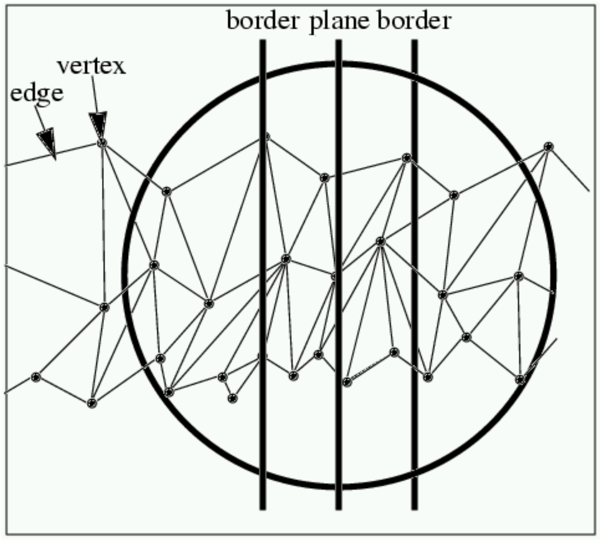

Figure 4: Side view of vertices at a small distance of the plane

Due to outliers (vertices inside the sphere but not belonging to the vessel), it is possible that this new sphere center is not in the neighborhood of the vessel's central axis. This initial estimate is therefore improved by the following iterative algorithm:

1. The selected vertices are projected onto the aligned plane and are subdivided into four sectors, with the new u and v axes as the coordinate system and the current sphere center as the origin (see figure 5):
$$s_1 = [315°, 45°],\ s_2 = [45°, 135°] \quad (7)$$
$$s_3 = [135°, 225°],\ s_4 = [225°, 315°] \quad (8)$$
2. For each sector, the number of projected vertices in that sector (used in vessel tracing discussed later on), the minimum distance and the maximum distance to the origin are computed.
3. The sphere center is moved in the direction of the larger minimum distances.

This shift is repeated until it is negligible or until a maximum number of shifts have been performed.

Finally, the minimum and maximum vessel radii are estimated using the minimum of the four minimum and the minimum of the four maximum distances of the projected vertices to the final sphere center.

Sectors with no projected vertices, for example in case of a lateral aneurysm, are not included in the determination of the minimum of the maximum distances. The minimum radius gives a good estimate for the maximum inside diameter. The quotient between the maximum radius and the minimum radius gives an estimate of the accuracy of the estimated minimum radius.

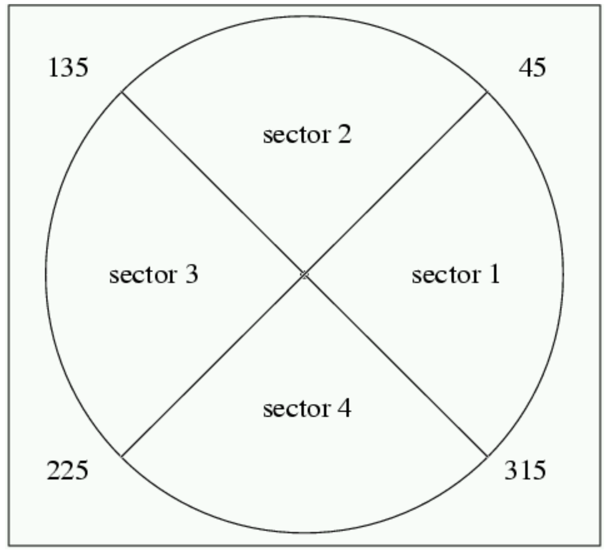

Figure 5: The four sectors

It is also possible to approximate the outline of the vessel by an ellipse or a closed spline instead of a circle. Outliers are automatically ignored when only the selected vertices between the two radii are used.

An example of an aligned probe is depicted in figure 9 and an example of an endoview is depicted in figure 10. A slice with the selected vertices (black dots), the final center position (small black square) and the two black radii are depicted in figure 11. This figure also shows that the neighboring vessel is correctly ignored.

Second Alignment

If the probe was not aligned before, it aligns itself twice. Twice because the initial orientation of a plane of a non-aligned probe can be completely wrong. Therefore, the surface vertices close to the plane in this initial orientation can differ very much from the surface vertices close to the plane of the aligned probe. The second alignment corrects a possible residue error in the position of the probe and/or in the orientation of its plane.

In addition, some decisions (like the "end of vessel detected" decision of the shape extraction functions described in the section "Shape Extraction") are based on the ratio between the selected vertices at the front side of the plane and the total number of selected vertices. However, whether a vertex is at the front side or at the back side is decided by the plane orientation before the alignment. So, if the plane orientation has changed a lot, many vertices are at the other side of the aligned plane. Moreover, changing the plane orientation can also change the set of selected vertices.

Therefore, to guarantee proper alignment, the probe aligns itself twice if it was not already aligned. Likewise, when the angle with the old plane normal is more than 30 degrees, the alignment is repeated after a move.

Shape Extraction

We have developed a set of semi-automatic shape extraction functions based on the self- adjusting probe. The first one is 'single step vessel tracing', with automatic alignment afterwards. This function is suitable for interactive exploration of the vessel shape.

The other two shape extraction functions are 'continual vessel tracing' and 'vessel tracing to a second probe'. These two functions activate the single step vessel tracing function repeatedly while collecting the shape parameters.

Single Step Vessel Tracing

The single step vessel tracing function enables interactive exploration of the vessel shape. First, the probe is positioned in the neighborhood of the vessel where the tracing should start. The probe aligns itself twice (as explained in the section "Second Alignment").

After successful alignment, the probe moves one step (the step size can and should be chosen before this function is activated) in the direction of the normal. If the probe goes in the wrong direction, the direction of the normal should be reversed and the single step vessel tracing function should be activated again.

If the move would result in a position outside the volume, the intersection(s) with the corresponding bounding plane(s) is used as the end position of the move: the position is clipped.

Finally, the probe aligns itself at the new position. After aligning, a possible end of vessel is detected by comparing the number of selected vertices in front of the plane with the total number of selected vertices.

If the position was clipped or if an end of vessel was detected, the plane normal and u axis are reversed.

The result can be inspected using a normal view as in figure 9, an endoview as in figure 10 or a slice view as in figure 11.

Continual Vessel Tracing

An application of the single step vessel tracing is continual vessel tracing. First, the probe is positioned in the neighborhood of the vessel where the tracing should start. The probe aligns itself twice.

Next, the single step vessel tracing function is activated to check whether the probe goes in the correct direction. If not, the direction of the normal should be reversed.

Subsequently, the continual vessel tracing function can be activated. The probe moves in the direction of the normal and aligns itself repeatedly. The probe position and the plane normal can be used to create an endoview while tracing (see figure 10).

During tracing, the probe data (not only the position of the sphere and the orientation of the plane but also other vessel properties such as the radius of the vessel) is collected in a tube (an array of probes), giving an approximate shape of the vessel. This can be used to compute the vessel volume, for example. An example of a tube is depicted in figure 12. The center line gives the estimated center positions. The concentric circles are the circles with the minimum and the maximum radius, respectively.

The continual vessel tracing function stops when one of the following conditions is detected:

1. The probe has moved a maximum number of steps.
2. The tube is full.
3. The probe could no longer be aligned.
4. An end of the vessel is detected.
5. An open vessel is detected: the vessel seems to

have a hole in the cylinder-like wall.

An open vessel (for example the beginning of an aneurysm) is detected by taking the *inclusive or* of the results of two rules. The first rule compares the number of selected vertices in front of the plane with the number of selected vertices behind the plane, per sector.

The number of selected vertices in front of and behind the plane are computed during the adjustment of the probe position. In fact, the total number of vertices are computed from these numbers.

The second rule compares the number of selected vertices in front of the plane with those of all other sectors, per sector.

6. The distance to the previous probe position is too small.
7. The distance of the previous probe position to the new plane is too small.

Vessel Tracing to a Second Probe

A variant for continual vessel tracing, as described above, is vessel tracing to a second probe. Both probes should be aligned. Tracing stops when the first probe has passed the plane of the second probe, or when a special stop condition is detected. The plane normal of the first probe should be so that the first probe stepping along the vessel in the direction of its plane normal should arrive at the second probe. If not, the plane normal of the first probe should be reversed before this function is activated.

The special stop conditions are:

1. The probe has moved a maximum number of steps.
2. The tube is full.
3. The probe could no longer be aligned.
4. An end of the vessel is detected.

Comparing these special stop conditions with the stop conditions for continual vessel tracing reveals that vessel tracing to a second probe avoids unwanted stopping, for example in the case of an open vessel condition caused by small surface irregularities.

The "face to face" function facilitates vessel tracing to a second probe by setting the plane normals of two probes such that the center of the first probe is at the positive side of the plane of the second probe, and vice versa. This function can be used to correct the plane normals before vessel tracing to a second probe is activated, but only when the intermediate vessel lies completely between the centers of both probes.

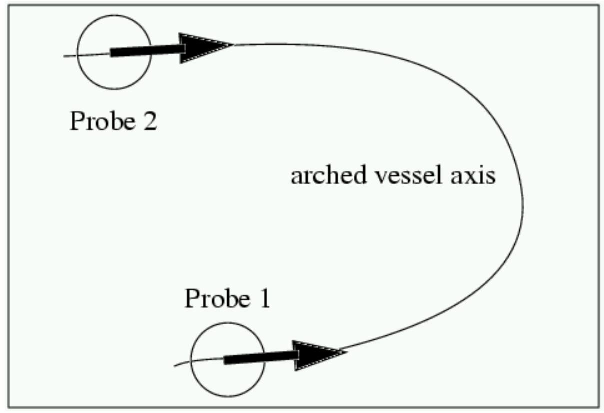

Figure 6: Two probes on a arched vessel

This "face to face" function is also needed to reliable detect whether the first probe has passed the second probe. This is detected by testing whether the center of the first probe is at the negative side of the plane of the second probe. The two probes must be "face to face" to measure this. But if for example the second probe is to the left of the first probe but the vessel first goes to the right of the first probe and then bends to the left in the direction of the second probe (see figure 6), the two probes should not be "face to face" at the start. Also, the pass test should not be applied as long as the first probe is moving away from the second probe. Therefore, the "face to face" function is applied when the two probes overlap for the first time. Next, passing is checked after each step. The resulting algorithm (without the other special stop conditions) is:

1. While the two probe spheres do not overlap, trace the vessel.
2. Apply the "face to face" function.
3. While the first probe is at the positive side of the plane of the second probe and the distance between the two probe's centers decreases, trace the vessel.

The extra condition in the last step is needed to stop vessel tracing in case the second probe is at a branch. Indeed, in this case it is possible that the first probe exhibits a random walk without passing the plane of the second probe because the normals at a branch are pointing in all directions.

Results

The self-adjusting probe and the shape extraction functions have been incorporated in a demo program for testing. Four gray value voxel volumes, acquired with the 3D Integris [1, 3], are transformed to 3D triangle surface representations of the boundary of the vessels using a marching cubes algorithm [2] as shown in figure 7. The number of vertices varied between 123681 and 180435 vertices.

Probe aligning at a vessel was never a problem, even when the probe's plane was almost parallel to the vessel before alignment. After one or two moves, the plane flipped to the orthogonal orientation.

The continual vessel tracing function performed well, even for the noisy 3D triangle surface representations. But, often an open vessel was not detected at a branch or the beginning of an aneurysm. In addition, this function sometimes stopped too early in case of an open vessel condition caused by small surface irregularities. Vessel tracing to a second probe avoids this problem. Using a begin and end probe also gives precise control over the vessel part from which the shape has to be extracted.

Vessel tracing takes about one second per step on an Onyx/Rea2. Rendering a slice (see figure 11) instead of the whole surface (see figure 9) hardly accelerated vessel tracing. So, this one second per step is required for moving and aligning the probe, not for rendering the 3D triangle surface representation. Space partitioning the vertices will probably accelerate vessel tracing considerable, but, for our purposes, accuracy of the extracted shape parameters is more important than speed. Besides, we have implemented a simple and fast method to extract a part of the 3D mesh. A mesh of about 10000 vertices takes about 40 milliseconds per step on an Onyx/Rea2.

The noise in the extracted shape parameters depends on the noise in the 3D triangle surface representation of the boundary of the vessels. An example of the quality obtained is shown in figure 12. Note that the noise can be further reduced by filtering the collected shape parameters.

The self-adjusting probe and the shape extraction functions have been incorporated into a prototype of the 3D Rotational Angiography Integris. This prototype was demonstrated on the LINC'99 in Paris and the RSNA'99 in Chicago. A number of these prototype systems are now being evaluated in the daily clinical routine.

Conclusions

The following conclusions can be drawn from the results, the pictures and the experience gained:

1. The self-adjusting moving probe is a very suitable tool for semi-automatic shape extraction. It makes shape (center positions and radii) extraction of vessels much easier to use. The possible endoview also gives a clear visualization of the inside shape of the vessel.

   The extracted shape parameters, especially the positions of the central axis, can be used to compute the curvature and tortuosity of the vessel part.
2. The accuracy of the extracted shape parameters depends on the quality of the 3D triangle surface representation of the boundary of the vessels. The classification of the voxels into tissue and vessel voxels should therefore be as reliable as possible.
3. Connectivity information of the 3D triangles is not needed. All that is needed is a set of possibly intersecting and/or possibly disconnected 3D triangles which together form a reasonable approximation of the boundary of the vessels.
4. It is also possible to use the normalized gradient vectors and the positions of the boundary vessel voxels. These boundary vessel voxels can be found easily: a vessel voxel is a boundary vessel voxel if and only if at least one of its neighboring voxels is a tissue voxel.
5. this case, the threshold for selection should be based on the voxel size.
6. Further investigation is necessary to improve the continual vessel tracing function, not only for improved detection of the connections of vessels to an aneurysm, but also to be able to detect the vessel structures at branches. Automatic reconstruction of branch structures, as described in [5], is a necessary (and maybe sufficient) con-

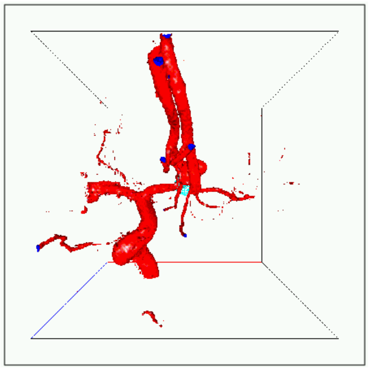

Figure 7: Triangle surface geometry

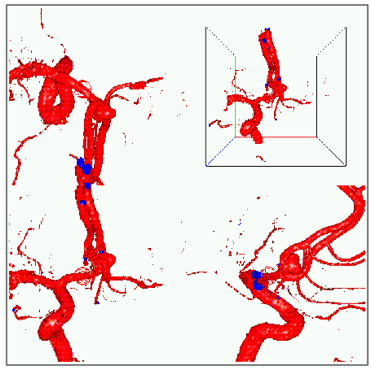

Figure 8: Three orthogonal views

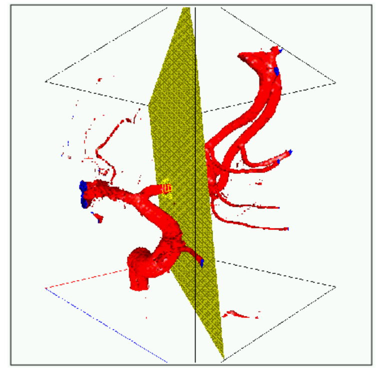

Figure 9: Aligned probe

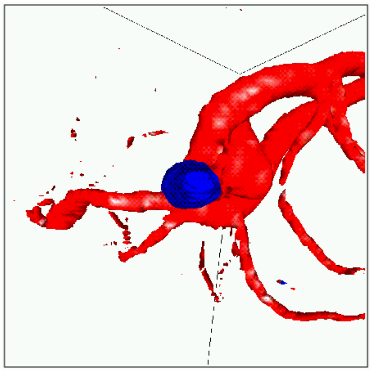

Figure 10: Endoview of a vessel

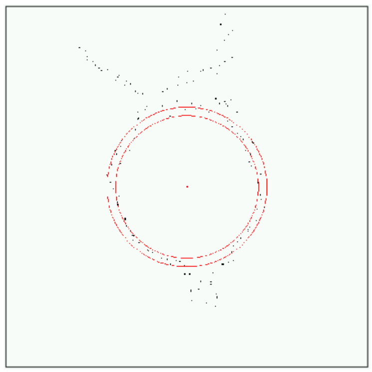

Figure 11: Slice view of a vessel

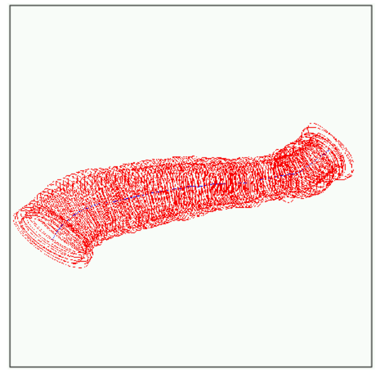

Figure 12: Vessel tube

dition for automatic extraction of the complete vessel structure and its shape parameters.

References

[1] R. Kemkers, J. Op de Beek, H. Aerts, R. Koppe, E. Klotz, M. Grasse, J. Moret *3D-Rotational Angiography: First Clinical Application with use of a Standard Philips C-Arm System* in Proceedings of CAR '98, 1998.

[2] William E. Lorensen, Harvey E. Cline *Marching Cubes: A High Resolution 3D Surface Construction Algorithm* Computer Graphics, Vol. 21, No. 4, July 1987.

[3] J. Moret, R. Kemkers, J. Op de Beek, R. Koppe, E. Klotz and M. Grass *3D rotational angiography: Clinical value in endovascular treatment* Medicamundi, Volume. 42, Issue 3, November 1998.

[4] William H. Press, Saul A. Teukolsky, William T. Vetterling, Brian P. Flannery *Numerical Recipes in C. The Art of Scientific Computing. Second Edition* Cambridge University Press, 1992.

[5] C. Zahlten, H. Juergens, H.-O. Peitgen *Reconstruction of Branching Blood Vessels From CT-Data* in Proc. of the Eurographics Workshop on Visualization in Scientific Computing, Rostock, June 1994.

[6] J. Bruijns *Semi-Automatic Shape Extraction from Tube-like Geometry* in Proc. Vision Modeling and Visualization 2000, Saarbruecken Germany, pp. 347–355, November 200