

CoWRadar: Visual Quantification of the Circle of Willis in Stroke Patients

Haichao Miao¹, Gabriel Mistelbauer¹, Christian Našel² and M. Eduard Gröller^{1,3}

¹Technische Universität Wien, Austria

²Karl Landsteiner University of Health Sciences, Austria

³VRVis Research Center, Austria

Abstract

This paper presents a method for the visual quantification of cerebral arteries, known as the Circle of Willis (CoW). The CoW is an arterial structure that is responsible for the brain's blood supply. Dysfunctions of this arterial circle can lead to strokes. The diagnosis relies on the radiologist's expertise and the software tools used. These tools consist of very basic display methods of the volumetric data without support of advanced technologies in medical image processing and visualization. The goal of this paper is to create an automated method for the standardized description of cerebral arteries in stroke patients in order to provide an overview of the CoW's configuration. This novel display provides visual indications of problematic areas as well as straightforward comparisons between multiple patients. Additionally, we offer a pipeline for extracting the CoW from Time-of-Flight Magnetic Resonance Angiography (TOF-MRA) data sets. An enumeration technique for the labeling of the arterial segments is therefore suggested. We also propose a method for detecting the CoW's main supplying arteries by analyzing the coronal, sagittal and transverse image planes of the data sets. We evaluated the feasibility of our visual quantification approach in a study of 63 TOF-MRA data sets and compared our findings to those of three radiologists. The obtained results demonstrate that our proposed techniques are effective in detecting the arteries of the CoW.

Categories and Subject Descriptors (according to ACM CCS): I.4.0 [Image Processing and Computer Vision]: General—Image processing software; I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms; I.3.4 [Computer Graphics]: Graphics Utilities—Software support; J.3 [Computer Applications]: Life and Medical Sciences—Health

1. Introduction

The human brain is a very delicate structure that is highly dependent on a well-functioning blood supply. A vascular disease in the brain can lead to a stroke, which is the second most common cause of death and the major cause of acquired disability in the developed world [The15b]. Stroke treatment relies on the application of imaging techniques and the investigations done by the radiologists. The focus lies hereby on the arterial blood supply, which is guaranteed by the Circle of Willis (CoW). This arterial circle is depicted in Figure 1. The radiologist inspects the TOF-MRA data set (later simply referred to as data set) using traditional display methods, such as slice-by-slice views and Maximum Intensity Projection (MIP) to identify the cause of a dysfunction. This diagnostic process is complicated by its urgent nature and the time-constraints of the stroke treatment. However, the process

is time-consuming and every patient is viewed as a case of its own without considering preexistent cases. Side-by-side comparisons are impeded by the traditional display methods.

3D Time-of-Flight Magnetic Resonance Angiography (TOF-MRA) is typically used to acquire high-resolution data sets from stroke patients. This data acquisition modality is always applied during routine as well as acute cases if time is sufficiently available, due to its relatively long acquisition time. In contrast to conventional Magnetic Resonance Angiography (MRA), the magnetization is applied outside the scanned object. Consequently, the inflowing blood exhibits a higher magnetization than the saturated stationary tissue, resulting in a significantly higher signal for blood vessels. This method has a high signal-to-noise ratio but turbulent or slow flow can cause signal loss from the vessel. As a result, these vessels are represented by only low intensity values.

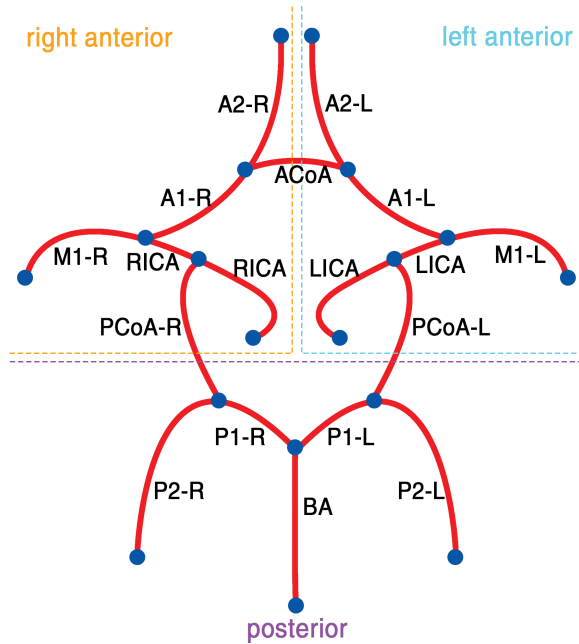


Figure 1: A schematic representation of the standard configuration CoW with anatomically-correct labeled arteries and a subdivision into three subtrees.

As the CoW exhibits a high variability across different patients, its shape and topology are highly relevant for stroke treatment. To support stroke assessment in a timely manner, we propose a standardized visualization of the CoW, rectifying the above shortcomings. The main contributions of our work are:

- Visual indication of problematic areas of the CoW.
- A standardized visualization of the CoW.
- Efficient comparison of multiple patients.
- Systematic description of the CoW.
- A fully automatic CoW extraction pipeline.

2. Related Work

In this paper we use common image processing methods but also introduce novel methods that are required for the visual quantification of the CoW. The standardized visualization of the CoW relies on a vascular model (we call it vessel graph) that has to be extracted first. The most closely related methods and algorithms for the proposed pipeline, as shown in Figure 2, are described in this section. A general introduction to medical image processing and visualization is given by Preim and Botha [PB13]. The CoW received much attention recently. Hartkamp and Van der Grond [HvdG00] investigated morphological variations of the CoW using MRA. The standard configuration of the CoW is given in Figure 1. The authors describe different variations of the standard config-

uration. Bullitt et al. [BGP*03] evaluated different types of the tortuosity metric on intracerebral arteries to discriminate normal from abnormal vessels. Kirbas and Quek [KQ03] reviewed vessel extraction methods targeted at neurovascular structures. They divided the algorithms into six main categories: pattern-recognition techniques, model-based approaches, tracking-based approaches, artificial intelligence-based approaches, neural network-based approaches, and miscellaneous tube-like object detection approaches.

Pock described in his diploma thesis [Poc04] an automated segmentation method using a level-set technique for tubular structures. His work also includes a skeletonization and graph construction approach. We use a segmentation approach that is based on the hysteresis thresholding method by Canny [Can86]. Hysteresis thresholding requires two parameters: A high and a low threshold value, which can be estimated by histogram analysis, as described by Conrache and Aach [CA05].

The CoW is supplied by three main arteries and can be divided into three subtrees, each responsible for the blood transport to a separate area of the brain. The segmentation result has to be separated into three clusters, in order to reflect this natural division of the CoW. In the work of Bullitt et al. [BMJ*05] the authors described a method that is based on the subdivision of the intracranial circulation into four vessel clusters.

Vascular structures are commonly modeled as a vascular graph with edges and nodes that represent segments and branching points. The centerline extraction or skeletonization is a first step in vessel modeling. A skeletonization approach based on topological thinning was introduced by Lee et al. [LKC94]. A detailed vessel model is described by Mistelbauer [Mis13], which is also used in this work.

An automated labeling approach of the CoW is proposed by Bogunovic [Bog12]. His approach is concerned with identifying the anatomically correct names of the bifurcations by using a maximum a posteriori estimation. He evaluated his automated approach on a set of 50 images of healthy patients and reported to have labeled 60% of the cases entirely correct. However, his approach favors sensitivity over specificity and rather tends to find a false bifurcation than to miss one, which is a potential weak point.

The CoW can be described by multiple trees (i.e. one tree per supplying artery) that are connected at their leaves. A radial graph layout is a fitting choice to display the branching structure and topology of the CoW comprehensibly. Furthermore, the radial graph is also motivated by the circular layout of the CoW. Draper et al. [DLR09] described radial design methods as visualizations that arrange data in an elliptical fashion and identified different design patterns.

The main purpose of our visual quantification of the CoW is to create a representation to facilitate human recognition. Beside the radial graph, there are other approaches that could

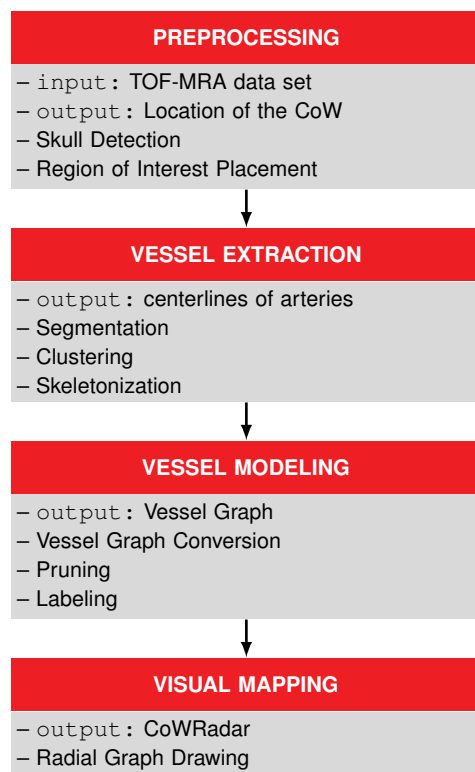


Figure 2: Our fully automatic pipeline for visual quantification of the CoW. It consists of a chain of methods that are subsequently executed. First the ROI is placed inside the data set to cover a small area around the CoW. Then the arteries are segmented and the centerlines are extracted. Next, we create a vessel graph that is visually mapped to a radial graph layout in order to display the CoW in a standardized manner.

be used for the visualization of the arterial circle. Information visualization is the field that investigates the computer-aided creation of abstract representations of otherwise complex data. It takes the human perception into consideration and describes different tools for the visual representation that allows the viewer to understand complex or large amounts of data. The geometrical properties of the CoW can be visualized by parallel coordinates, which are described by Inselberg and Dimsdale [ID90]. Another method is a radar chart that displays multivariate data on different axes aligned in a circular fashion. To the best of our knowledge, the visual quantification of the CoW has not been attempted before.

3. Medical Background

The human brain has a high oxygen and nutrition demand and is, therefore, critically dependent on a well-functioning blood supply. The CoW is responsible for blood circulation to the brain. An ongoing reduction can lead to a stroke, which is a medical emergency possibly causing permanent neurological

damage or even death. Radiologists are therefore interested in the blood supply of the brain tissue and, consequently, in the CoW. The identification of the collateral blood flow is especially important for stroke treatment. Investigating the medical data sets, doctors can find out the source of a stroke and determine the treatment.

The CoW can be naturally separated into three parts that we refer to as *subtrees*. See Figure 1 for the standard configuration of the CoW. Each subtree is primarily supplied by one main artery. The Left Internal Carotid Artery (LICA) supplies the left anterior subtree, the The Right Internal Carotid Artery (RICA) supplies the right anterior subtree and the Basilar Artery (BA) supplies the posterior subtree. Other important arteries are the A1 segment of the Anterior Cerebral Artery (A1), the M1 segment of the Middle Cerebral Arteries (M1) and the P1 segment of the Posterior Cerebral Artery (P1). The subtrees are connected into a circle by the communicating arteries in order to create a collateral blood supply. The anterior subtrees are connected by the Anterior Communicating Artery (ACoA) and the posterior subtree is connected to the anterior subtrees by the left and right Posterior Communicating Artery (PCoA). The communicating arteries (ACoA and PCoA) provide valuable information about the collateral blood circulation and the main arteries ensure the supply of a major part of the brain.

4. Methodology

We propose the visual indication of problematic areas and a simple standardized display of a patient's CoW. In order to alleviate radiologists from tediously inspecting the entire data set, we propose a fully automated pipeline that processes TOF-MRA data sets to create appropriate standardized visualizations of the CoW (see Figure 2). We developed a software solution for the clinical practice that does not introduce additional working steps.

First and in a preprocessing step, we define the Region of Interest (ROI) and extract the main arteries of the CoW. In a second step, we convert the intensity-based representation of the blood vessels into a graph structure. Our work introduces a novel method for the systematic labeling of this vessel graph. Finally, we visually map the vessel graph to a radial layout, called CoWRadar. It features a standardized visualization of the CoW and offers a visual indication of problematic areas as well as a comparison across multiple patients.

4.1. Preprocessing

First, we need to preprocess the TOF-MRA data sets in order to extract the blood vessel of the CoW. One problem is that the intensity range of the skull overlaps with the intensity values of arteries, which leads to problems when using hysteresis thresholding for segmentation. Initially, we segment the skull by shooting rays from the outside of the data sets to

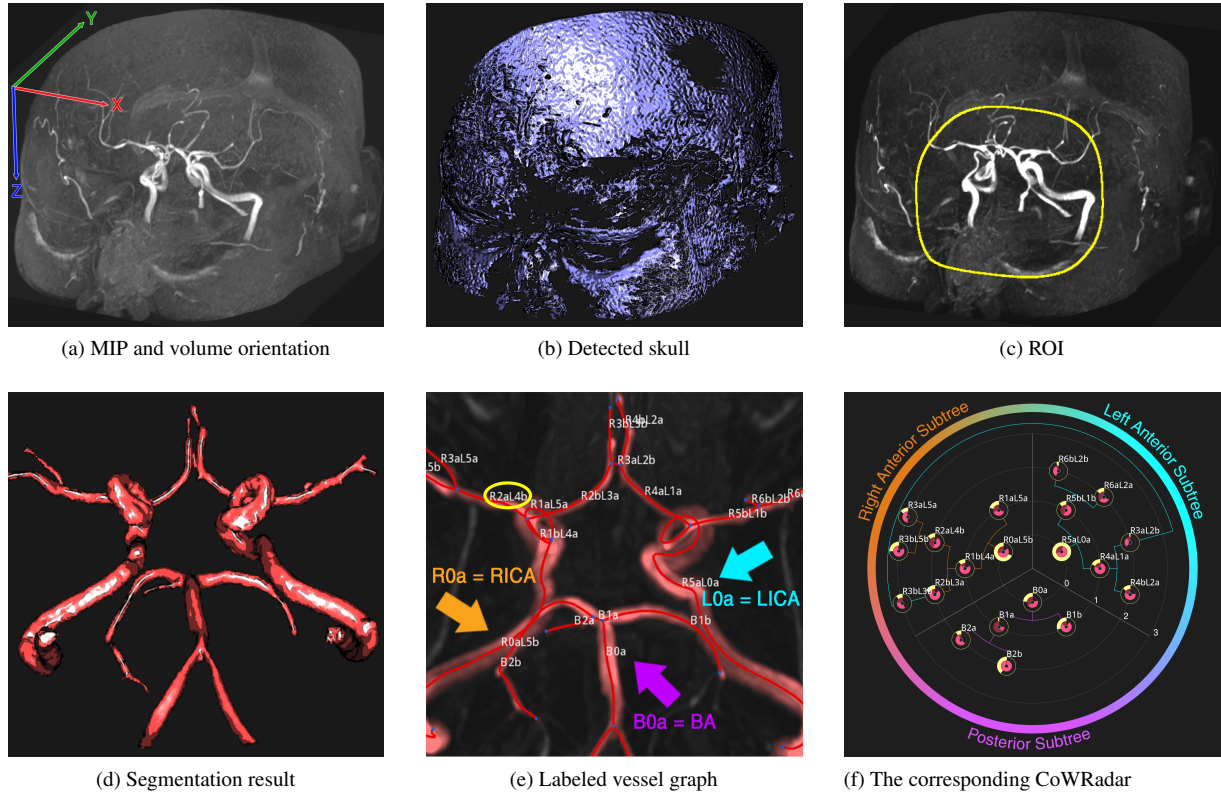


Figure 3: The figures display the intermediate results of our proposed pipeline. (a) displays the volume orientation. (b) shows the segmented skull. (c) displays the defined ROI. The segmented arteries inside the ROI are shown in (d). (e) displays the labeled vessel graph and (f) shows the corresponding CoWRadar.

the inside. The first intensity peak along these rays identifies the bone, which is then removed (see Figure 3b).

Then, we place a ROI in the shape of a super-ellipsoid, including the CoW (see Figure 3c). The location and size of the ROI are based on an empirically defined ratio determined from the skull. The center position (c_x, c_y, c_z) of the skull-voxels is computed as the mean of the skull-voxel positions. The center of the ROI (p_x, p_y, p_z) is shifted from (c_x, c_y, c_z) by using the following empirically motivated factors: $p_x = c_x \cdot 1.02$, $p_y = c_y \cdot 0.78$ and $p_z = c_z \cdot 1.2$.

This heuristic guarantees the placement of the ROI in close proximity to all arteries of the CoW. Next, the shape of ROI has to be adjusted. The goal is to cover the entire CoW, but not too much of the surrounding tissue. This is done by setting the semi-axes A , B and C of the suggested super-ellipsoid. For this purpose, the average Euclidean distance r between (c_x, c_y, c_z) and the skull-voxel positions are computed. r is thereby regarded as an approximative radius of the skull. The semi-axes are then empirically calculated as follows: $A = r \cdot 0.52$, $B = r \cdot 0.46$ and $C = Z \cdot 0.52$, where Z defines the extent of the data set along the z -axis. Since the skull

is not entirely covered by the scan, the C -parameter of the super-ellipsoid is adjusted by using Z instead of r .

The detected ROI covering the CoW specifies the scope of our subsequent approach. By adjusting the axes-lengths, size and roundedness of the super-ellipsoid the differently shaped CoWs can be covered.

4.2. Vessel Extraction

The arteries are represented by voxels and their intensity values inside the ROI. In order to create a vascular model, we distinguish between artery and background voxels. We extract the vessel centerlines, which are good abstractions for the arteries due to their tubular structure.

The arteries are segmented using hysteresis thresholding and the result is shown in Figure 3d. The segmentation result is then partitioned into three clusters in order to reflect the natural subdivision of the CoW into three subtrees. We thereby project the segmentation result along the x -, y - and z -axis, reducing the data sets to three image planes, which are shown in Figure 4. The orange part is associated with

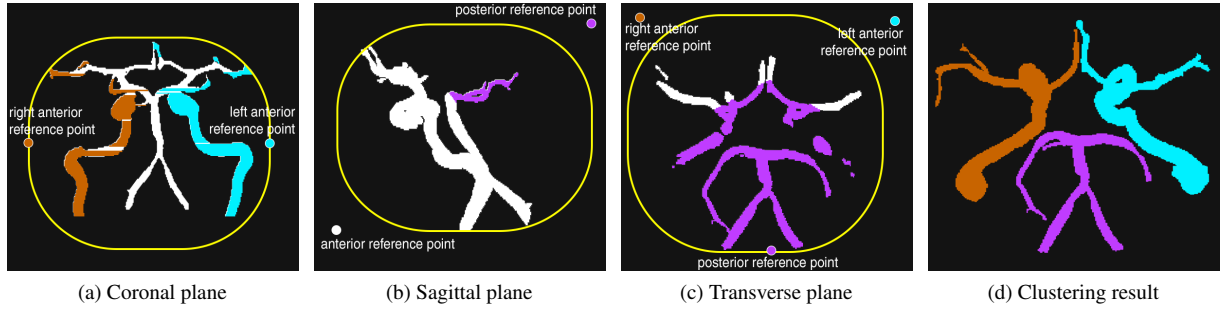


Figure 4: The three image planes of the segmentation result. (a) displays the coronal plane, where the left and right anterior parts are marked. (b) displays the sagittal plane, which is used to mark the P1 segments in the posterior subtree. (c) shows the transverse plane where the posterior part is marked. (d) shows the result of the clustering.

the right anterior subtree, the cyan part with the left anterior subtree and the violet part is associated with the posterior subtree. The separation into clusters is performed in these three planes by using reference points, which are based on the shape and location of the super-ellipsoidal ROI.

p_x , p_y and p_z describe the respective coordinates of the super-ellipsoid's center and A , B and C describe the semi-axes of the super-ellipsoid in the x , y and z direction. The pixels in the image planes are assigned to the closest reference point. In the coronal plane the right and left anterior reference points are at $(p_x + A, p_z)$ and $(p_x - A, p_z)$. Using the image planes, a rough estimation of the BA's location can be done. As the coronal plane in Figure 4a shows, the BA is normally in between the two carotid arteries. This circumstance is exploited to get a better approximation of the location of the BA on the x -axis. For this reason, a ray casting method is additionally applied. Rays are shot from left and right. After the first object is hit, the ray terminates if it hits a background point. Finally, the center c_{BA} of the left (cyan) and right (orange) areas are calculated as an approximation of the BA's location on the x -axis. The sagittal plane is used to assign the P1 segments to the posterior cluster. The reference points are at $(p_y + B, p_z - C)$ and $(p_y - B, p_z + C)$ and the result is shown in Figure 4b. In the transverse plane the posterior reference point is defined by $(c_{BA}, p_y + B)$, the left anterior reference point is at $(c_{BA} + A, p_y - B)$ and the right anterior reference point is at $(c_{BA} - A, p_y - B)$. The result is shown in Figure 4c.

The segmentation result is then clustered according to the results from these image planes. However, some small parts are incorrectly assigned. To address this issue, we apply the following approach: We determine the largest connected components of each of the three parts and set them as the initial areas. Then, these initial areas grow in a breadth-first manner and absorb parts that do not belong to another initial area. The result is displayed in Figure 4d, which shows the correct separation of the segmentation result into the three subtrees.




Figure 5: The detection of an absent main artery using the coronal plane. (a) shows the transverse MIP with an arrow pointing to the location of the missing LICA. (b) shows that the cyan area is much smaller than the orange area. (c) displays the CoW separated into two clusters instead of three.

This approach allows us to quickly spot the absence of a main artery. In such a case, a major part of the brain is under-supplied with blood, which is a highly relevant information for the domain experts. The main arteries are the largest arteries of the CoW. Their presence can be detected by comparing the different (colored) parts in the image planes with each other. The transverse MIP in Figure 5a shows an example data set with a missing LICA according to the domain expert. If the left or right part in the coronal plane is relatively small, then we assume that a main artery is missing, as illustrated in Figure 5b. Compared to the example in Figure 4a, this CoW is not symmetric since the LICA is missing. The area of the cyan left anterior part is smaller than the area of the orange right anterior part as shown in Figure 5b. As a result, the CoW is separated into two clusters instead of three in order to reflect this irregular blood circulation (see Figure 5c).

In order to provide a sufficiently smooth segmentation for the subsequently performed skeletonization, we apply morphological operations such as closing. We use the skeletonization approach by Lee et al. [LKC94] to extract the centerlines of the arteries of the CoW.

4.3. Vessel Modeling

Until now, the arteries are still represented by voxel intensity values. In this section we describe methods to create a model that represents the arteries and their branching points. This model is then used as the basis for the visualization.

The centerlines are converted into a graph representation of the vascular system using the rules of Pock [Poc04]. The output is a representation of the CoW, which we refer to as vessel graph. The vessel graph has to be postprocessed in order to remove noise introduced during the segmentation and skeletonization steps. This assures a more accurate representation of the vasculature by the vessel graph.

Labeling usually refers to the assignment of the anatomically correct names to the segments of the vessel graph. However, we developed a different approach in collaboration with our domain expert. We know that during the diagnosis the radiologist is less interested in the anatomical names of the segments but more in how the brain is supplied and consequently in the topology of the CoW. Neuro-radiologists can easily derive the identity of an artery by considering its connection to the supplying main artery. Therefore, we propose a systematic labeling of the vessel graph, starting at the main arteries, which we refer to as root segments. Each segment can be labeled from each main artery as long as there is a connection. A label consists of zero to three terms, depending on its connections to the main arteries. A labeling term is thereby described by the regular expression:

$$[R|L|B][0-9]^+[a-z] \quad (1)$$

The letters R , L and B specify the root segment where R stands for the RICA, L for the LICA and B for the BA. The subsequent numerals specify the number of branching points between a segment and its root segment. Lower case characters at the end indicate the branch index, which enumerates the child branches. This approach is demonstrated on an example tree in Figure 6, assuming that the tree is the right anterior subtree. Per definition, the root segment is the $R0a$ segment. With each branching point, the numeral increases.

Figure 3e shows a vessel graph labeled with our approach. The effectiveness of our labeling approach is illustrated on the example of the segment with the label $R2aL4b$. The $R2a$ term indicates that this particular segment is two branching points away from the RICA. The $L4b$ allows the viewer to assume that there is a connection between the left and right anterior subtree and the segment is four branching points away from the LICA. Therefore, it can be assumed that there is an ACoA to connect the anterior subtrees. Since the enumerator of the right labeling term is smaller than the left one, it indicates that this segment is located on the right anterior subtree. Furthermore, the fact that this segment is not labeled from the BA is a clear indication that there is no connection between the anterior and posterior part. The PCoAs are absent on both sides. This relevant information can be extracted from the single segment label.

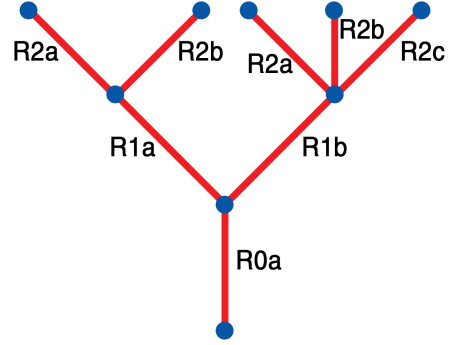


Figure 6: Illustration of the labeling method on a tree example. Assuming this is the right anterior subtree, the root segment is per definition $R0a$. The numeral increases with the number of branching points. Additionally, the labeling direction can be interpreted as the primary direction of the blood flow.




Figure 7: Detailed view of the attributes shown inside the RICA node of the CoWRadar from Figure 3f. These attributes are displayed in diagrammatic form inside the nodes with multiple circular arcs. The inner arc depicts the intensity range of the segment and the outer arc depicts its length.

Finding the starting point of the labeling is crucial in our approach. For each subtree, the root segment must be detected. We determine the root arteries heuristically by analyzing the following seven attributes of each segment:

1. **Length l :** The main arteries are longer than the other segments.
2. **Position on the z-axis u :** The main arteries supply the CoW from below. Therefore, their position on the z-axis is likely to be lower than the other arteries of their subtrees.
3. **Centrality on the x-axis v :** This attribute is used for identifying the BA since it has a medial position on the x-axis. We thereby use the c_{BA} value from the coronal plane. The centrality is calculated from the average distance of a segment to the c_{BA} by only considering the x-components of the positions.
4. **Distance to the centroid of the cluster d :** The cluster centroid is usually close to its main artery since the main artery is the largest one in the subtree.
5. **Intensity values b :** Due to the properties of MRAs, the main arteries have the highest intensity values. b is the average intensity value along the pathway of the segment.

6. **Vertical alignment a :** The main arteries run vertically upwards until they bifurcate into the CoW. A segment is approximated by a vector and a is calculated from the angle between this vector and the z-axis.
7. **Affiliation to the respective subtree s :** We use the clustering result to determine the corresponding subtree of a vessel segment. This is a three-dimensional vector where each component defines the affiliation to the three subtrees.

Finally, we calculate the rank for each segment in the vessel graph to be selected as one of the root segments. For each subtree, the segment with the highest associated rank is selected as the root segment by using the following rank-function:

$$C_{i,j} = \text{sgn}(s_{i,j}) \cdot (l_i \cdot L_j + u_i \cdot U_j + v_i \cdot V_j + d_i \cdot D_j + b_i \cdot B_j + a_i \cdot A_j) \quad (2)$$

$C_{i,j}$ is the rank of segment i to be selected as the root of subtree j , for $j = 1, 2, 3$. The variable $s_{i,j}$ is the component j of the affiliation vector s of segment i . It describes to which subtree the segment i can be assigned to. The above calculated attributes have different ranges. Therefore, we standardize the ranges between 0 and 1 in order to equalize their influence.

$\text{sgn}(\cdot)$ is the sign function. $s_{i,j}$ is only positive, if a segment is part of a subtree. This guarantees that the RICA can only be considered as a candidate for the root segment of the right anterior subtree, the LICA for the left anterior subtree and the BA for the posterior subtree. The influence of the attributes is regulated by the weights $L_j, U_j, V_j, D_j, B_j, A_j$. The three main arteries LICA ($j = 1$), RICA ($j = 2$) and BA ($j = 3$) differ in size, location, alignment and shape and are therefore not determined by the same attributes. For this reason, the weights have to be adjusted accordingly. We empirically determined $L_1 = L_2 = 0.4, U_1 = U_2 = 0.2, V_1 = V_2 = 0, D_1 = D_2 = 0.5, B_1 = B_2 = 0.5, A_1 = A_2 = 0.2$ as weights for the LICA and RICA root segments. Hereby, the centrality attribute v is eliminated since the LICA and the RICA are in the lateral part of the ROI. For the BA root segment, we use the following weights $L_3 = 0.3, U_3 = 0, V_3 = 1, D_3 = 0.3, B_3 = 0.4, A_3 = 0.8$. Hereby, u is eliminated since the BA is not the lowest segment in the posterior subtree. Usually, the two vertebral arteries are below the BA.

4.4. Visual Mapping

The vessel graph contains information about different properties of the CoW, which have to be communicated to the viewer. Figure 3e displays the labeled vessel graph together with the MIP from the transverse view. This display is disadvantageous since the content can not be fully perceived without changing the viewing direction. In this section we propose a visual mapping to abstract the CoW in such a way that it can be easily comprehended by the radiologist. Our visualization supports the physician to observe the overall

configuration of the CoW, while still retaining sufficient details for an extensive analysis.

Our proposed visual abstraction displays the CoW in a similar way as before, namely in a radial graph layout. The CoW consists of subtrees that are connected with each other at the leaves. We chose a circular layout that offers an effective way of displaying the subtrees inside different sectors. Figure 3f demonstrates our approach, which we refer to as Circle of Willis Radar/Radial Visualization (CoWRadar).

We convert the labeled 3D vessel graph into a dual 2D graph representation, i.e., segments are represented by nodes and branching points by edges. Neuro-radiologists are primarily interested in the blood supply and collateral circulations, which can be easily observed in a dual graph. Furthermore, other important attributes of the segments are visually better encoded inside a node than an edge.

The main arteries are located at the center of the radial graph, at the zero level. Increasing level numbers indicate the direction of the blood flow. Every level is represented by a concentric circle, starting from the center and simultaneously reflecting the distance of a segment to its corresponding main supplying artery. In essence, we visually encode the subsequent three different Levels of Detail (LODs) that successively provide more information:

1. **Low:** The overall blood supply to an affected region of the CoW can be derived from the global arrangement of the node within a single sector. An empty sector can be immediately spotted, representing a problematic blood circulation in this region. This is a time-crucial and life-saving aspect and is, therefore, encoded at the lowest LOD to be as fast perceivable as possible (see Figure 9).
2. **Medium:** At this LOD, the branching structure can be observed from the nodes and edges in the CoWRadar. Important are the edges connecting two sectors, since these indicate collateral blood circulation. This is a major aspect to conclude if a certain region is still supplied with blood despite being not connected to its spatially closest main artery [HvdG00]. Again, this can be easily observed in our CoWRadar (see Figure 3f).
3. **High:** The highest LOD offers the possibility to inspect the attributes inside the nodes, which are shown in diagrammatic form (see Figure 7). They allow the radiologist to compare attributes of different segments with each other. For example, length and intensity values are displayed to provide additional information. The length allows the radiologist to distinguish between the different segments. Arteries are further differentiated by intensity values because of various reasons. The main arteries are the brightest in the CoW. Low intensity values could indicate a stenosis.

The CoWRadar shown in Figure 3f can be interpreted as follows. The connection between the left anterior and right anterior sector is established by the ACoA, which we labeled as $R3aL2b$. This label consists of two labeling terms, the $R3a$ term indicates that this segment is three branching points

away from the RICA segment and the *L2b* term indicates that it is two branching points away from the LICA. Furthermore, each segment in the anterior subtrees is labeled from both, the RICA and LICA root segment and therefore, carries two labeling terms. The missing third label implies that the posterior subtree is not connected to the anterior subtrees. Consequently, we can assume that the PCoAs are missing on both sides.

The attributes used for the identification of the main arteries as described above can be visualized within the nodes – circular charts – normalized to the respective maximum value of all segments. This layout offers a compact representation of information, while minimizing the covered space and retaining the overall topology of the CoW. An example is given in Figure 7, showing the root segment *R0aL5b* of the right anterior tree. The inner arc displays the range of the intensity values, whereas the outer arc shows the length of the segment. It can be seen that the RICA is a fairly long segment, encoded in the outer circle. The inner circle conveys the minimum and the maximum intensity value of the segment. These two attributes play an essential role in selecting the root segment, since they characterize the main arteries. The individual circular arcs could encode different attributes, depending on the information the radiologist is interested in.

5. Implementation

All steps in the proposed pipeline were implemented in a software solution for the radiologists. We developed the software as an extension to the AngioVis framework [The15a]. The methods have been implemented on the CPU. We use OpenGL and Qt to render the geometrical elements in the visualization. The software processes the data sets automatically and requires the radiologist only to read-in the data set and start the pipeline. Currently, a data set is processed in 55 seconds on average. The measurements were done on an Intel Core i5 with 3.4 GHz and 16 GB system memory.

6. Results and Discussion

The data sets were acquired across patients with various diseases such as brain tumors or other cerebrovascular diseases, but the majority being stroke patients. The sizes of these data sets are between $448 \times 512 \times 64$ and $512 \times 512 \times 156$ voxels. All data sets have been automatically processed using our proposed approach, without any manual interventions or adjustments.

The first example (see Figure 8) scored good results and demonstrates the effectiveness of our visualization. In this data set, according to the gold-standard, all arteries of the CoW are present except for the left PCoA. As the transverse MIP (see Figure 8a) shows, the left PCoA is the only artery missing as indicated by the yellow arrow. The separation into the three subtrees is shown in Figure 8b and Figure 8c presents the CoWRadar. Nearly all segments are correctly

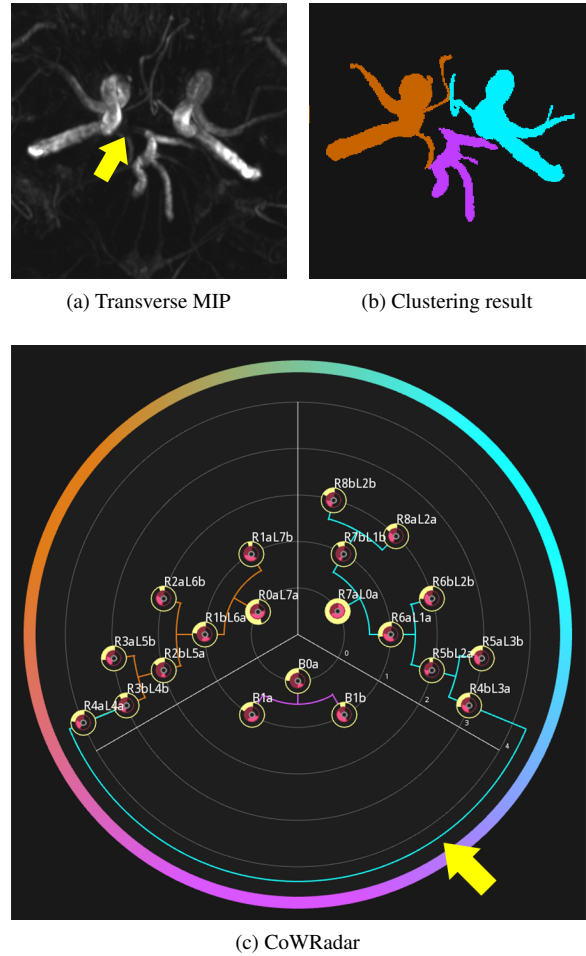


Figure 8: In this patient 91% of the arteries could be correctly identified using our proposed method. (a) shows a transverse MIP indicating a missing left PCoA (yellow arrow). (b) illustrates the clustering result and (c) shows the CoWRadar with the ACoA (yellow arrow) connecting both anterior subtrees.

identified. The arc-shaped connection between the two anterior subtrees indicates the presence of the ACoA (yellow arrow in Figure 8c). The right P1 segment could not be identified due to low intensity values. The right PCoA is correctly detected but not represented by a connection to the posterior sector since the right P1 segment is not correctly identified.

The second example (see Figure 9) demonstrates the limitations of our proposed method. According to the gold-standard, the LICA is missing as shown in Figure 5a. The only other missing artery is the ACoA. We correctly identified the missing LICA, but our approach does not detect the remaining arteries of the left subtree since the root segment is missing. This means that the left M1 and A1 segments will not be detected by default.



Figure 9: The figure displays the CoWRadar of the example data set in Figure 5. In this patient only 58% of the arteries could be correctly identified with our proposed method.

7. Evaluation


We conducted a study consisting of 63 TOF-MRA data sets that were investigated by an expert neuro-radiologist and three volunteering radiologists. The neuro-radiologist, being much more experienced than the three radiologists, created the gold-standard. The task was to identify the presence of the twelve cerebral arteries of the CoW. We compared the findings of our approach and the findings of the radiologists to the gold-standard and calculated the sensitivity, specificity and negative predictive value to compare the results.

The separate evaluation of the main arteries and the entire CoW is motivated by the way our approach works. We proposed a method that automatically identifies the main arteries by certain attributes. However, the remaining arteries are not directly identified but labeled with our systematic labeling approach. Our method displays their connections to the main arteries and the branching structure that lies in between, hence the arteries are visually described.


The sensitivity metric, or true positive rate, measures the proportion of existing segments that could be correctly identified. The specificity metric, or true negative rate, measures the proportion of missing arteries that could be detected as absent. Finally, the negative prediction rate is the proportion of the correctly classified absent arteries. In order for our proposed method to be considered to have good performance, all three of these metrics have to be analyzed. The visualizations of the 63 data sets were presented to the domain expert for interpretation and verification if the arteries could be correctly detected by our approach.



(a) Sensitivity



(b) Specificity



(c) Negative predictive value

Figure 10: The evaluation shows that the performance of our automatic approach is slightly below the radiologists'. (a) Sensitivity: Almost all main arteries could be correctly identified. (b) Specificity: A specificity of 93.04% for the entire CoW indicates that our approach performs better in detecting missing arteries than radiologists. (c) Negative predictive value: It seems in general to be quite low.

The **sensitivity** values are shown in Figure 10a. The sensitivity value demonstrates the ability of our proposed method to detect the presence of arteries. 80.66% of all present arteries could be detected correctly, which is slightly below the performance of the radiologists. The value is mainly decreased by arteries that are represented with low contrast.

Figure 10b shows the **specificity** values. According to the gold-standard, the RICA is absent in two and the LICA is absent in one data set. All three cases have been correctly classified by our proposed method. The specificity is therefore 100%. The specificity of the participating radiologists

is in general lower compared to our method. This is because of increased false positives. The radiologists have the highest number of false positives in the communicating arteries, which means they tend to identify the presence of a communicating artery when the gold-standard states the opposite. The main reason for the slight decrease of specificity for all arteries is caused by a false connection between the left and right anterior subtree.

Our approach seems to favor specificity over sensitivity. This means, a segment is rather detected as missing than as being present. However, these two values do not demonstrate how precise our approach is in detecting the missing arteries by taking the false negatives into account. This is given by the **negative predictive value**, which is shown in Figure 10c. The reason for the low value is the same as for the sensitivity value. Since the number of absent arteries is relatively small compared to the number of present arteries, the influence of the false negatives is much larger on the negative predictive value. However, this value is in general quite low.

8. Limitations

According to the evaluation our proposed method performs well in the detection of absent arteries. This is potentially of great interest for the diagnostic process since they are often the cause of problems.

Our approach mostly deviates from the gold-standard in those cases where the arteries are represented by relatively low intensity values. In such cases, the arteries are not segmented and consequently not represented in the CoWRadar. Another limitation of our approach is given by the way how our labeling works. If a segment is disconnected due to locally low intensity values, our approach stops the labeling at this point. This causes the remaining segments in the respective subtree to be left unlabeled or assigned to the wrong subtree.

9. Conclusion and Future Work

We proposed an automated pipeline for the visual quantification of the CoW in stroke patients. Thereby we developed a novel method for the systematic labeling of the vessel graph. In addition, we proposed heuristics for the identification of the main arteries based on seven attributes of each segment in the vessel graph. The CoW is finally visually represented in a standardized manner in order to provide a preliminary assessment of the CoW's configuration as well as a visual indication of problematic areas. This can be used as an interface for comparisons across multiple patients.

Concluding, the evaluation demonstrated the feasibility and practicability, especially considering the heterogeneity of the samples in our study. Our domain expert stated that the findings of our proposed fully automatic method are as good as those of a beginner radiologist. This shows that our approach is worth to be investigated further.

Acknowledgments

The work presented in this paper has been supported by the Austrian Science Fund (FWF) grant no. TRP 67-N23 (Knowledge Assisted Sparse Interaction for Peripheral CT-Angiography – KASI). The data sets are courtesy of the Karl Landsteiner University of Health Sciences, Austria.

References

- [BGP*03] BULLITT E., GERIG G., PIZER S. M., LIN W., AYLWARD S. R.: Measuring tortuosity of the intracerebral vasculature from MRA images. *IEEE Transactions on Medical Imaging* 22, 9 (2003), 1163–1171. 2
- [BMJ*05] BULLITT E., MULLER K. E., JUNG I., LIN W., AYLWARD S.: Analyzing attributes of vessel populations. *Medical Image Analysis* 9, 1 (2005), 39–49. 2
- [Bog12] BOGUNOVIC H.: *Geometric Modeling and Characterization of the Circle of Willis*. PhD thesis, Universitat Pompeu Fabra, 2012. 2
- [CA05] CONDURACHE A., AACH T.: Vessel segmentation in angiograms using hysteresis thresholding. In *Proceedings of the IAPR Conference on Machine Vision Applications 2005* (2005), pp. 269–272. 2
- [Can86] CANNY J.: A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 8, 6 (1986), 679–698. 2
- [DLR09] DRAPER G. M., LIVNAT Y., RIESENFELD R. F.: A survey of radial methods for information visualization. *IEEE Transactions on Visualization and Computer Graphics* 15, 5 (2009), 759–776. 2
- [HvdG00] HARTKAMP M. J., VAN DER GROND J.: Investigation of the Circle of Willis using MR angiography. *MedicaMundi* 44, 1 (2000), 20–27. 2, 7
- [ID90] INSELBERG A., DIMSDALE B.: Parallel coordinates: a tool for visualizing multi-dimensional geometry. In *Proceedings of IEEE Visualization* (1990), pp. 361–378. 3
- [KQ03] KIRBAS C., QUEK F. K. H.: Vessel extraction techniques and algorithms: A survey. In *Proceedings of the IEEE Symposium on Bioinformatics and Bioengineering* (2003), pp. 238–245. 2
- [LKC94] LEE T.-C., KASHYAP R. L., CHU C.-N.: Building skeleton models via 3-D medial surface/axis thinning algorithms. *Graphical Models and Image Processing* 56, 6 (1994), 462–478. 2, 5
- [Mis13] MISTELBAUER G.: *Smart Interactive Vessel Visualization in Radiology*. PhD thesis, TU Wien, 2013. 2
- [PB13] PREIM B., BOTHA C. P.: *Visual Computing for Medicine*. Elsevier, 2013. 2
- [Poc04] POCK T.: *Robust Segmentation of Tubular Structures in 3D Volume Data*. Master's thesis, Graz University of Technology, 2004. 2, 6
- [The15a] THE ANGIOVIS FRAMEWORK.: <http://www.angiervis.org/>, April 2015. 8
- [The15b] THE TOP 10 CAUSES OF DEATH, WORLD HEALTH ORGANIZATION.: <http://www.who.int/mediacentre/factsheets/fs310/en/>, April 2015. 1