Eurographics Conference on Visualization (EuroVis) 2019
M. Gleicher, H. Leitte, and I. Viola
(Guest Editors)

Volume 38 (2019), Number 3

A Geometric Optimization Approach for the Detection and Segmentation of Multiple Aneurysms

Kai Lawonn[1], Monique Meuschke[2], Ralph Wickenhöfer[3], Bernhard Preim[2], and Klaus Hildebrandt[4]

[1]University of Koblenz-Landau, Germany; [2]University of Magdeburg, Germany;
[3]Heart of Jesus Hospital Dernbach, Germany; [4]Delft University of Technology, The Netherlands

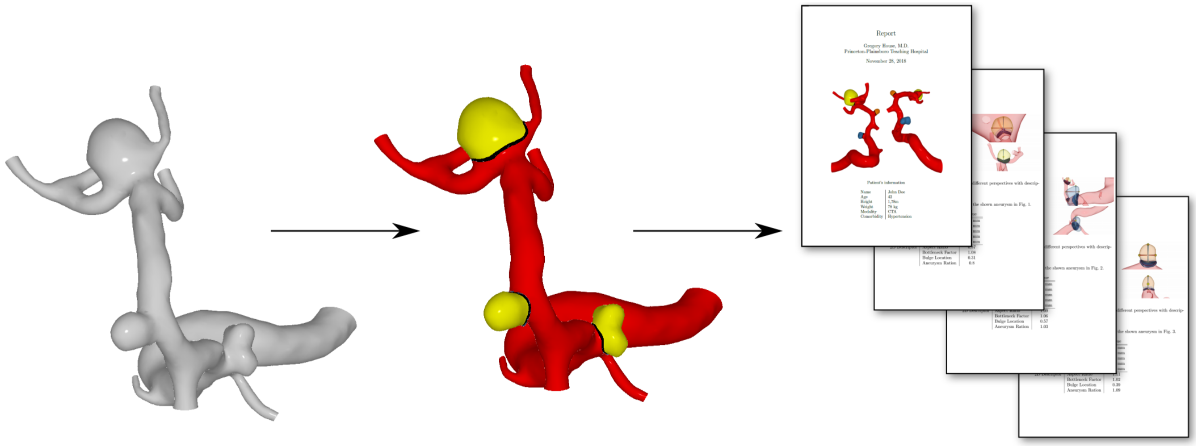

Figure 1: From an input mesh representing the surface of a vessel, our algorithm detects and segments the aneurysms in the vessel. Afterwards, a report is generated including meta information on the patient as well as summaries of characteristics of the aneurysms, e.g., their widths and heights.

Abstract

We present a method for detecting and segmenting aneurysms in blood vessels that facilitates the assessment of risks associated with the aneurysms. The detection and analysis of aneurysms is important for medical diagnosis as aneurysms bear the risk of rupture with fatal consequences for the patient. For risk assessment and treatment planning, morphological descriptors, such as the height and width of the aneurysm, are used. Our system enables the fast detection, segmentation and analysis of single and multiple aneurysms. The method proceeds in two stages plus an optional third stage in which the user interacts with the system. First, a set of aneurysm candidate regions is created by segmenting regions of the vessels. Second, the aneurysms are detected by a classification of the candidates. The third stage allows users to adjust and correct the result of the previous stages using a brushing interface. When the segmentation of the aneurysm is complete, the corresponding ostium curves and morphological descriptors are computed and a report including the results of the analysis and renderings of the aneurysms is generated. The novelty of our approach lies in combining an analytic characterization of aneurysms and vessels to generate a list of candidate regions with a classifier trained on data to identify the aneurysms in the candidate list. The candidate generation is modeled as a global combinatorial optimization problem that is based on a local geometric characterization of aneurysms and vessels and can be efficiently solved using a graph cut algorithm. For the aneurysm classification scheme, we identified four suitable features and modeled appropriate training data. An important aspect of our approach is that the resulting system is fast enough to allow for user interaction with the global optimization by specifying additional constraints via a brushing interface.

1. Introduction

Aneurysms are abnormal weak areas in the walls of blood vessels that can bulge out and balloon. They bear the risk of rupture, which leads to a subarachnoid hemorrhage (SAH) causing high risks of mortality. Though most aneurysms will never rupture, the potential risk of a SAH makes the detection and risk-assessment of aneurysms an important issue. Imaging methods are used for the detection and localization of aneurysms. Decisions on whether or not to treat aneurysms have to be taken carefully as interventions bear the danger of a fatal outcome [CVS*15]. In clinical routine, the risk assessment and treatment planning are based on the aneurysm's size and location, where aneurysms larger than 7 mm are usually treated. Physicians manually extract the aneurysm's size either using the acquired image data or a surface mesh reconstructed from the image data. However, the size alone is not reliable enough for evaluating rupture risk and treatment options. In current medical research, the influence of other morphological descriptors, such as the height and width, on the aneurysm state is analyzed. Unfortunately, the manual extraction of these descriptors is an error-prone and time-consuming process [MGW*18]. Accurate measurements are necessary to identify meaningful risk factors.

We present a novel system that assists physicians in the detection and risk assessment of aneurysms. The system can be used to detect and segment aneurysms in blood vessels, extract the ostium curves separating the aneurysms from the healthy vessel, and estimate the relevant morphological parameters. Finally, it generates a report that summarizes the relevant results, including a list of morphological descriptors, and combines them with renderings providing visual impressions of aneurysms.

Our approach for detecting and segmenting aneurysms splits into three stages.

- In the first stage, a global optimization problems is solved to create list of candidates regions on the input vessel tree.
- In the second stage, the candidates are classified as either *aneurysm* or parts of the *healthy vessel* tree.
- The third stage is optional and allows users to adjust or correct the results of the previous stages by brushing on the surface. The strokes are interpreted as additional constraints that are integrated to the global optimization.

For creating the list of candidate regions in the first stage a combinatorial optimization problem is solved. It provides every triangle of the mesh with a binary label indicating whether or not the triangle could be part of an aneurysm. The objective, we model for this problem, combines two terms. One term that based on an analysis on the local geometry associates a costs with two possible labels for each triangle. The underlying idea is that cylindrical or saddle-like regions are more likely to be parts of the vessels whereas spherical regions indicate aneurysms as these bulge out. Since an aneurysm can not be detected by analyzing at the local geometry only, a second term is added to the objective. This term promotes consistent labeling among nearby triangles by penalizing switching labels between nearby triangles. This leads to consistently labeled regions as either aneurysm or vessel and creates short borderlines between aneurysms and vessels. The resulting optimization problem can be solved in a fraction of a second using graph cut algorithms. The result of this stage is a list of segmented regions that are candidates of aneurysms. After this first stage, we employ a 2nd stage to reduce the candidate set and thus the number of false positives (FPs) with an appropriate classifier. For training and validation, we used the aneurysm database provided by Pozo et al. [PSFC17]. Moreover, we experimented with artificially created training data.

The resulting method is an efficient approach for the detection and segmentation of aneurysms. Once aneurysms are detected, ostium curves and morphological descriptors are computed. To improve the quality of the generated of ostium curves, we developed a specific smoothing scheme that produces better aligned and smoother curves. In summary, the main contributions of this paper are:

- A fast detection and segmentation algorithm for finding multiple aneurysms in vessels that combines a combinatorial optimization problem with a specific classifier.
- A brushing tool that allows users to corrected or adjusted the results of the optimization problem.
- An approach for smoothing the ostium curve based on a specific anisotropic diffusion of an indicator function.
- A report generation summarizing all important facts about the data, where additional renderings of the aneurysms are provided.

2. Related Work

In this section, we first discuss prior work on the (semi-)automatic extraction of aneurysms and the estimation of their morphological descriptors. Then, we summarize approaches for the generation of reports of clinical data, which evaluates the state of a patient.

2.1. Aneurysm Detection

Most algorithms for detecting aneurysms employ the vascular centerline. In a nutshell, these algorithms reconstruct the diameter of the vessel above the centerline and search along the curve for areas with enlarged diameter as this indicates the location of an aneurysm.

Based on clinical image data, Karmonik et al. [KAB*04] proposed a technique to determine points on the centerline in each image slice. After the points are constructed, circles are fitted to the cross-sections of the parent vessel. An analysis of the variation of the radii of the circles is used to locate the ostium. The approach yields erroneous results if the parent vessel deviates strongly from a tubular shape, which often happens nearby an aneurysm. Lauric et al. [LMFM10] segmented vessel regions in clinical image data and computed their medial axis. Based on the writhe number, they distinguished tubular and non-tubular regions to identify aneurysms. Their approach fails if an aneurysms has a more elongated shape similar to the shape of vessels. Hassan et al. [HHFP11] suggested an approach that localizes the centerline in clinical image data and fits a quadratic function along the contour of the vessel. If a certain coefficient of the quadratic function exceeds a user-provided threshold, the corresponding part of the vessel is considered to be an aneurysm. The algorithm requires that the centerline runs only along the vessel and does not branch and enter the aneurysms. Hentschke et al. [HBNT11] developed a method for the detection of aneurysms based on multimodal angiographic images using a blob-enhancing filtered image. Another image-based approach for detecting cerebral aneurysms was introduced by Chandra et al. [CM16]. The core idea of their method is a combination of automated thresholding for vessel segmentation and morphological filtering for detecting aneurysms in the segmented vessels. In principle, image-based methods are able to detect multiple aneurysms. However, current methods struggle with missing detections and false positives. Moreover, the reliable ostium extraction from 2D image slices is challenging [KAB*04].

In addition to image-based techniques, 3D surface representations of the vasculature are used for aneurysm detection. In this setting, a combination of Voronoi diagrams and distance transformations was used for the construction of the centerline by Antiga et al. [APB*08] and the detection of aneurysms by Cárdenes et al. [CPB*11]. Ford et al. [FHP*09] analyzed the part of the Voronoi diagram that is inside the vessel for the detection of the aneurysm. The outgoing vessels are identified as the regions in which the maximally inscribed sphere intersect the centerline. The aneurysms are detected by an analysis of the distance of the vessel wall and the Voronoi regions. Regions with maximum distance are classified as aneurysms. Neugebauer et al. [NDSP10] introduced an interactive tool for the construction of the ostium curve. After a start and end point on the centerline as well as a point on the aneurysm are specified, their algorithm generates a candidate ostium curve and provides the user with four control points that allow for adjusting and correcting the suggested curve. Mohamed et al. [MSM*10] determined all vessel parts that are further away from the centerline than a given threshold (1 mm). Then, the largest connected component is used as aneurysm surface. Larrabide et al. [LVUC*11] employed the skeleton of the vessel surface to identify potential aneurysm regions. Parts of the skeleton going from an endpoint to the first branch point are used as candidates for aneurysms. These parts are analyzed to identify the aneurysm. A problem of this approach is that irregularities of an aneurysm, for example additional bulges on the aneurysm—so-called *blebs*, can lead to branches of the skeleton within the aneurysm. In such a case, the assumption that only the part between the endpoint and the first branch correspond to the aneurysm is violated, which results in incomplete or wrong detection. Moreover, the user needs to specify how many aneurysms are present in the data. Jerman et al. [JPLŠ15] proposed a ray-casting approach to identify aneurysms. From the centerline, rays are emitted towards the vessel wall. The distance from the start point to the first intersection is used to detect abnormal distances, which are then used to detect the aneurysm. In this approach, the ostium curve is constructed by applying *Otsu thresholding* [Ots79]. Recently, Meuschke et al. [MGW*18] present a scheme for the detection of aneurysms that does not require the centerline as input. Based on the skeleton of the vessel surface, parts with the highest *mean shape index* are considered to be an aneurysm. Besides the mentioned limitations of the individual approaches, the following two limitations are common to all surface-based approaches. First, the methods always locates an aneurysm, even in the case that the input is a completely healthy vessel. Second, they are restricted to finding one aneurysm in a vessel and would therefore miss further aneurysms [FHP*09, MSM*10, CPB*11, JPLŠ15, MGW*18].

2.2. Report Generation

In clinical routine, reports give a detailed insight into the patient's state, diagnostic findings, and recommendations for treatment. In addition, they serve as a basis for discussions of the involved physicians, e.g., in a stroke unit or tumor board. Often, automatically-generated documentations are text-based descriptions that lack a clear structure, which hampers the readability and interpretation of the reports. Much effort is dedicated to improving clinical reports by developing systems that support the generation of structured reports. The approach taken is to provide predefined data entries, which are then filled with information. For example, Karim et al. [KFB*13] developed a web-based application that produces a structured report in the field of abdominal aortic aneurysms. For this purpose, different graphical widgets, such as drop-down menus, diagrams and predefined images are used. Wibmer et al. [WVS*14] presented a standardized diagnostic certainty lexicon for reporting the likelihood of prostate cancer. With this the number of expressions used by radiologists to indicate their levels of diagnostic certainty was reduced, which supports clinical decision-making.

More advanced methods use visual analytics techniques to support the collaborative analysis of medical data sets. Pankau et al. [PWN*15] developed a 3D documentation system that allows standardized reporting for head and neck cancer including aspects like tumor staging. The benefits of visual elements in a report, such as standardized and individualized screenshots, were emphasized by the physicians. Ai-Awami et al. [AABH*16] presented *NeuroBlocks*, a multi-user web-based visualization system for managing and performing very large volumetric segmentations in neuroscience. Users groups having different access rights for editing and verifying results can be defined. Bannach et al. [BBJ*17] developed a system that combines medical image analysis with visual analytic techniques to build, analyze, and evaluate patient cohorts. They used *radiomics* that are quantitative image features to enrich the patient's diagnosis with further information. These are complemented with additional meta data like the patient's age and gender. Filter masks are provided to define cohorts that should be statistically analyzed.

While existing systems concentrate on the analysis and structured communication of clinical image data, for the investigation of cerebral aneurysms other data such as morphological descriptors based on geometrical information need to be included. Our method provides a consistent report structure that augments estimated values of morphological descriptors with additional depictions. This provides the physician with a rich and easily accessible source of information and is intended to support the analysis and treatment planning.

3. Requirement Analysis

The basis of our approach is a careful analysis of the requirements our method needs to satisfy to be effective in practice. We defined the requirements in close cooperation with one medical expert, who has 16 years of work experiences. He regularly treats cerebral aneurysms and is involved in the diagnosis and treatment planning. To define requirements, we discussed the workflow in clinical routine. For every patient, different types of data are acquired including clinical image data as well as patient-specific meta information, such as gender and age. These data are analyzed collaboratively by different domain experts. While currently no standardized report is used for their communication, they see it as something they could potentially benefit from. Physicians search manually for aneurysms in the acquired image data. Morphological descriptors of the detected aneurysms, in particular their widths and diameters, are estimated and form the basis for decisions on whether or not to treat an aneurysm. For treatment planning, additional morphological properties, such as the width of the ostium, are considered. The morpho-

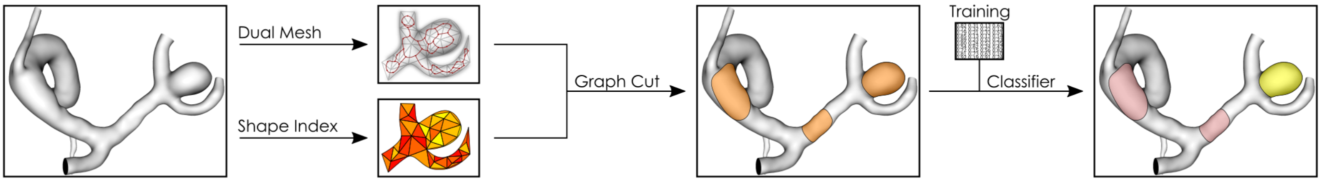

Figure 2: Illustration of our pipeline. First, an input mesh is loaded, a data structure storing the mesh's dual edges is constructed, and the shape indices of all triangles are computed. Then, an optimization problem is solved to generate a list of candidate patches on the surface. Finally, a classifier identifies the aneurysms in the candidate list.

logical descriptors are manually estimated using either clinical software that operates on 2D-image slices or the open-source program *ParaView* that operates on 3D-surface meshes. In both cases, the estimation of the descriptors is a time-consuming and error-prone process. For example, it is difficult to find the maximum extension manually. For the construction of the ostium curve, which is done manually, Paraview is used. Again, this is a time-consuming task, in which 2D-planes need to be specified to cut the surface mesh along the ostium. Based on these discussions, we defined the following requirements for our method:

Req. 1. An approach for detecting and segmenting aneurysms in vessels that avoids users having to specify values for parameters for which they have no intuition.

Req. 2. To adapt to the anatomical diversity, the approach must allow manual correction of the results of the detection and segmentation scheme.

Req. 3. For the detected aneurysms, morphological descriptors should be automatically estimated.

Req. 4. A report that lists the values of morphological descriptors and augments this information with renderings of the aneurysms should be generated.

4. Methods

In this section, we describe how we combine the shape index, which characterizes the local surface geometry, a global combinatorial optimization problem, a k-means classifier, a brushing interface for user interaction and a smoothing scheme for the ostium curve to form an effective system for detecting and segmenting aneurysms in vessels. Furthermore, we discuss the generation of a report, which summarizes and illustrates the results including the morphologic properties of the aneurysms. While we address the whole pipeline from 3D-image data to the generation of the final report summarizing the results, our focus in this section is on aneurysm detection and segmentation in vessels described by triangular surface meshes and the report generation. The conversion of 3D-image data to a surface mesh representation is discussed in Section 7.5.

4.1. Preliminaries

We start the technical description of our method by introducing some notation. The triangular surface mesh representing the surface of the vessel is denoted with $\mathcal{M}$. The mesh is a simplicial complex consisting of vertices $V = \{1,2,\dots,n\}$, edges $E = \{(i,j) \mid i,j \in V\}$, and triangular faces $F = \{(i,j,k) \mid (i,j),(j,k),(k,i) \in E\}$. To describe the first stage of our algorithm, we will use dual edges $E^\star = \{(i,j) \mid i,j \in F\}$. A pair (i,j) is a dual edge if triangles i and j are neighbors, i.e., if the triangles share a common edge. For every vertex $i \in V$, we denote with $\mathbf{p}_i \in \mathbb{R}^3$ the vector encoding the position of the vertex in $\mathbb{R}^3$ and with $\mathbf{n}_i \in \mathbb{R}^3$ and $\mathbf{n}^j \in \mathbb{R}^3$ the unit surface normal vectors at vertex i and triangle j. Using double indices, we refer to the differences, *e.g.*, $\mathbf{p}_{ij} = \mathbf{p}_j - \mathbf{p}_i$ and $\mathbf{n}_{ij} = \mathbf{n}_j - \mathbf{n}_i$.

Our algorithm for detecting aneurysms makes use of the curvatures of the surface of the vessel. Various schemes for the approximation of the curvatures of a surfaces from an approximating triangle mesh have been proposed [CSM03, Rus04, PWY*07, HP11]. For a recent quantitative evaluation curvature estimation scheme, we refer to [VVP*16]. For our experiments, we use the technique by Rusinkiewicz [Rus04], which yields an approximation of the two principal curvatures, κ_1, κ_2, at every triangle of a mesh. Instead using the principal curvatures directly, we use them to compute the shape index S, which is defined as

$$S = \frac{1}{2} - \frac{1}{\pi}\,\mathrm{atan}\left(\frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2}\right). \tag{1}$$

The shape index of a point, in our case of a triangle, is a number from the unit interval that describes the local shape of the surface. It depends only on the principal curvatures, and, therefore, is invariant to rigid motion. Moreover, the shape index is invariant to scaling of the shape as it is a quotient of the curvatures. For example, spheres of different radius have the same shape index. A shape index of 1 corresponds to a convex spherical region, in which both principal curvatures take the same positive value, and, a shape index of 0.5 corresponds to saddle-like shape, in which both principal curvatures have the same magnitude but opposite signs. A shape index of 0 corresponds to a concave spherical region where both principal curvatures have the same negative value. To provide some intuition on the shape index, Fig. 3 shows some example surfaces with varying shape index. For more background on the shape index, we refer to [KvD92].

4.2. Candidate Aneurysms

The first stage of our algorithm detects and segments regions in the vessels that serve as candidate aneurysms. In this stage, every triangle is assigned a binary label indicating whether or not the triangle is part of a candidate aneurysm. The values of the labels are determined by solving a combinatorial optimization problem. Our modeling of this optimization problem is based on two heuristics. The

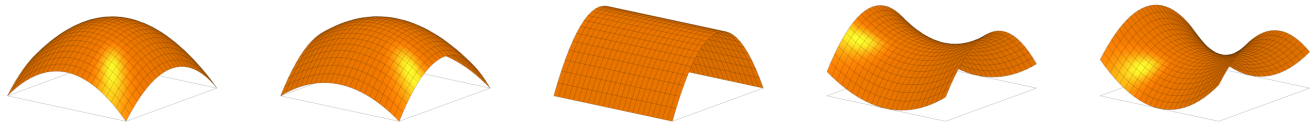

Figure 3: Different shape index values for the depicted surface.

first is that we think of a vessel as a curved tube. Therefore, its local geometry is rather that of a cylinder or a saddle than that of a sphere. On the other hand, the aneurysms are areas that bulge out and balloon, which means that their local geometry is rather spherical than cylindrical. The second heuristic is that at the border between an aneurysm and the vessel the local geometry is rather saddle-like. This is also a consequence of the ballooning of the aneurysm.

The optimization problem assigns a binary label l_i to every triangle of the mesh. If $l_i = 1$, then triangle i is labeled as part of an aneurysm, and, if $l_i = 0$, it is considered part of the vessel. The binary vector $l \in \mathcal{Z}^{|F|}$ stacks all labels l_i. The objective

$$E(l) = \sum_{i \in F} D_i(l_i) + \gamma \sum_{(i,j) \in E^\star} V_{i,j}(l_i, l_j), \tag{2}$$

to be minimized consists of a data term, which is a sum of the unary terms D_i, a regularization term, which is a sum of the binary terms $V_{i,j}$, and a smoothing parameter γ.

The data term reflects our two heuristics. We use the shape index as a rigid motion- and scale-invariant measure of the deviation of a local geometry from being spherical, which indicates an aneurysm, as well as from being saddle-like or cylindrical, which corresponds to a vessel. The data term is univariate, thus providing a cost for the label given to a triangle. If the triangle i is labeled *aneurysm*, the cost depends on the difference of the shape index S_i and 1, which is the shape index of a sphere. If the triangle is labeled *vessel*, the cost depends on the difference of S_i and 0.5, the shape index of a saddle. The penalty function grows exponentially, meaning that labeling a saddle-shaped region as *aneurysm* or labeling a spherical region as *vessel* yields a high penalty. Explicitly the data term is defined as:

$$D_i(l_i) = \begin{cases} -\log(+2 - 2 \cdot S_i) & \text{if } l_i = 1 \\ -\log(-1 + 2 \cdot S_i) & \text{if } l_i = 0. \end{cases} \tag{3}$$

The regularization term is a binary term that specifies a cost for the labeling of pairs of triangles. In our case, only neighboring triangles with different labels produce a cost. Explicitly the regularization term is

$$V_{i,j}(l_i, l_j) = \begin{cases} 0 & \text{if } l_i = l_j \\ -\log(|S_{ij}|) & \text{if } l_i \neq l_j. \end{cases} \tag{4}$$

The term S_{ij} is the difference of the shape indices of the two triangles. To explain our modeling of the binary term, we want to mention that $-\log(|S_{ij}|) = -\log(1/2) - \log(|2S_{ij}|)$. This means that the penalty for switching labels consists of two terms. The first is the constant term $-\log(1/2)$. This term discourages switching labels, and, thereby, helps to create consistent aneurysm structures. For example, though aneurysms tend to have spherical local geometry, due to irregularities there can be small areas in the aneurysm which are cylindrically shaped. Though the data term would result in a lower cost for labeling the area as *vessel*, the regularization term encourages to still label the areas as *aneurysm* because switching the label produces additional costs. Moreover, the penalty for switching labels encourages creating short borderlines between aneurysms and healthy vessels regions.

The second term, $-\log(|2S_{ij}|)$, depends on the difference of the shape indices of the triangles. A large difference in shape index results in low cost and vice versa. The reason for this choice relates to our second heuristic, which states that the border region between aneurysm and vessel is rather saddle-like shaped. Since the aneurysm is spherical, there is strong variation of the shape index across the border and only little variation of the shape index along the borderline. This means that associating little cost to pairs of triangles with large difference of shape index encourages the scheme to actually place the borderline, as desired, orthogonal to the direction of variation of the shape index. The factor of 2 in the second term appears since the value of S_{ij} is between 0 and 0.5 and the argument for the $-log$ function should be in the range $[0,1]$. The smoothing parameter γ was assigned the same value for all reported results. The value was determined experimentally and set to 4.1.

The resulting optimization problem can be efficiently solved using *graph cut* algorithms. We refer to [BVZ01, YBVRZ01, BK04] for some background on this type of optimization problem. The final step of the first stage of the algorithm, is a connected component analysis of all triangles that have the label *aneurysm*. The result after this step is a list of regions on the surface that are candidates of aneurysms. In our experiments, the optimization was able to identify all aneurysms on the vessels we tested with, which includes the The problem, however, is that also parts of vessels were labeled as aneurysms (false positives). Therefore, we added a second stage to the algorithms that separates the false from the true positives.

4.3. Classification

The goal of the second stage of our algorithm is to remove the false positives from the list of candidate aneurysms that was constructed in the first stage. This is achieved by a binary classification, *aneurysm* or part of the healthy *vessel*, of all candidates. In contrast to the first stage of the algorithm, in which the individual triangles received labels, in this stage entire regions, which result from the connected component analysis step of the first stage, receive a label. For the classification, we use four carefully chosen features. By a feature, we mean a statistical value for the whole candidate region. To select the features, we tested different possible features for statistical significance.

The first feature, denoted by f_1, is the mean shape index of the

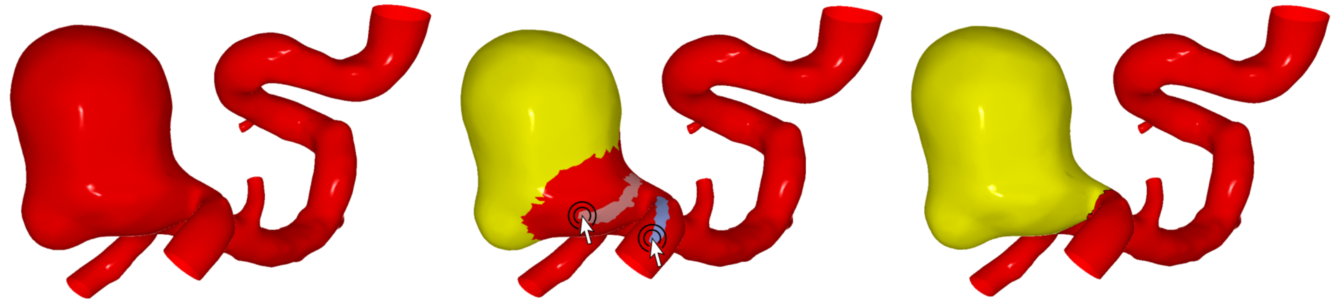

Figure 4: Left the input mesh is shown. After the graph cut algorithm was applied the aneurysm was detected, but the ostium needs to be corrected. Thus, the user brushes aneurysm and vessel parts in light red and blue, respectively. The result is shown right.

region. Since we expect the aneurysms to be rather spherical and the vessels to be rather cylindrical or saddle-like, the mean shape index of an aneurysm will tend to be larger than that of the false positive. The second feature relates to the shape of the region. We use the vertices of the region as a point sample, construct the covariance matrix of the point sample, and compute the singular values $\lambda_1 \geq \lambda_2 \geq \lambda_3$ of the covariance matrix. The feature is

$$f_2 = 3\frac{\lambda_3}{\lambda_1+\lambda_2+\lambda_3}, \tag{5}$$

which is one of Westin's measures [WEMM*02]. Our interpretation of why this feature is useful is that the aneurysms tend to be rather roundish whereas the parts of vessels tend to stretch out along the vessel. As a result, f_2 tends to be larger for the aneurysms than for the false positives.

The third and fourth features relate to the shape of the boundary of the region. We use the vertices of the boundary as a point sample and compute the singular values of the corresponding covariance matrix. The two features are the two Westin's measures

$$f_3 = 2\frac{\lambda_2-\lambda_3}{\lambda_1+\lambda_2+\lambda_3} \qquad f_4 = \frac{\lambda_1-\lambda_2}{\lambda_1+\lambda_2+\lambda_3}. \tag{6}$$

Our explanation for why these features help to detect the false positives is that the boundary of the aneurysms tends to be circular. This means that the difference of the first two singular values is rather small and the third singular value is small compared to the other two. In contrast, the false positives do not have this structure. They rather stretch in one direction. Therefore, aneurysms tend to have a larger value of f_3 and a smaller value of f_4 than the false positives, see Fig. 6 for examples of the f values on *case 43*.

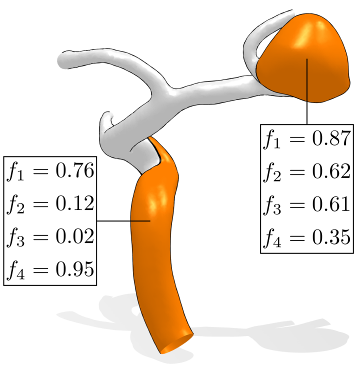

Figure 6: Two candidate regions on a vessel identified by the first stage of our algorithm and the values of the four features used for classification are shown.

To train and evaluate the classifier, we used the aneurysm database provided by Pozo et al. [PSFC17]. To reduce computation time, we reduce every surface to 20,000 triangles using the mesh coarsening algorithm of Garland and Heckbert [GH97]. Afterwards, we ran the first stage of our algorithm on all vessels in the data set to create a list of candidate regions. We computed all four features for every candidate and manually classified them. The resulting data consists of 105 candidate regions, 62 aneurysms and 43 vessel patches. We experimented with different classification techniques, including decision trees, discriminant analysis, logistic regression, variants of support vector machines, and nearest neighbor classifiers, and found the weighted k-nearest-neighbor technique [HS04] to perform well for our problem. We used a 5-cross validation to evaluate the classifier and achieved an accuracy of 99%.

Since our data set is quite small and we have an expectation of what values the four features will take for aneurysms and for vessel regions, we decided to also experiment with a classifier trained on purely artificial data. We generated an artificial list of 10,000 4-tuples of feature values each equipped with a label, either *aneurysm* or *vessel*. The 4-tuples are generated by drawing four values from certain random distributions. A tuple that receives the label *aneurysm* is drawn from a Gaussian distributions with mean values $\mu_{f_1} = 0.85$, $\mu_{f_2} = 2/3$, $\mu_{f_3} = 2/3$, and $\mu_{f_4} = 1/3$, and, for a tuple with label *vessel*, the mean values are $\mu_{f_1} = 0.75$, $\mu_{f_2} = 1/3$, $\mu_{f_3} = 1/3$, and $\mu_{f_4} = 2/3$. We set the standard deviation σ of the Gaussian distributions to be $1/20$ for the first feature and $1/6$ for the other three features. The mean values represent our expectation of the feature values for the two classes. After training on the artificial data set, the classifier was evaluated on our data set comprising 105 candidates. Interestingly, the resulting performance was also 99%, hence matched the performance of the classifier trained on real data. This indicates that the four features are well chosen and the resulting classification problem is not difficult.

4.4. Manual Correction

Since the aneurysm detection is fast, a user can interact with the method to modify or correct results. We developed a brushing interface that can be used to adjust or correct the segmentation of the aneurysms. Additionally, it can be applied to added missing aneurysms to the results or remove misclassified regions. The user can brush on the surface of the vessel in two colors, magenta to indicate that a region should receive the label *aneurysm* and blue to

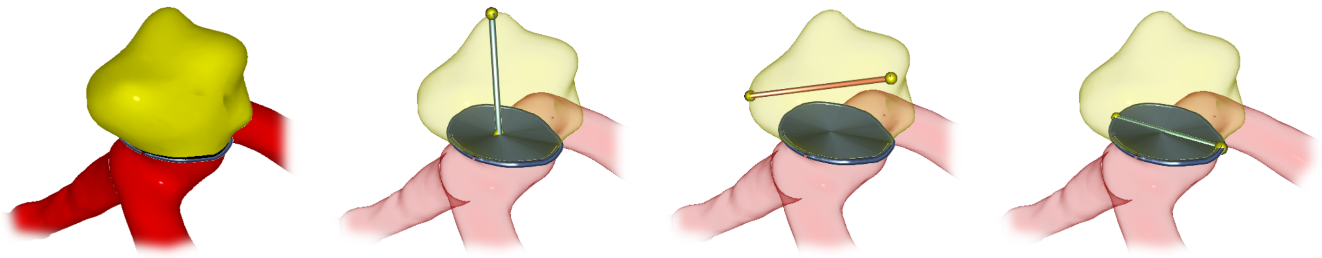

Figure 5: Left the input mesh is shown. Based on the ostium surface, different morphological descriptors are calculated and rendered within the semi-transparent aneurysm surface. The start and endpoints of each descriptor are depicted as spheres and the connecting lines are visualized as tube. From left to right the aneurysm's height, its width, and the width of the ostium curve are shown.

define healthy vessel regions. The user input provides constraints to the optimization. Depending on the color used, the brushed triangles are constrained to be aneurysms or parts of the healthy vessel, respectively. The constraints are implemented by setting the shape index of the corresponding triangles to 1 for magenta strokes and 0.5 for blue strokes. We want to emphasize that the user only needs to indicate the desired changes by some strokes. During the optimization the labels of the strokes will be propagated to the neighboring regions. The process is illustrated in Figure 4 on the vessel *case 21* from the data set [PSFC17].

4.5. Smoothing Scheme

In addition to the optional correction of the labeling, we developed a post processing step for improving the segmentation. Since the optimization problem results in binary labeling of the triangle, the borderline between aneurysms and healthy vessel, the ostium curve, is restricted to run along edges of the mesh, which results in jagged borderlines. A second issue that we observed is that the optimization sometimes constructs aneurysms that are concave regions in the vessel. Figures 7 and 8 show two examples of segmentation where regions classified as vessels protrude into the aneurysms. We developed a smoothing scheme that counteracts both effects. The first step is to convert the label of the triangles into a continuous function that is a linear polynomial on every triangle. Then we use a threshold, which we set to 0.95, to separate the aneurysm from the vessel. The resulting ostium curve is no longer constraint to run along the edges. The idea is to smooth the continuous function in a specific way to obtain smoother regions and ostium curves. Functions on the mesh that are continuous and linear polynomials over the triangles can be described by listing function values at the vertices. To convert the labels at triangles to such a function, we simple the set values at the vertices to 1 if they are adjacent to a triangle with label *eneurysm* and 0 otherwise. For smoothing of the resulting function, which we denote by $\mathbf{u}$, we use an anisotropic diffusion scheme. The anisotropy depends on the norm of the gradient of $\mathbf{u}$. We compute the gradient of $\mathbf{u}$ in every triangle and set the weight $\mathbf{w}_i$ at a vertex to be the maximum of the norms of the gradients in all adjacent triangles. The effect our choice of weights is that the smoothing process is concentrated to the border region between aneurysm and vessel. We use an implicit integration of the anisotropic diffusion process, which requires solving the following sparse linear system in every iteration

$$(I - \tau \mathbf{W}^k \mathcal{L})\mathbf{u}^{k+1} = \mathbf{u}^k. \tag{7}$$

Here $\mathbf{u}^k$ and $\mathbf{u}^{k+1}$ list the function values at the k^{th} and $(k+1)^{st}$ iteration, I is the identity matrix, $\mathcal{L}$ is the cotan Laplace matrix [BKP*10], $\mathbf{W}^k$ is a diagonal matrix that lists the weights $\mathbf{w}_i$, and τ is the step size. The scheme is initialized with $\mathbf{u}^0 = \mathbf{u}$ and the weights $\mathbf{w}_i$ are updated after each iteration. Before starting the

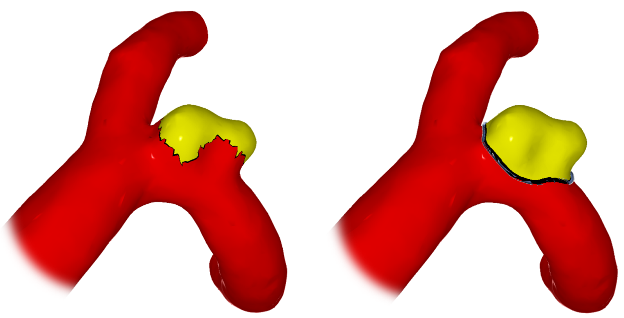

Figure 7: A result of the smoothing process: On the left the segmented aneurysm before and on the right after smoothing are shown.

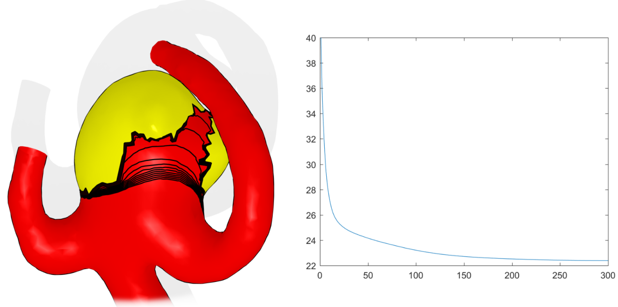

Figure 8: Illustration of the evolution of the ostium curve during the smoothing process. The plot on the right shows the decrease of length of the ostium curve during smoothing.

Table 1: Summary of the performance test with the number of triangles, the number of vertices, the calculation of the shape index, the computation of the dual mesh, the graph cut, the connected component analysis (CCA), the smoothing scheme with 20 steps of iterations, and the extraction of morphological parameters in seconds depending on the number of triangles of the aneurysm.

Name	$\lvert F\rvert$	$\lvert V\rvert$	Shape in s	Dual in s	Graph cut in s	CCA in s	Smoothing in s	Morph. in s (# tri)
Case1	281.932	141.130	12.5	3.1	1.19	2.27	51.5	12.7 (44.479)
	140.965	70.602	6.3	1.4	0.5	1.24	10.6	5.3 (22.225)
	70.482	35.341	3.1	0.7	0.24	0.64	15.1	2.5 (11.120)
	35.241	17.702	1.6	0.3	0.12	0.31	5.9	1.1 (5.535)
Case40	75.434	37.861	3.2	0.8	0.26	0.66	1.6	2.8 (2.801)
	37.716	18.982	1.6	0.3	0.13	0.33	0.5	1.1 (1.206)

smoothing process, we mark all vertices that are adjacent to a triangle with the label *aneurysm*. The function values of these vertices are excluded from the smoothing process by constraining the function values to be 1. This ensures that the aneurysm does not shrink. Instead the aneurysm grows in particular in the concave regions, which addresses the second issue discussed above. The smoothing process can be started and stopped by a user, or, when used as an automatic post process, we found 20 iteration and a step size of $\tau = 1$ to deliver convincing results in our experiments. The supplementary material show some experiments with different values for the steps size and indicates that by setting $\tau = 5$ and the iterations to 4, the smoothing process can be accelerated without degrading quality. Figures 7 and 8 show results of the smoothing process for *case 2* and *case 19* from [PSFC17]. Fig. 8 illustrates the evolution of the ostium curve with the iterations of the smoothing process.

4.6. Report Generation

After we identified the aneurysms and smoothed the ostium, we can generate the report that includes the relevant information about the data set. Since the aneurysms are segmented from the vessel, the morphological properties can be estimated directly from the mesh. We refer to Meuschke et al. [MGW*18] for an overview of important properties and how they can be estimated. In addition to the estimated values, renderings of the aneurysms that include illustrations of the morphological descriptors are added to the report. Example images are shown in Fig. 5. For every aneurysm, the report includes four renderings, which provides the physicians with depictions from different anatomical perspectives and is intended to help them in assessing the anatomical conditions. To generate the renderings, we define four camera settings c_1 to c_4, each comprising a viewpoint, a position and an up vector of the virtual camera. The view direction of the camera can be calculated from the viewpoint and the position of the camera. The first two settings c_1 and c_2 show the aneurysm from the front and back sides and c_3 and c_4 show the aneurysm from the left and right. The same viewpoint $\mathbf{p}_v$, which is a point in the center of the aneurysm, is used for all camera settings. Explicitly, the viewpoint is the intersection point of the aneurysm's height vector $\mathbf{v}_h$, which is the vector between the center of the ostium and the aneurysm's dome position, and the aneurysm's width vector $\mathbf{v}_w$, which is the vector defining the maximal aneurysm extent orthogonal to the height vector. As up vector for all camera settings, we use the aneurysm's height vector $\mathbf{v}_h$. The position of the camera varies for the four camera settings. For c_1 and c_2, we define a direction vector $\mathbf{v}_d$ orthogonal to the aneurysm height and width $\mathbf{v}_d = \mathbf{v}_h \times \mathbf{v}_w$. Then, the camera position $\mathbf{p}_{c1}$ of c_1 is $\mathbf{p}_{c1} = \mathbf{p}_v + (diameter \cdot \mathbf{v}_d)$ and the camera position of c_2 is $\mathbf{p}_{c2} = \mathbf{p}_v - (diameter \cdot \mathbf{v}_d)$, where $diameter$ is the maximum extent of the aneurysm. For the remaining two camera settings, we use the width vector $\mathbf{v}_w$, which the difference vector of the two points $\mathbf{p}_{ws}$ and $\mathbf{p}_{we}$ of the aneurysm. The camera position $\mathbf{p}_{c3}$ of c_3 is $\mathbf{p}_{c3} = \mathbf{p}_{we} + (diameter \cdot \mathbf{v}_w)$, and the camera position of c_4 is $\mathbf{p}_{c4} = \mathbf{p}_{ws} - (diameter \cdot \mathbf{v}_w)$. To avoid occlusions of the aneurysm by the parent vessel geometry, we depict the surface mesh semi-transparently. Moreover, we depict the morphological descriptors within the aneurysm, where we restrict the rendering to the height and width descriptor, the two most important aspects to avoid visual clutter. To obtain an overview screenshot of the whole surface, we apply the technique for the automatic viewpoint selection proposed by Neugebauer et al. [NLBP13]. Finally, LaTex code is generated and used to produce a PDF that contains an overview image, the four screenshots of every aneurysm, and a table listing the morphological descriptors.

5. Performance

The performance was tested on a Microsoft Surface Book 2 with 16 GB RAM, Intel Core i7-8650U Quad-Core 4,2 GHz, NVIDIA GeForce GTX 1050 with 2 GB GDDR5. Our algorithm comprises the steps listed in Sec. 4. We used surface meshes with different resolutions and timed every step 10 times. The timings were then averaged. A list of timings can be found in Tab. 1. For this, again we used the data provided by Pozo et al. [PSFC17] and reduced to different resolutions. Especially for the shape index a parallelizable version of the curvature estimation could be used to improve the result. However, we only determined the values once and saved them. Moreover, the values for other surface scale according to the other values, thus, we decided to list only a few values in the table.

6. Evaluation of Requirements

We evaluated our program with a neuroradiologist, who has 16 years of work experience and who regularly treats cerebral aneurysms. This evaluation should assess whether we fulfilled the requirements (recall Sec. 3) and provide advice on how our program could be improved. In the first phase of this evaluation, we introduce our framework and acquaint the physician with all important features. Then, we gave him the opportunity to load data himself and explore them. For every requirement, we asked him to

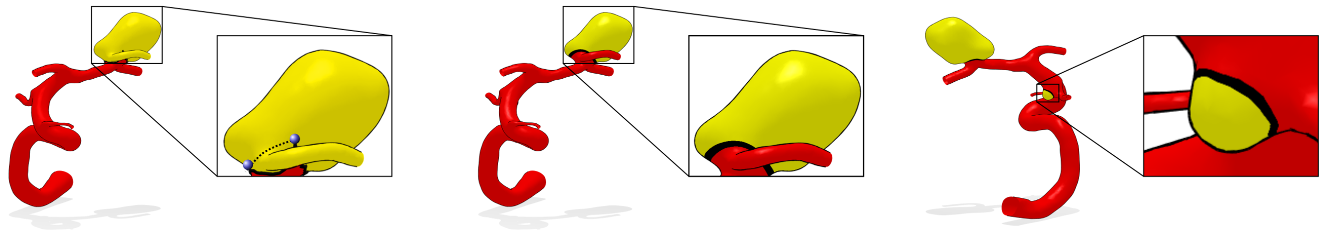

Figure 9: Left the result by Meuschke et al. [MGW*18] is shown. The ostium curve should follow the dashed line as shown by our approach. Additionally, with our method the smaller aneurysm is detected.

test it on a few data sets and to give us his opinion. During the evaluation, we asked the participant to think-aloud such that we could take notes.

6.1. Evaluation of Req. 1

In order to test for requirement Req. 1, our domain expert loaded every data set from the aneurysm database provided by Pozo et al. [PSFC17] and checked whether our algorithm correctly detects and segments the aneurysm. During the session, the smoothing parameter γ was set to the default value stated in 4.2. In addition to the segmented aneurysms, the expert also inspected the candidate list constructed in the first stage of our algorithm. The expert found our algorithm to detect and segment the aneurysms correctly and confirmed that our classification of the candidates has high accuracy, as reported in Sec. 4.3. The expert suggest that the program should make the candidate list available to the user on request as in case the algorithm misclassified a candidate, the user could easily correct the result. We incorporated this feature request to our program. We met again with the neuroradiologist and he was satisfied with the revised program.

6.2. Evaluation of Req. 2

The typical workflow, as observed during the evaluation, was that the domain expert first visually inspected the surface mesh to identify the aneurysm himself and then checked the results of our algorithm. As first feedback the results after the classification stage are shown to the user. If multiple aneurysms are detected each receives a different color. As can be seen from the timings list in Tab. 1, this process requires only few seconds. The expert used the brushing tool to modify and adjust the results. Since the optimization and classification stages are fast, the algorithm incorporates the user's strokes in less than a second. Once he was satisfied with the result, he pressed the button to determine the ostium curve, which triggers smoothing process. The expert was surprised that the smoothing process took longer than the other stages of the algorithm. For many examples, he reduced the number of smoothing iteration to ten to shorten the execution time. In summary, he was positively surprised with the performance.

6.3. Evaluation of Req. 3

After the aneurysm identification and the smoothing process, the ostium curve is extracted, which is needed for the estimation of the morphological parameters. The expert mentioned that the automatic estimation saves much time as he has to do this manually during the clinical workflow. Moreover, he emphasized the value of the automatic computation as the manual computation is error-prone since it is difficult to estimate a line from the height vector, which is perpendicular. He positive mentioned the visualizations of the morphological parameters, see Fig. 5 for an example. He found these to be very helpful to acquaint himself with the data. Altogether, he rated the automatic estimation of the morphological descriptors as very helpful and intuitive. In addition, he made one suggestion, which we consider as future work. He was curious to know how the morphological parameters change when the ostium curve slightly varies. We agree that it would be interesting to experimentally explore the stability of the morphological parameters with respect to variations of the ostium curve and to extend our program to support such exploration.

6.4. Evaluation of Req. 4

The domain expert also looked at the automatically generated reports for all inspected data sets. He wanted to check whether the important information was provided by the reports. He was very happy with the reports and stated that he considers this very helpful support for board meetings in which cases are discussed. As a feature request, he asked to allow users to add additional screenshots and annotate them. The resulting illustrations should be integrated to the report. Furthermore, he pointed out that he would like to be able to add notes to the report, e.g., concerning comorbidity. In summary, he was satisfied with the report.

6.5. Comparison with State of the Art

In addition to the evaluation with a domain expert, we compared our approach with the recent technique proposed by Meuschke et al. [MGW*18]. Their technique is based on a thinning method. First, they thin the mesh and extract the mesh skeleton. From every skeleton end node a path to the closest junction node is determined. For every path, the corresponding vertices on the surface mesh can be obtained. Afterwards, the region with the highest mean shape index is a designated aneurysm region. This gives the aneurysm region but only if an aneurysm occurs in the mesh. Afterwards, four characteristic points on the ostium are determined and connected using the the shortest edge path on the mesh, which is computed with Dijkstra's algorithm. Fig. 9 (left), shows the result of the approach on *case 1* from [PSFC17]. The ostium curve is wrongly

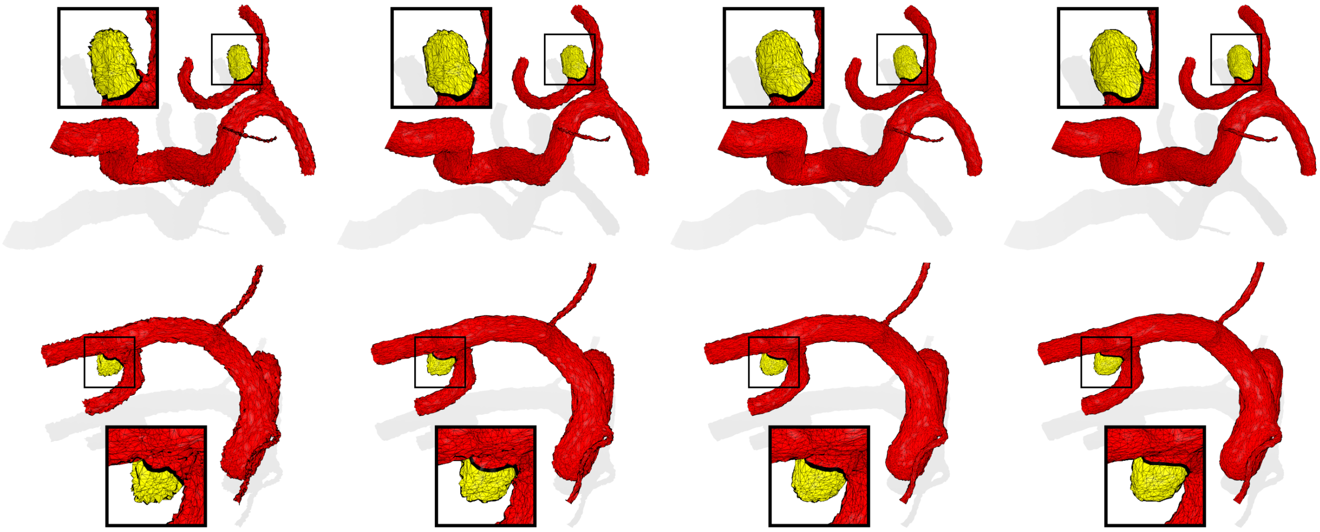

Figure 10: The results of our algorithm tested with different level of noise. The first row shows *case 31* with an average edge length of 0.3713 and the second row shows *case 37* with an average edge length of 0.3620. We randomly shifted the vertex positions along normal direction with a random value of $[0,1/2],[0,1/3],[0,1/4],[0,1/5]$ for the column from left to right.

constructed, as it is below and not above the outgoing vessel. One would expect to have a curve similar to the one indicated by the dashed lines. Our approach constructs the ostium curve correctly. Moreover, our approach additionally detects the smaller aneurysm on the same vessel.

7. Discussion and Further Evaluation

In this section, we provide further evaluation and discussion of our method. We discuss the effect of variation of the smoothing parameter γ on the results, evaluate the classifier, and test our method on input meshes with different mesh resolutions, geometric noise in the vertex positions and on meshes directly extracted from 3D-image data with standard software tools.

7.1. The Smoothing Parameter γ

To increase usability of our method, we specified a default value for the method's main parameter, the smoothing parameter γ, see Section 4.2. We experimentally found a standard setting for this parameter, $\gamma = 4.1$, that was used for all examples shown in the paper

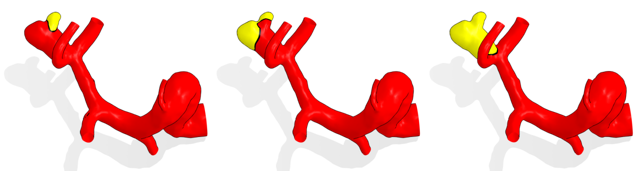

Figure 11: The influence of the parameter γ on the results is illustrated. The values of γ from left to right are $1,3,4$.

and the supplementary material. Still, we want to discuss the effect of this parameter on the results. The lower γ, the stronger the influence of the data term $D_i(l_i)$. For example, for $\gamma = 0$, the data term dominants. Since it is a unary term, the label of a triangle depends only its shape index. It is set to *aneurysm* if the shape index is larger than 0.75 or to *vessel* if the shape index is less than 0.75. Increasing γ increases the influence of the regularization term. Then, the decision for every label depends on the global context. For example a set of triangles whose shape index indicates an aneurysm may be assigned the label *vessel* if they are surrounded by triangles with label *vessel*. The other way around, a connected set of triangles that have the label aneurysms may grow in order to reduce the length of the boundary curve, which is penalized by the binary term. Results obtained with different γ values on *case 48* are shown in Fig. 11. When the value of γ is decreased to 3 and 1, the aneurysm is still detected but only a part of the aneurysm is segmented.

7.2. The Classifier

Our main source of data is the collection of aneurysms provided by Pozo et al. [PSFC17]. The results for all data sets of the collection are shown in the supplementary material. These results were all generated without user interaction and the same default value for the smoothing parameter in the optimization and 20 iterations for the smoothing process. Moreover, we used the classifier that was trained on artificial data for the shown results, hence, no data from the collection [PSFC17] was used for training. The results was evaluated by a domain expert as discussed in the previous section. Our classification achieved an accuracy 99%, which we want to discuss. There was one wrongly classified patch, which is was a false positive, i.e., we had a vessel patch as ground truth, which was predicted as an aneurysm patch. The wrong classified data is *case 33*. The problem can be solved by setting γ to 4.625. A second issue was that in *case 17*, the optimization missed to identify an aneurysm. By setting γ to 1, the aneurysm can be found identified. However, it is connected to large parts of the outgoing vessel. For this case, additional user input is necessary to segment the aneurysm only.

7.3. Mesh Resolution

For the results on 50 test cases shown in the supplementary material, we decreased the number of triangles to 20,000 to reduce the computation time. To justify this step, we conducted experiments to evaluate the effect of mesh coarsening on the results. We compared results of our method on meshes approximating the same vessel surfaces with different mesh resolutions. Results for *case 12* are shown in Fig. 13. In our experiments, variations of more than an order of magnitude in the number of triangles resulted in visibly similar results and negligible differences in the estimated morphological parameters.

7.4. Noisy Surfaces Meshes

Another aspect to evaluate is how our method can deal with noise in the data. When a substantial level of noise is present, we need to be careful with estimating curvature from the noisy surface. This is a known problem and a quantitative comparisons of the performance of different curvature estimation schemes on noisy data can be found in [VVP*16]. The comparison indicates that schemes that estimate the curvature over larger neighborhoods, like [HP11], exhibit best accuracy on noisy data. For our experiments we opt for an alternative strategy, which is to explicitly smooth the surface prior to curvature computation. Since the curvature estimation is the only step that suffers from the geometric noise on the surface, we use the smoothed mesh only for the curvature computation and discard it after the computation. Only the estimated shape indices at the triangles are kept and used for next step of our algorithm. To keep our approach simple, we evaluate our method with a simple Laplace smoothing process and the same curvature estimation scheme as used for all other experiments. Laplace smoothing is an iterative procedure and requires users to stop the smoothing process when the desired level of smoothness is achieved. Since we try not to hand such tasks to the users, we used a heuristic stopping criterion for the smoothing process in our experiments. The stopping criterion computes the local deviation of the vertex normals to assess the level of smoothness of the surface. Before the first and after every iteration of Laplace smoothing, we compute the mean of the dot products of the vertex normals of all pairs of vertices that are connected by an edge. The smoothing process is terminated if either the mean dot product exceed 0.98 or the mean value decreases compared to the mean value computed in the previous iteration. We want to mention that an alternative to the discussed simple heuristic stopping criterion would be to use an optimization-based smoothing scheme, like [HP07]. This scheme minimizes a fairness energy subject to spatial constraints that are derived from accuracy margins of the data which often can be estimated, e.g., they are provided by the manufacturer of a measuring device or relate to the spatial resolution used for 3D-imaging. Fig. 10 shows results of an experiment, in which normal noise was added to the surfaces of the vessels (*case 31, 37*) used as input to our method. For both vessel surfaces, we added random noise with four different levels of strength (random numbers drawn from the intervals $[0,1/2],[0,1/3],[0,1/4],[0,1/5]$). In case of the strongest noise level, the smoothing algorithm took less than 2 seconds for both cases (both with about 10k vertices). As the results demonstrate, the noise had just little effect on the aneurysm detection and segmentation.

7.5. 3D-Image Data

So far our evaluation focused on the aneurysm database of Pozo et al. [PSFC17]. In this paragraph, we extent the evaluation and report on how our approach can be applied to raw image data, which often is the starting point in a medical workflow. To convert the raw image data to a surface mesh, we use the publicly available software Slicer (`www.slicer.org`) [FBKC*12]. Our medical partner provided us with 3D-magnetic resonance angiography time-of-flight (3D-MRA-TOF) image data. After loading the provided Dicom data, our first step is to segment the vessels. This procedure is done automatically with Dicom's default segmentation module. Due to the high contrast of the vessels compared to the surroundings, the module could perform the segmentation without a need to change the default parameters. Slicer allows us to convert the resulting segmented vessels to a 3D surface mesh, which is shown in Fig. 12. The Dicom data set has a resolution of 384×224 and consists of 185 images. The 3D-surface mesh produced by Slicer can directly be used as input for our algorithm. The result of our method on this data is shown in Fig. 12. Though this is large data set containing many vessels, the aneurysm was correctly detected and segmented. The surface mesh contains 94,249 vertices and 174,264 triangles. The computation time of the segmentation and the surface extraction in Slicer was less than 1 second. Results for a second data set are shown in Fig. 14. In this case, the raw data was produced with digital subtraction angiography. The image data dimensions are $256 \times 256 \times 222$. The extracted surface mesh has $37,296$ vertices and $65,112$ triangles, extraction again took Slicer about 1 second.

8. Conclusion and Future Work

In this paper, we present a system for detecting and segmenting multiple aneurysms on vessels. The system combines a local shape descriptor and a combinatorial optimization problem to generates a list of candidate aneurysms with a classifier that separates the true aneurysms from the false positives in the candidate list. The optimization problem is modeled such that it can be efficiently solved with a graph cut algorithm. The low computation time allows us to integrate a brushing tool that allows users to adjust and correct results of the optimization to our system. Furthermore, we developed a smoothing scheme that allows to improve the segmentation and results in a smooth ostium curve. Finally, a report, which includes renderings of the aneurysms from different perspectives and values as well as illustrations of the morphological descriptors, is automatically generated. The importance and usefulness was shown in

Figure 12: Results when testing our system on surfaces meshes directly extracted from 3D MRA TOF image data using the Slicer software.

our evaluation and the supplementary materials shows the results of our algorithm on the data set provided by Pozo et al. [PSFC17].

For future work, we like to include flow information in the report, since hemodynamic information such as specific flow patterns also seem to influence the aneurysm rupture risk and treatment success [CVS*15]. One possible direction for achieving this would be to follow the approach of Meuschke et al. [MVPL18, MOJB*18] to cluster path lines and to classify the resulting blood flow patterns. An arising question is how such additional information, like the wall shear stress and flow velocities, could be effectively integrated to the generated report. In the case of multiple detected aneurysms, we want to further support the decision-making process. Several questions arise from a medical point of view. First, the physician has to decide whether an aneurysm should be treated and if so, which aneurysm should be treated first. It may happen that the treatment of one aneurysm affects the flow behavior, leading to different flow patterns in the other aneurysms. One could try to visualize the change such that the physician can inspect the difference and the effects of the treatment. For these questions also

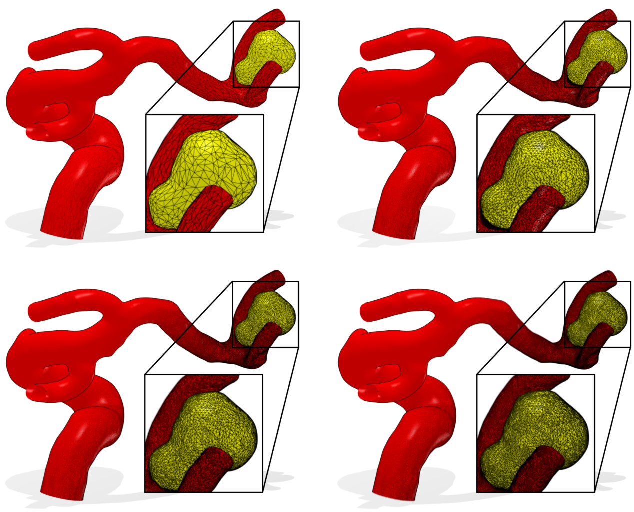

Figure 13: The results of our algorithm tested with different resolutions of the vessel surface mesh. From top left to bottom right: meshes with 10.000, 70.000, 150.000, 200.000 triangles.

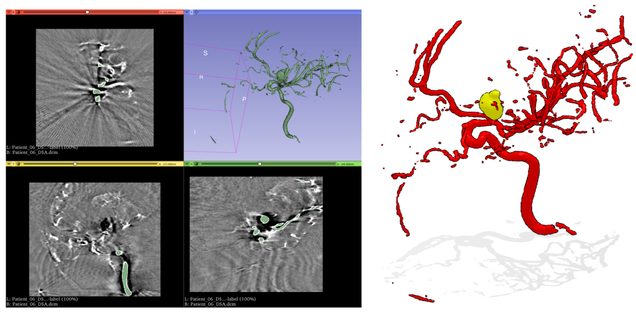

Figure 14: Results of our method tested on surface meshes extracted from digital subtraction angiography image data using Slicer.

the visualization of different treatment options plays an important role. For the screenshot generation, we plan to add a more sophisticated approach, that also analyzes the occlusion of outgoing vessel as well as the visibility of certain quantities [MEB*17]. Moreover, it would be interesting to explore the detection and segmentation of fusiform shaped aneurysms in future work.

References

[AABH*16] Al-Awami A. K., Beyer J., Haehn D., Kasthuri N., Lichtman J. W., Pfister H., Hadwiger M.: Neuroblocks–visual tracking of segmentation and proofreading for large connectomics projects. *IEEE Trans Vis Comput Graph 22*, 1 (2016), 738–46. 3

[APB*08] Antiga L., Piccinelli M., Botti L., Ene-Iordache B., Remuzzi A., Steinman D. A.: An image-based modeling framework for patient-specific computational hemodynamics. *Medical & Biological Engineering & Computing 46*, 11 (Nov 2008), 1097. 3

[BBJ*17] Bannach A., Bernard J., Jung, Florian Kohlhammer J., May T., Scheckenbach K., Wesarg S.: Visual Analytics for Radiomics: Combining Medical Imaging with Patient Data for Clinical Research. In *Workshop on Visual Analytics in Healthcare* (2017), p. to appear. 3

[BK04] Boykov Y., Kolmogorov V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. Pattern Anal. Mach. Intell. 26*, 9 (Sept. 2004), 1124–1137. 5

[BKP*10] Botsch M., Kobbelt L., Pauly M., Alliez P., Levy B.: *Polygon Mesh Processing*. AK Peters, 2010. 7

[BVZ01] Boykov Y., Veksler O., Zabih R.: Fast approximate energy minimization via graph cuts. *IEEE Trans. Pattern Anal. Mach. Intell. 23*, 11 (Nov. 2001), 1222–1239. 5

[CM16] CHANDRA A., MONDAL S.: Amalgamation of iterative double automated thresholding and morphological filtering: a new proposition in the early detection of cerebral aneurysm. *Multimedia Tools and Applications* (2016). 2

[CPB*11] CÁRDENES R., POZO J. M., BOGUNOVIC H., LARRABIDE I., FRANGI A. F.: Automatic aneurysm neck detection using surface voronoi diagrams. *IEEE Trans Med Imaging 30*, 10 (2011), 1863–76. 3

[CSM03] COHEN-STEINER D., MORVAN J.-M.: Restricted Delaunay triangulations and normal cycles. *ACM Symposium on Computational Geometry* (2003), 312–321. 4

[CVS*15] CEBRAL J. R., VAZQUEZ M., SFORZA D. M., HOUZEAUX G., ET AL.: Analysis of hemodynamics and wall mechanics at sites of cerebral aneurysm rupture. *J Neurointerv Surg 7*, 7 (2015), 530–36. 2, 12

[FBKC*12] FEDOROV A., BEICHEL R., KALPATHY-CRAMER J., FINET J., FILLION-ROBIN J.-C., PUJOL S., BAUER C., JENNINGS D., FENNESSY F., SONKA M., BUATTI J., AYLWARD S., MILLER J. V., PIEPER S., KIKINIS R.: 3d slicer as an image computing platform for the quantitative imaging network. *Magnetic Resonance Imaging 30*, 9 (07 2012), 1323–1341. 11

[FHP*09] FORD M., HOI Y., PICCINELLI M., ANTIGA L., STEINMAN D.: An objective approach to digital removal of saccular aneurysms: technique and applications. *Br J Radiol 82*, 1 (2009), 55–61. 3

[GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques* (1997), pp. 209–216. 6

[HBNT11] HENTSCHKE C. M., BEUING O., NICKL R., TÖNNIES K. D.: Automatic cerebral aneurysm detection in multimodal angiographic images. In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)* (2011), IEEE, pp. 3116–3120. 2

[HHFP11] HASSAN S., HÉTROY F., FAURE F., PALOMBI O.: Automatic localization and quantification of intracranial aneurysms. In *Computer Analysis of Images and Patterns* (2011), pp. 554–62. 2

[HP07] HILDEBRANDT K., POLTHIER K.: Constraint-based fairing of surface meshes. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Geometry Processing* (2007), pp. 203–212. 11

[HP11] HILDEBRANDT K., POLTHIER K.: Generalized shape operators on polyhedral surfaces. *Computer Aided Geometric Design 28*, 5 (2011), 321–343. 4, 11

[HS04] HECHENBICHLER K., SCHLIEP K.: Weighted k-nearest-neighbor techniques and ordinal classification. *Discussion paper, LudwigsâĂŞMaximilians University Munich 399* (01 2004), 16. 6

[JPLŠ15] JERMAN T., PERNUŠ F., LIKAR B., ŠPICLIN Ž.: Computer-aided detection and quantification of intracranial aneurysms. In *Medical Image Computing and Computer-Assisted Intervention* (2015), pp. 3–10. 3

[KAB*04] KARMONIK C., ARAT A., BENNDORF G., AKPEK S., KLUCZNIK R., MAWAD M. E., STROTHER C. M.: A technique for improved quantitative characterization of intracranial aneurysms. *AJNR 25*, 7 (2004), 1158–61. 2

[KFB*13] KARIM S., FEGELER C., BOECKLER D., SCHWARTZ L. H., KAUCZOR H.-U., VON TENGG-KOBLIGK H.: Development, implementation, and evaluation of a structured reporting web tool for abdominal aortic aneurysms. *JMIR Res Protoc 2*, 2 (2013). 3

[KvD92] KOENDERINK J. J., VAN DOORN A. J.: Surface shape and curvature scales. *Image and Vision Computing 10*, 8 (1992), 557–564. 4

[LMFM10] LAURIC A., MILLER E., FRISKEN S., MALEK A. M.: Automated detection of intracranial aneurysms based on parent vessel 3d analysis. *Medical image analysis 14*, 2 (2010), 149–159. 2

[LVUC*11] LARRABIDE I., VILLA-URIOL M. C., CÁRDENES R., POZO J. M., MACHO J., ROMAN L. S., ET AL: Three-dimensional morphological analysis of intracranial aneurysms: A fully automated method for aneurysm sac isolation and quantification. *Medical physics 38*, 5 (2011), 2439–49. 3

[MEB*17] MEUSCHKE M., ENGELKE W., BEUING O., PREIM B., LAWONN K.: Automatic Viewpoint Selection for Exploration of Time-dependent Cerebral Aneurysm Data. In *Bildverarbeitung fÃijr die Medizin (BVM)* (March 2017), Springer, pp. 352–357. 12

[MGW*18] MEUSCHKE M., GÜNTHER T., WICKENHÖFER R., GROSS M., PREIM B., LAWONN K.: Management of Cerebral Aneurysm Descriptors based on an Automatic Ostium Extraction. *IEEE Computer Graphics and Applications 38*, 3 (2018), 58–72. 2, 3, 8, 9

[MOJB*18] MEUSCHKE M., OELTZE-JAFRA S., BEUING O., PREIM B., LAWONN K.: Classification of Blood Flow Patterns in Cerebral Aneurysms. *IEEE Transactions on Visualization and Computer Graphics* (2018). 12

[MSM*10] MOHAMED A., SGOURITSA E., MORSI H., SHALTONI H., MAWAD M. E., KAKADIARIS I. A.: Computer-aided planning for endovascular treatment of intracranial aneurysms (capeta). *SPIE: Medical Imaging 2010 7625* (2010). 3

[MVPL18] MEUSCHKE M., VOSS S., PREIM B., LAWONN K.: Exploration of Blood Flow Patterns in Cerebral Aneurysms during the Cardiac Cycle. *Computers & Graphics 72* (2018), 12–25. 12

[NDSP10] NEUGEBAUER M., DIEHL V., SKALEJ M., PREIM B.: Geometric Reconstruction of the Ostium of Cerebral Aneurysms. In *VMV* (2010), pp. 307–14. 3

[NLBP13] NEUGEBAUER M., LAWONN K., BEUING O., PREIM B.: Automatic Generation of Anatomic Characteristics from Cerebral Aneurysm Surface Models. *International Journal of Computer Assisted Radiology and Surgery 8(2)* (March 2013), 279–289. 8

[Ots79] OTSU N.: A Threshold Selection Method from Gray-level Histograms. *IEEE Transactions on Systems, Man and Cybernetics 9*, 1 (1979), 62–66. 3

[PSFC17] POZO SOLER J., FRANGI A. F., CONSORTIUM T. N.: Database of cerebral artery geometries including aneurysms at the middle cerebral artery bifurcation, May 2017. URL: https://figshare.com/articles/Database_of_Cerebral_Artery_Geometries_including_Aneurysms_at_the_Middle_Cerebral_Artery_Bifurcation/4806910/1. 2, 6, 7, 8, 9, 10, 11, 12

[PWN*15] PANKAU T., WICHMANN G., NEUMUTH T., PREIM B., DIETZ A., STUMPP P., BOEHM A.: 3D model-based documentation with the Tumor Therapy Manager (TTM) improves TNM staging of head and neck tumor patients. *International journal of computer assisted radiology and surgery 10*, 10 (2015), 1617–1624. 3

[PWY*07] POTTMANN H., WALLNER J., YANG Y., LAI Y., HU S.: Principal curvatures from the integral invariant viewpoint. *Computer Aided Geometric Design 24*, 8-9 (2007), 428–442. 4

[Rus04] RUSINKIEWICZ S.: Estimating curvatures and their derivatives on triangle meshes. In *Symposium on 3D Data Processing, Visualization, and Transmission* (Sept. 2004). 4

[VVP*16] VÁSA L., VANECEK P., PRANTL M., SKORKOVSKÁ V., MARTÍNEK P., KOLINGEROVÁ I.: Mesh statistics for robust curvature estimation. *Comput. Graph. Forum 35*, 5 (2016), 271–280. 4, 11

[WEMM*02] WESTIN C.-F., E. MAIER S., MAMATA H., NABAVI A., JOLESZ F., KIKINIS R.: Processing and visualization of diffusion tensor mri. *Medical Image Analysis 6* (06 2002), 93–108. 6

[WVS*14] WIBMER A., VARGAS H. A., SOSA R., ZHENG J., MOSKOWITZ C., HRICAK H.: Value of a standardized lexicon for reporting levels of diagnostic certainty in prostate mri. *American Journal of Roentgenology 203*, 6 (2014), W651–W657. 3

[YBVRZ01] Y. BOYKOV Y., VEKSLER R. O., ZABIH R.: Efficient approximate energy minimization via graph cuts. *IEEE Trans. Pattern Anal. Mach. Intell. 20* (01 2001). 5