

VesselMap: A Web Interface to Explore Multivariate Vascular Data

Jun Tao^a, Xiaoke Huang^b, Feng Qiu^c, Chaoli Wang^a, Jingfeng Jiang^d, Ching-Kuang Shene^e, Ye Zhao^b, Daphne Yu^c

^aDepartment of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN 46556

^bDepartment of Computer Science, Kent State University, 800 East Summit Street, Kent, OH 44240

^cSiemens Corporate Technology, 755 Colledge Road East, Princeton, NJ 08540

^dDepartment of Biomedical Engineering, Michigan Technological University, 1400 Townsend Drive, Houghton, MI 49931

^eDepartment of Computer Science, Michigan Technological University, 1400 Townsend Drive, Houghton, MI 49931

Abstract

Hemodynamics plays a key role in the pathological evolution of vascular diseases such as vascular stenosis, plaques and aneurysms. Augmented visualization of relevant hemodynamic and geometric data accelerates engineering designs of diagnostic and therapeutic methods for those vascular diseases. However, four dimensional (4D) hemodynamic data are intrinsically complex. Consequently, exploration of this information has not been streamlined, despite its importance. In this paper, we propose a web interface to explore multivariate hemodynamic data, thereby facilitating the visualization of spatial and temporal relationships among those hemodynamic parameters of interest. The main thrust of this web system is the proposed VesselMap methodology, a 2D representation of vessel structures that provides an overview of all query results without any visual occlusions. Furthermore, in this framework, quantile-quantile plot was also used to compare spatial variations in parameter distributions, thereby greatly aiding in extracting local characteristics of hemodynamic parameters.

Using “patient-specific” computational fluid dynamics (CFD) simulations, realistic flow velocity fields were obtained. Based on those CFD-simulated velocity fields, we performed an empirical evaluation to confirm the usefulness of our approach for biomedical engineering applications. We concluded that the proposed VesselMap methodology allowed users to rapidly interact with hemodynamic data of interest, enhancing efficacy of information retrieval.

1. Introduction

Visualization of medical data has already been used extensively to show anatomical structures topologically using photo-realistic surfaces or volume rendering. Moving forward, a more integrated visualization approach is needed in order to quantitatively connect processes of a disease with abnormalities at various spatial and temporal levels in biomedical research. Integral to this approach is an efficient way to retrieve quantitative information from complex medical data. In this work, visualization of vessel data in and around cerebral aneurysms was used to demonstrate a novel web-based visualization approach of this kind. Vascular data in this paper are referring to complex three-dimensional geometric data, hemodynamic data (e.g., time-resolved 3D blood velocity) and other parameters derived from hemodynamics.

We envision that visual exploration of hemodynamic parameters and their spatial distributions facilitates engineering design of therapeutic methods toward treating cerebral aneurysms. Benefits of having such a web-based visualization tool are twofold. First, the proposed tool provides a flexible but systematic way for biomedical engineers to explore flow data with their own criteria. For instance, some vascular stents were designed to open up arteries (i.e., structural integrity of vessel) while other stents were designed to divert blood flow. Second, interdisciplinary studies of this kind often involve researchers located at different geographical locations. A web-based tool will enable collaborative reviewing of experimental or computer simulation

results, reducing burdens on engineers for conducting a large-scale study. Furthermore, we intend to provide an unobscured viewing of 3D and 4D vessel data which may enhance the collaborative reviewing process over the Internet. It is recognized that labeling a 3D volume in a consistent way and communicating it with peers over the Internet could be a difficult task, because the relative positions of vessel branches may vary due to the change of viewpoints. Now there are a large amount of time-resolved flow data spatially registered with the vessel anatomy. Web-based collaborations involving in quantitative flow analysis without the right tools would be much more difficult.

Toward this end, we propose VesselMap, a novel web-based solution to assist biomedical experts in reviewing vessel data and analyze the relationships among different properties of blood flow. The VesselMap is centered on a scheme that enables 2D illustrative visualization of parameters for true 3D or 4D data. Using this scheme we first flatten the 3D vessel structure and its corresponding parameters onto a 2D plane. Then, all subsequent interactions can be operated in 2D. This provides an easier way to specify and label regions of interest and eliminates potential (angle-dependent) viewing occlusions. It also serves as a clearer and paper-friendly overview, since rotation and zooming are not available on a printed report. Furthermore, a segmentation scheme is used to divide data into sub-groups for any parameters of interest, in order to provide comparative investigations of flow characteristics. Data from sub-groups

generated from the aforementioned segmentation scheme can be effectively displayed on the same 2D illustration, and their inter-group differences are evaluated and displayed in a difference matrix. It is worth noting that an inherited benefit of the proposed visual exploration method stems from the availability of statistical information, such as histograms of different properties. Users can query flow characteristics within the entire volume or a region of interest by simply interacting with histograms. The proposed system is flexible and can be adapted for different applications. Particularly, in this paper, some benefits of the aforementioned visualization strategy are going to be demonstrated through visualization of particle transport in and around cerebral aneurysms. Clinical applications include deliveries of chemical agents to either dissolve clots in the blood stream [1], or occlude vascular aneurysms [2, 3] and stop hemorrhage from ruptured aneurysms or arteries [4]. Flow visualization is also important for designing flow-diverting stents to occlude cerebral aneurysms or bleeding sites. Finally, the web-based environment places a minimum amount of effort to setup and requires only the displayed information to be transferred.

2. Related Work

Computed tomography angiography (CTA) is commonly visualized with *maximum intensity projection* (MIP) and *direct volume rendering* (DVR) techniques. To enhance the perception of vessel structures, researchers often approximate the vessel surface using model-based or model-free surface rendering approaches. The model-based surface rendering approach utilizes the information of centerline and radius, and approximates the vessel surface using models, such as truncated cones [5], B-splines [6], subdivision surfaces [7], or convolution surfaces [8]. The model-free surface rendering approach extracts the iso-surface using algorithms such as marching cubes [9] based on a given threshold. Instead of approximating the vessel surface, Lathen et al. [10] proposed spatially varying transfer functions. This approach locally shifts the transfer function to enhance the perception of low intensity structures. Mistelbauer et al. [11] used halo rendering to enhance the lumen of a vessel structure. Schumann et al. [12] used the multi-level partition of unity implicits (MPUI) approach to reconstruct the surfaces.

Unlike approaches based on MIP or DVR, other approaches flatten the vessel structures and map the corresponding information to 2D images. One of the commonly used approaches in this category is *curved planar reformation* (CPR). Kanitsar et al. [13] introduced CPR as a curved cutting through the data set along the centerline of a single vessel. Then, they extended the CPR approach to multi-path CPR (mpCPR) that supports multiple vessel branches and spiral CPR that flattens the vessel along a spiral to show its interior [13]. Kretschmer et al. [14] extended the mpCPR approach and used a bilateral filtering to remove undesired depth discontinuities. Mistelbauer et al. [15] proposed an approach based on CPR that aggregates the information around the centerline along circular rays. Borkin et al. [16] introduced a 2D vessel visualization method that uses a tree diagram to represent the structure of a coronary artery tree.

Each branch is straightened and displayed as a tape with varying widths, which represents the diameter of the vessel. Zhu et al. [17] presented a work that produces a flattened visualization of vessel branches. Two algorithms were proposed in this work. The first one is a conformal mapping algorithm by minimizing two Dirichlet functionals, and the second one adjusts the conformal mapping to produce a flattened representation that preserves areas. Marino and Kaufman [18] flattened treelike structures based on their extracted skeletons. A radial planar embedding was adopted to layout the skeleton on a plane. Neugebauer et al. [19] designed a multi-perspective 2D projection map. Combined with a standard 3D visualization, a correlation tool was provided to make correspondence between the 2D map and 3D model. Won et al. [20] presented an uncluttered single-image visualization of vascular structures. The centerlines were first extracted, so that the vascular structures could be represented by multiple tubes. Each tube was then laid out on 2D by minimizing a score function to avoid occlusion. Instead of flattening the vessel structures, Neugebauer et al. [21] selected appropriate viewpoints to reveal the anatomic characteristics of cerebral aneurysms.

Other than scalar volumes, vascular data sets often come with a simulated blood flow field as well. Recently, different flow visualization techniques have been developed specifically for these data sets. van Pelt et al. [22] used various techniques to depict the blood flow and associated characteristics in different styles, together with an evaluation to measure the value of those visualization styles. Köhler et al. [23] extracted vortices in blood flow data sets using line predicates and highlighted the corresponding regions. van Pelt et al. [24] proposed to semi-automatically place and align a probe in the blood flow field, which serves as a seeding basis. Then, particles, integral lines and integral surfaces are used to convey distinct characteristics of the flow field. Born et al. [25] found the representatives of a line bundle, and used streamtapes with arrow heads to visualize the line bundles. The tape-like structure provides a clear picture of how the representatives diverge and merge. Oeltze et al. [26] proposed to cluster the streamlines and use the cluster representatives for a clear view. They conducted a qualitative study on using different similarity measures, including geometry-based similarities and attribute-based similarities. Angelelli and Hauser [27] presented a solution to straightening tubular flow for side by side visualization, which facilitated the comparison of different visualizations or time steps of the same flow.

Due to the presence of multiple fields, some vessel visualization approaches also provide contextual information. Straka et al. [28] proposed VesselGlyph which combines both DVR and CPR. It depicts the vessels using CPR which is naturally placed in a DVR anatomic context. Mistelbauer et al. [11, 15] provided optional context rendering that displays the volume outside the lumen of vessel as well. Gasteiger et al. [29] presented a focus+context approach called FlowLens that uses some predefined lens templates to combine visualization results of different properties. The property of focus and property in context are both selected by the users. Lawonn et al. [30] introduced an adaptive surface visualization of vessels to highlight

the animated pathlines and emphasize nearby surface regions. Lawonn et al. [31] combined the blood flow animation with wall thickness visualization. The wall thickness was mapped onto the vessel surface, which could be cut away dynamically to display a glyph-based blood flow animation.

Among the above approaches, those that flatten the vessel and produce 2D images are closest to our approach. However, unlike those approaches, our 2D visualization serves not only as an overview but also as an interface for convenient interaction. We provide very concise information on the 2D visualization as a guideline to explore the data. Our approach does not require the centerline and radius information as CPR-based techniques [32, 13, 15], the triangulated vessel surface as the conformal mapping technique [17], or others do [20, 18]. In addition, these approaches only visualize scalar information of the vessel, without providing the ability to trace particles in the vessels, which is important for discover hemodynamic characteristics, such as particle transport information. In terms of exploring multiple fields, our approach also differs from those focus+context techniques [28, 29]. Instead of visualizing multiple fields, we provide an interface guided by histograms of different properties to enable users to discover the dynamic statistical information for regions of interest during particle tracing. It also facilitates the process of finding a region with a certain feature, or exploring relationships among different regions.

3. Algorithm

Our exploration is guided by statistical information, mostly histograms of particles with different properties. The property values of particles are either interpolated in scalar volumes, e.g., pressure and wall shear stress (WSS), or derived from particle tracing, e.g., age. Users can freely interact with the histograms to select a combination of bins in one or more histograms to filter particles and highlight the corresponding regions. In addition, we propose an approach to map the 3D vessels to a 2D *VesselMap*, while preserving its perceptual structure. Users can brush the 2D *VesselMap* or select a group from our segmentation to specify a region of interest, and the histograms computation will be constrained to the user-specified region. The viewpoint of the 3D particle rendering can be automatically adjusted to focus on that region and reveal more details in the original 3D space. Users can further interact with histograms or select regions to gradually refine their query. They can also brush the 2D *VesselMap* to place a seeding plane from which the particles will first be released and then be traced.

We provide a segmentation of the vessel structure to compare the histograms in different regions, so that local behavior can be observed. The segmentation can be guided by different mapping functions to provide a variety of results. The differences among regions are derived from their histograms. In addition to the traditional measures which usually compute a single value to represent the difference, we visualize the differences using quantile-quantile (Q-Q) plots. For global investigation, we develop an interface that shows all differences between every pair of groups.




Figure 1: Mapping from the 3D volume space to a 2D *VesselMap* representation using the VDS1 data set. (a) shows the mesh structure of the *VesselMap*. (b) shows streamline visualization with a semi-transparent vessel structure.

Our system consists of a web front-end client for visualization and user interface, and a back-end server for computation. The server is responsible for histogram computation, *VesselMap* generation, and flow visualization. While the web front-end displays visualization results, it also interacts with users and requests the server to update visualization results using messages.

3.1. *VesselMap*: a 2D Representation

The entire volume is evenly partitioned into small blocks (e.g., $3 \times 3 \times 3$ in this paper), and only blocks containing the vessel structure are loaded. These 3D blocks approximate the vessel structure. We map the blocks to points in 2D, so that they can be displayed on a webpage more desirably. The mapped points are then triangulated to form a mesh structure representing the vessel structure. To build the connection between the original 3D volume and the 2D image, the mapping should preserve the local shape of vessels. We formulate this as a 2D graph layout problem, and achieve the desired layout using a minimization approach. Assuming the blocks that contain some of the vessel structure are $B = \{b_1, b_2, \dots, b_n\}$, the neighboring blocks of b_i are those whose distances to b_i are smaller than or equal to a given threshold δ , and the non-neighboring blocks are those whose distances to b_i are larger than δ . The energy of the mapping is defined as follows:

$$\sum_{b_i \in B} \sum_{b_j \in B} c_{ij}, \text{ where } c_{ij} = \begin{cases} w_n |d_{ij} - e_{ij}|, & e_{ij} \leq \delta, \\ w_{nn} |d_{ij} - g_{ij}|, & e_{ij} > \delta. \end{cases} \quad (1)$$

where c_{ij} is the cost between b_i and b_j , w_n and w_{nn} are weights for neighboring and non-neighboring blocks, respectively, d_{ij} is the distance between b_i and b_j in the 2D image, e_{ij} is their Euclidean distance in the original 3D volume and g_{ij} is their geodesic distance in the 3D vessel structure. Intuitively, for neighboring blocks, we preserve their original distances to maintain the local shape of vessels. For non-neighboring ones, using the geodesic distance will separate two blocks in different vessel branches to avoid occlusion. Figure 1 (a) shows the mapped points with triangulation, we observe that the shape is still similar to that in 3D, but occlusion is removed. Each point is displayed as a node whose color represents the distance from the corresponding data block to the vessel wall. A block that is in proximity of the vessel is colored in green. The color transits from green to red for blocks further away from the vessel

wall. Since the widths for most vessel branches are small when measured in number of blocks, we can only observe a few different colors. Visually, we find that red nodes are likely to stay around the centerline of the flattened vessel, consistent with our intuition. This indicates visualization through the flattened VesselMap can still be used to infer spatial characteristics, although spatial continuities among nodes is no longer valid.

Considering block centers as point clouds, we approximate the geodesic distance using the shortest path between two points. Given a threshold ϵ , we initialize the geodesic distance matrix as follows:

$$g_{ij} = \begin{cases} e_{ij}, & e_{ij} \leq \epsilon, \\ \infty, & \text{otherwise.} \end{cases} \quad (2)$$

where e_{ij} is the Euclidean distance between the centers of blocks b_i and b_j . Running an all-pair shortest distance algorithm, e.g., the Floyd-Warshall algorithm, on this distance matrix, we can obtain the approximated geodesic distance in $O(n^3)$ time, where n is the number of blocks. In our implementation, we use $\epsilon = 1.5$, meaning that at the initial stage, only the distances between neighbors are known, and the distances between non-neighboring blocks are then approximated by the shortest paths.

To obtain the positions of vertices $V = \{v_1, v_2, \dots, v_n\}$, where v_i is a 2D vertex that represents block b_i , the equations are represented as a linear system, so that the minimization is performed by solving this linear system in a least-squares sense. Since in these equations $d_{ij} = \sqrt{(v_i - v_j)^2}$ is not linear, we propose a scheme to solve them in multiple iterations. Let $V = \{v_1, v_2, \dots, v_n\}$ be the desired positions and $V' = \{v'_1, v'_2, \dots, v'_n\}$ be the current positions. Then a target distance t_{ij} is achieved by moving the two vertices along the line segment connecting them. Note that the target distance t_{ij} is e_{ij} for neighboring blocks and g_{ij} for non-neighboring blocks. In this way, the equation $|d_{ij} - t_{ij}|$ is rewritten in the form of

$$\left| (v_i - v_j) - t_{ij}(v'_i - v'_j) / (|v'_i - v'_j|) \right|.$$

Note that v_i and v_j are variables to solve, and $t_{ij}(v'_i - v'_j) / (|v'_i - v'_j|)$ is a constant. Therefore, the equation is a linear combination of the variables. All equations can then be formulated in a linear system $AV = C$, where A is an adjacency matrix, V is a vector of all vertices, and C is a vector of all constants. Each equation is represented by one row in the linear equation. For example, assume that $(v_i - v_j - c)^2$ is the k -th equation. Then the k -th row in matrix A is constructed as the following: its i -th element is set to 1, its j -th element is set to -1, and all the other elements are set to 0. And the k -th element in the vector C is set to c . Applying the least square method will provide us a new set of positions based on the current positions. We will repeat this procedure until the positions do not change or some predefined number of iterations is achieved. In our implementation, if the change of energy is smaller than 1×10^{-6} , the positions are considered to be stable, with the maximum number of iterations being 1500.

Viewpoint selection. When users brush a group of blocks to query, the viewpoint of the 3D rendering can be changed accordingly to focus on that region. We adopt a heuristic scheme




Figure 2: The automatic viewpoint selection when users brush the 2D VesselMap. (a) shows the selected blocks and their principal components together with the camera. (b) shows the rendered result under the automatically selected viewpoint.

to automatically select the viewpoint. Each block is considered as a data point located at the center of that block. We compute three principal components of the data points, and the variations along all principal component directions. The viewing direction is given by

$$var_3 \cdot p_1 + var_2 \cdot p_2 + var_1 \cdot p_3,$$

where p_1 , p_2 and p_3 are the first, second and third principal component directions, and var_1 , var_2 and var_3 are the variations along the three directions, respectively. Intuitively, if the variation of point positions along one direction is large, the viewing direction is desired to be perpendicular to that direction in order to preserve the variation. Therefore, the three principal component directions are weighted by their variations in the reversed order, so that the most significant component gets the smallest weight to increase its angle difference from the viewing direction. In Figure 2 (a), each block is represented by a sphere located at its center. The first, second and third principal components are indicated by the red, green and blue arrows, whose lengths are proportional to the variations along those directions. Since the variations along all directions are similar to each other, the camera tends to look at them in a balanced way. The projected plane is shown by the blue rectangle, whose center is determined by averaging the center positions of all selected blocks. The rendering result under the automatically selected viewpoint is shown in Figure 2 (b).

Seeding plane placement. In this application, all particles are released from an inlet plane specified by a user-defined plane in a vessel branch. Our 2D VesselMap can be used to place the seeding plane conveniently. Users can simply brush the VesselMap to select a group of blocks, and the seeding plane will be placed at the center of those blocks. While the size of the seeding plane is specified by the users, the normal direction is given by the flow direction in order to maximize the effective area for particle releasing.

3.2. VesselMap Segmentation

In the next stage, the VesselMap is segmented into regions. Inspired by Singh et al. [33], we develop this segmentation scheme to divide the vessel structure at the block level based on a mapping function that assigns each block a scalar value.




Figure 3: The nerve of a covering using a torus as an example. The torus is colored by a height function. This set of images demonstrates how a covering of the height function becomes a covering of the torus, whose nerve captures the topology of the torus with one connected component and a hole.

The regions generated from the segmentation serve as a starting point to investigate their local behaviors. The connection among the segmented regions forms what is known as the nerve of the VesselMap. The nerve of the VesselMap provides an abstract overview of the vessel data. Depending on the selection of the mapping function, the nerve of the VesselMap conveys different information. For example, if the geodesic distance from a block to the inlet is used, it represents the shape of the vessel structure; while if a scalar property in a block is used, the nerve of the VesselMap represents the topological structure of that property.

Background. The construction described in this section is motivated by the concept of the nerve of a covering in the field of topology. Given a scalar field over the domain, this provides an abstraction of how different value ranges of the scalar field are related to each other. We refer readers to [34] for background information on topological space and simplicial complex, and only provide an intuitive description here with an example. In this example, we construct the nerve of a covering that captures the topology of a torus. A finite *covering* is a collection of subsets whose union is the entire space. In the left part of Figure 3, an interval $I \in \mathbb{R}$ is colored from red to green to blue, and three smaller intervals are indicated by three line segments to its right: the red one at the top, the green one in the middle, and the blue one at the bottom. The three small intervals represented by these three lines form a covering of I , since their union covers the entire interval of I . However, there is no such covering immediately available for a torus. In this example, we obtain the covering of the torus through the covering of an interval. Consider the height function of the torus $f: T \rightarrow \mathbb{R}$, that maps each point of the torus T to a scalar value of height. We see that the covering of the range of this height function naturally becomes a covering of the torus, with a blue region at the bottom, two green regions in the middle, and one red region on the top. The nerve of this covering is then generated by connecting the overlapping regions. The right part of Figure 3 shows that the nerve captures the topology of the torus.

Our construction. Similarly, we obtain a covering of the blocks B by considering a map $f: B \rightarrow \mathbb{R}$, which assigns each block a real number. This map could be the distance to the inlet plane, or the average velocity magnitude in a block, etc. Assuming the range of f is $[r_{\min}, r_{\max}]$, the subsets $\{[r_{\min}, r_1], [r_2, r_3], \dots, [r_k, r_{\max}]\}$ is a covering of $[r_{\min}, r_{\max}]$, and $[r_{i-1}, r_i]$ and $[r_{i+1}, r_{i+2}]$ are two overlapping subsets. Considering the interval $[0, 10]$ as an example, $\{[0, 4], [3, 7], [6, 10]\}$ is such a cov-

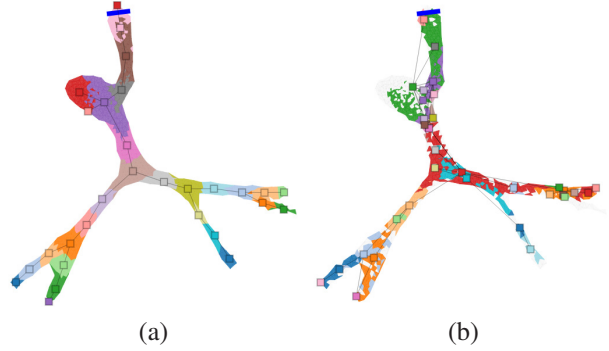


Figure 4: The segmentation and the nerve of the VesselMap using the geodesic distance to the inlet as the mapping function. The result is colored by groups of blocks. The rectangle nodes and the edges connecting them show the nerve of the VesselMap. (a) Blocks are grouped by their distances to the inlet. (b) Blocks are grouped by the vorticity values of the particles. In this case, the blocks that do not contain any particle are not grouped and are colored in white.




Figure 5: Q-Q plot examples. The correlation of quantiles is mapped onto the background color of the plot. The ranges of quantiles are mapped by two blue bars. The two distributions in (a) are very similar; those in (b) are less similar; and those in (c) are least similar.

ering of $[0, 10]$. In our implementation, we first determine the minimum and maximum of the mapped values, i.e., r_{\min}, r_{\max} . Then, we evenly divide the range into overlapping sub-ranges, and each block corresponds to one or more sub-ranges according to its mapped value. Finally, connecting the neighboring blocks that correspond to the same sub-range forms a set of connected components, which serve as the vertices in the nerve of the VesselMap, denoted as $V = \{v_1, v_1, \dots, v_m\}$, where m is the number of generated connected components. We name each connected component a group of blocks. Note that the connection relationship is computed from 3D blocks. Therefore, the topological structure will not be affected by the mapping from 3D to 2D.

Although the nerve of a covering may contain higher dimensional simplices, for the nerve of the VesselMap, we only consider simplices up to dimension one, i.e., vertices and edges. Therefore, only the connections between two vertices are taken into consideration. Following the previous definition, an edge (1-simplex) $e_{ij} = (v_i, v_j)$ is in the nerve of a covering, if and only if $U_i \cap U_j \neq \emptyset$, where U_i and U_j are subsets in the covering corresponding to v_i and v_j , respectively. In our context, if two groups of blocks share some blocks in common, an edge will be added between them in the nerve of the VesselMap.

However, this procedure generates segmentation results with overlapping. We further simplify the procedure and divide the range of f into non-overlapping sub-ranges, and therefore, the groups are non-overlapping as well. In this case, an edge $e_{ij} = (v_i, v_j)$ is in the nerve of the VesselMap, if a group represented

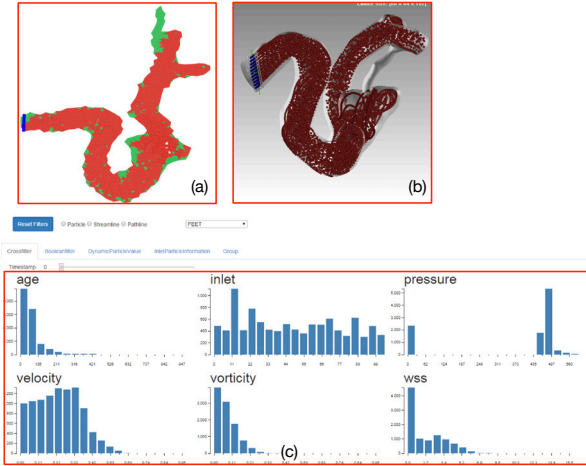


Figure 6: The VesselMap Interface. (a) shows the 2D VesselMap. (b) shows the 3D flow visualization. (c) shows the histogram visualization, which is a tab of the information interface.

by v_i contains a block b_p that is a neighbor of some block $b_q \in v_j$. Figure 4 shows two examples of the segmentation and the nerve of the VesselMap. In Figure 4 (a), the neighboring blocks whose geodesic distances to the inlet fall into the same range are grouped together. This segmentation produces similar results as those segmentation approaches considering the shape of the structure. In addition, the corresponding nerve demonstrates the topological structure of the vessel.

3.3. Comparing Regions and Properties

One of the major goals of our application is to guide users to explore the relationships among properties. However, we feel that the commonly used approaches, e.g., correlation coefficients, only provide one single value indicating the difference, which fails to answer how they differ from each other. In our approach, in addition to the traditional difference values, we use the Q-Q plot [35], to capture the details. Quantiles are points taken at regular intervals from the cumulative density function (CDF) of a random variable. The concept of quantiles differs from percentiles in the sense that quantiles are indexed by sample fractions instead of sample percentages. Generally, the p -th quantile $Q(p)$ of a random variable Z is the value z such that $P(Z \leq z) = p$. The Q-Q plot for two random variables Z_1 and Z_2 plots the quantile pairs $\langle Q_{Z_1}(p_1), Q_{Z_2}(p_1) \rangle, \langle Q_{Z_1}(p_2), Q_{Z_2}(p_2) \rangle, \dots, \langle Q_{Z_1}(p_n), Q_{Z_2}(p_n) \rangle$ on a 2D plane, where $p_1 < p_2 < \dots < p_n$ are equally spaced.

Figure 5 shows three examples of the Q-Q plot. The background color of a Q-Q plot represents the correlation of quantiles for the two corresponding groups. The green, brown and red colors indicate small, medium and large correlations, respectively. The color gradually transits in between. For a clearer view, the plotted points always occupy the entire space, i.e., they start from the bottom-left corner and end at the top-right corner. In addition, a blue bar is used to indicate the local value range in a group with respect to the global value range in the volume. In (a), the two distributions are almost identical and the points are mostly aligned along the diagonal. In (b), the two

distributions are less similar. The slope of the Q-Q plot is first steep and then becomes flat. This indicates that the distribution corresponding to the x -axis has slightly higher probabilities for the smaller values while the distribution corresponding to the y -axis has slightly higher probabilities for the larger values. In (c), the two distributions are least similar. The slope of the Q-Q plot is almost flat for a large portion of quantiles. This indicates that the distribution corresponding to the y -axis has very high probabilities for the smallest values. Thus, a Q-Q plot provides more information on how the two distributions differ from each other.

4. User Interface and Interaction

Our interface is shown in Figure 6. It consists of three components: the VesselMap in the top right region, the flow visualization in the top left region, and the information interface at the bottom. The VesselMap displays the 2D representation of the vessel structure. A blue bar is shown to indicate the inlet where particles are released, as shown in Figure 7 (b). If histogram, particle inlet information, or particle stay information is displayed, the triangles are colored in red or green: red indicates blocks in the triangle containing queried particles, and green indicates blocks in the triangle containing no queried particles. When group comparison is displayed, the triangles are colored according to the groups they belong to. Flow visualization displays the vessel surface as a semi-transparent mesh and the flow inside the vessel, as shown in Figure 7 (c). The flow can be visualized as particles, streamlines or pathlines. Flow visualization is computed by the server, and only the resulting images are sent to the webpage. Users can select six orthogonal directions, i.e., head, feet, anterior, posterior, left and right, to view the data, or a predefined direction which is specified on the server. The information interface contains four tabs: cross filter, boolean filter, particle inlet visualization, and group comparison. Cross and boolean filters both contain histogram visualizations, where users can interact with the histograms to filter the particles and highlight regions in the 2D VesselMap. Particle inlet visualization shows the particle distribution over the inlet plane that feeds the user-specified region. Note that a single seeding plane is used as the inlet of particles. When referring to a property of particles, we use the term “particle inlet”, or simply “inlet”, to indicate the position where a particle is released in the seeding plane, instead of a specific inlet among multiple inlets. Group comparison demonstrates the differences among groups of blocks generated from the segmentation.

4.1. Histogram Visualization and Filtering

Our method visualizes statistical information of different properties, such as velocity, vorticity, WSS, particle inlet and pressure, in real time. Users can specify a region of interest where the histogram will be generated. If it is not specified, the entire volume will be taken into consideration. As shown in Figure 7 (a), the histograms are visualized as bar charts and the selection of a certain range of property values to further refine the histogram computation can be directly performed on those bar charts. This will be very helpful for domain experts,

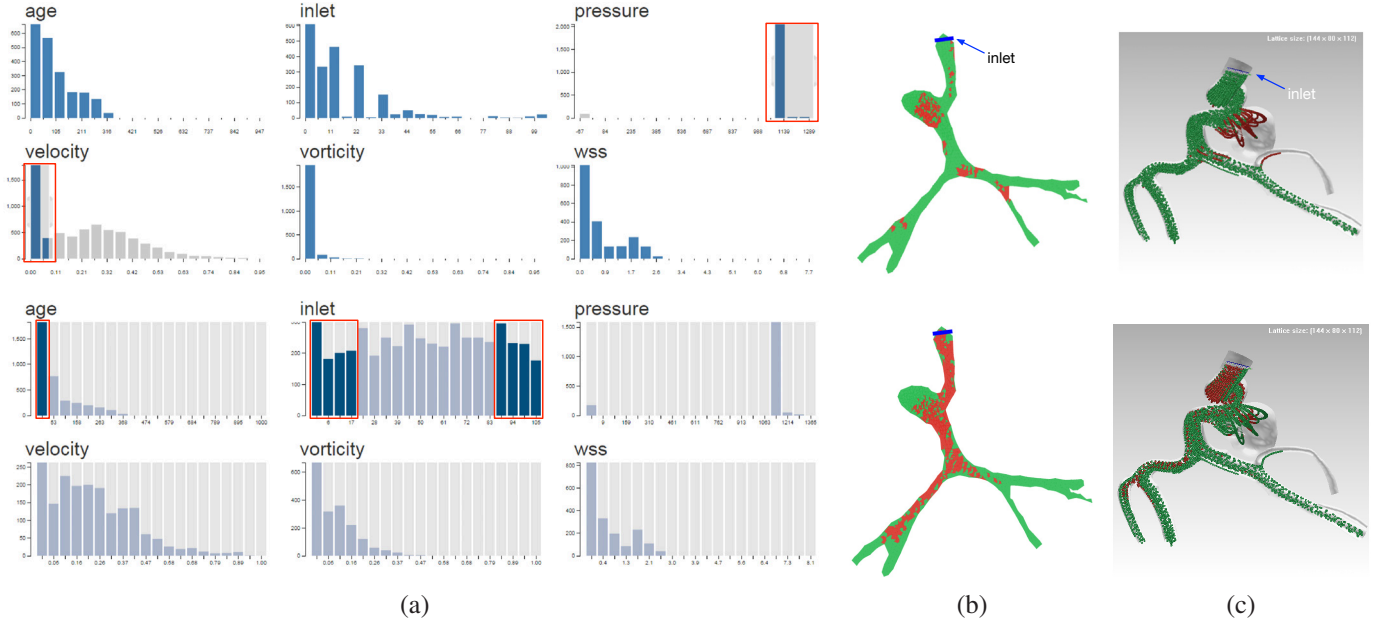


Figure 7: Cross filter and boolean filter for refining queries on histograms. The selected ranges in the histograms are highlighted in red rectangles. Column (a) shows the interface of visualizing histograms and selection of bins, where selection is performed by a cross filter in the first two rows and a boolean filter in the later two rows. Column (b) shows the highlighting results on the VesselMap, where red blocks contain particles fulfilling the selection criteria and the green ones do not. Column (c) shows the highlighting results of particles in the original 3D volume.

since finding the regions that meet certain criteria is a common task, e.g., finding the region with high pressure and low WSS. We provide two methods to specify the refining query criteria: *cross filter* and *boolean filter*.

For cross filter, users can mouse over multiple bins of a histogram, so the computation of all other histograms will be restricted to the particles falling into these selected bins. In other words, when computing other histograms, only the particles in these selected bins are counted. If multiple histograms are selected, the computation of all non-selected histograms will be restricted by these histograms; and any selected histogram will be restricted by other selected ones but not by itself. The first row of Figure 7 shows an example of a cross filter. In (a), the bins corresponding to lower velocity and higher pressure are selected using a cross filter. Histograms of other properties show the distributions of particles with low velocity and high pressure values, the velocity histogram shows the distribution of particles with high pressure values, and the pressure histogram shows the distribution of particles with low velocity. The blocks and particles are also highlighted according to the selected bins, as shown in (b) and (c), respectively. Those particles that belong to the selected bins are highlighted in red, while other particles are displayed in green.

The boolean filter behaves similarly. Nevertheless, instead of brushing across the bins, users need to click on multiple bins to specify the refining criteria. Considering the selection status of each bin as a boolean, the selected bins in the same histogram are connected by the *or* operator, while those in different histograms are connected by the *and* operator. The constraints on different histograms are the same as with the cross filter. The second row of Figure 7 demonstrates an example of a boolean filter. Eight bins of inlet are selected, which means only par-

ticles released from these eight positions will be counted. In addition, a bin of age is selected to further constrain that only the particles with small age are counted.

Note that using cross filter, users can select a portion of the range of a bin, and the corresponding value will be interpolated. But they are not able to select two disjoint ranges for query. Using boolean filter, users can select disjoint ranges. But for each selected bin, the entire range of the bin is used.

4.2. Particle Inlet Visualization

Particle inlet visualization displays the number of contributing particles in a user-specified region for each inlet position using a heat map. In our case, each small color block on the heat map corresponds to a position on the inlet plane. This plane is displayed as a blue bar in the VesselMap, as shown in Figure 7 (b). In flow visualization, the plane is shown as a transparent box with white boundaries, and the particles on the plane are colored in blue, as shown in Figure 7 (c) pointed by the label “inlet”. The seeding plane is evenly partitioned into small regions. Each small region corresponds to a color block on the heat map. The color indicates the number of particles originated from the corresponding region which are currently in the region of interest. Users can interact with the VesselMap to specify the region of interest, as shown in Figure 8. Note that some blocks in the selected region are still colored in green because these blocks do not contain any particle. Figure 8 shows two heat maps with respect to the aneurysm and a segment of a vessel branch, respectively. This information might help biomedical engineers identify the vessel branches that feed the aneurysm and the locations to release drugs that will be delivered to the desired location. For example, in the first row of Figure 8, we find that most particles in the aneurysm region are released from the upper right region of the seeding plane.




Figure 8: Contributing particle distribution at a flow inlet. (a) shows the user-selected region on the VesselMap. (b) is the heat map that indicates the number of contributing particles released from each position on the seeding plane. In the first row, the region of interest is the aneurysm. In the second row, the region of interest is a segment of a vessel branch highlighted in red.

4.3. Visualization of Particle Transportation Information

Visualization of particle transportation is important for clinical applications. More specifically, therapeutic outcomes could be better predicted if we can reliably assess how those chemical/biological agents interact with the blood stream, perhaps, in a “patient-specific” manner. Using a cerebral aneurysm as an example, users can specify a duration of time during which chemical agents or particles are released at different locations, as shown by the blue and orange line segments in Figure 9 (a). Then, particle transportation will be assessed and recorded between two time-intervals from the first particle being released to the last particle leaving the volume of interest. As shown in Figure 9 (b), time-intensity curves related to agent release can be stored and analyzed. Therefore, clinically-relevant parameters such as peak (particle) density, time to peak, and washout time can be obtained. In addition, we display the histogram of lengths of time that each particle stays in the selected region, as shown in Figure 9 (c) and (d). Our system can store multiple curves for comparison. In Figure 9 (b), a comparison between two releasing locations is provided. It is easy to imagine that those catheter-based therapies [36] would often become more effective if those injected agents would arrive in a large quantity and remain in the wounded site longer. By examining multiple releasing locations like that, it would help clinicians optimize the effects of agent delivery.

4.4. Group Comparison

In the segmentation, the vessel structure is divided into several groups of blocks guided by some property. The group com-

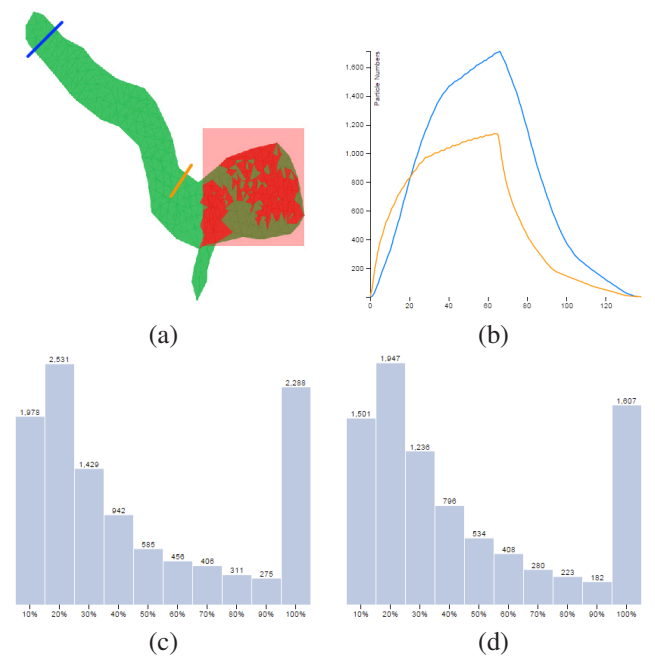


Figure 9: Visualization of particle transportation information using the VDS3 data set. (a) shows the VesselMap representation with the queried region, and two inlets marked by the blue and orange line segments. (b) shows the numbers of particles in the selected region over time. (c) and (d) show the histograms of lengths of time that the particles stay in the selected region released from the blue and orange inlets, respectively.

parison interface is designed to investigate the differences of histograms among these groups. Note that the property for histogram computation is specified by the users, and could be different from the one that guides the segmentation. In this way, the local behaviors and relationship between properties can be better studied. The differences are organized in an $n \times n$ matrix, where each cell visualizes a difference value together with a Q-Q plot indicating how the two distributions differ. All groups are sorted by decreasing numbers of particles in them, and arranged left to right and top to bottom in the matrix. Figure 10 (a) shows an example of the difference matrix interface. Users can use this interface to narrow down a region of interest. By clicking a cell in the difference matrix, the row and column containing that cell will be highlighted, as shown in (b). In addition, the corresponding groups and particles will be highlighted in the VesselMap and the particle visualization as well, as shown in (c) and (d), respectively.

5. Results and Discussion

Although our histogram interactions can be performed on unsteady data, the current segmentation is based on the average values computed at a given time step. Otherwise, segmentation may change over time and the relationships among groups can be difficult to perceive. Thus, we only used steady fields in this paper. Four data sets were used: vascular data set 1 (VDS1; one terminal aneurysm at middle cerebral artery), vascular data set 2 (VDS2; two closely-spaced lateral aneurysms at the internal carotid artery), vascular data set 3 (VDS3; one giant lateral aneurysm at the internal carotid artery), and vascular data set 4 (VDS4; one lateral aneurysm at the internal carotid

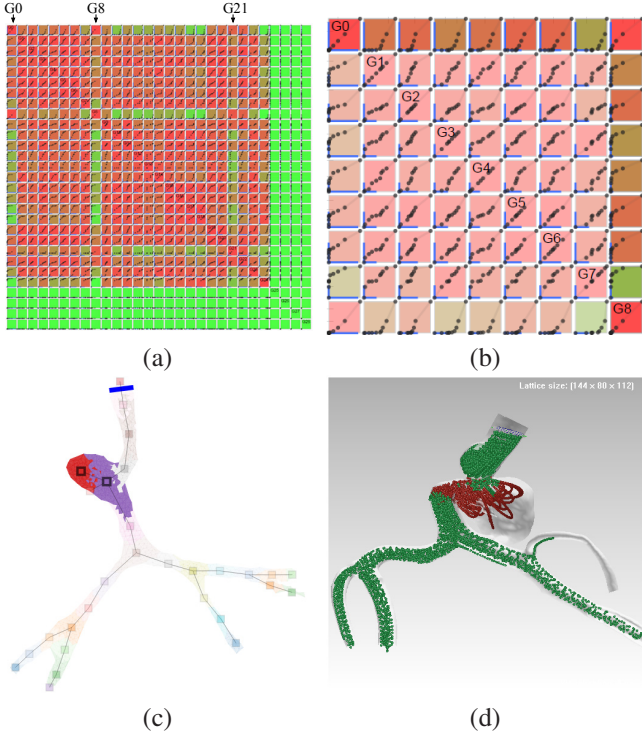


Figure 10: The relationships among velocity distributions of groups of blocks segmented according to their distances to the inlet. (a) shows the difference matrix. (b) is a zoom-in view focusing on the groups from G_0 to G_8 . Cell (G_0, G_8) is clicked, and row G_0 and column G_8 are highlighted. (c) and (d) are the segmentation result and particle visualization with groups G_0 and G_8 highlighted.

artery). Their dimensions were $108 \times 60 \times 84$, $45 \times 63 \times 78$, $72 \times 60 \times 72$, and $45 \times 92 \times 142$ respectively. The timing results were collected on a PC with an Intel Core i7-3820 CPU running at 3.6 GHz, 16 GB main memory, and an nVidia Geforce 670 graphics card with 2 GB graphics memory. Even leveraging the power of the GPU, the generation of the VesselMap dominates the timing cost, as it requires multiple linear systems to be solved. For the VDS1 containing 1095 blocks, the corresponding VesselMap took 55.8 seconds to compute. The time needed to calculate the VesselMaps of VDS2 (1482 blocks), VDS3 (863 blocks), and VDS4 (1278 blocks) was 137.6, 35.2 and 128.7 seconds, respectively. This performance is acceptable since the VesselMap only needs to be computed once for each data set. The histograms can be updated using the GPU in real-time.

5.1. Case Study: VDS1

We submitted a query to obtained particles with both large pressure and small velocity values. The results from the proposed system are shown in the first row of Figure 7. The VesselMap and the particle visualization both show that the queried particles mostly resided in the aneurysm sac and those regions where vessel branches bifurcate. We observed from the inlet histogram that most of these particles were released from the positions corresponding to four bins. In the first row of Figure 8, we also investigated the contribution of each position on the seeding plane to the aneurysm sac. Note that some blocks

in the selected region are still in green because they do not contain any particles (see Figure 7). Particle inlet visualization indicates that most particles were released in a small region, which is consistent with the inlet histogram. In the second row of Figure 8, a segment of a vessel branch was selected, and the contribution of each position in the inlet plane was somewhat uniform. In the second row of Figure 7, a bin with small age values and eight bins of inlet were selected. We observed that the inlet positions corresponding to the particles with small age values were mostly evenly distributed (see Figure 8). It is also obvious that the velocities of these particles with small age values are more uniformly distributed. Furthermore, among those particles with small age values, there were more particles with higher vorticity values, as compared with the histograms in the first row of Figure 7. From the VesselMap, we observed that most particles associated with small age were in one vessel branch. This indicates that particles in the above-mentioned vessel branch spent less time within the aneurysm sac.

In Figure 4 (b), the segmentation result of the VDS1 guided by the vorticity showed that more groups were generated around the neck of the aneurysm and around an inlet of a vessel branch. This indicates that the average vorticity values of blocks in these regions varied more dramatically.

The segmentation guided by the distances to the inlet better preserved the shape of the vessel structure, as shown in Figure 4 (a). We used the difference matrix to analyze the differences of velocity distributions among different groups. Note that the groups that do not contain any particle will be considered as entirely different with other groups, and the corresponding cells in the difference matrix will all be colored in green. In Figure 10 (a), we observed that most green cells correspond to groups G_0 , G_8 and G_{21} , except the empty groups. This indicates that the velocity in these three groups might be different from other groups. Clicking at cell (G_0, G_8) highlights row G_0 and column G_8 , as shown in Figure 10 (b). In cells (G_0, G_0) and (G_8, G_8) , points in the Q-Q plot were denser in the smaller value range. Moreover, cells in row G_0 and G_8 showed that the Q-Q plot lines mostly aligned horizontally when comparing these two groups with other groups. These facts imply that particle velocities in these two groups were smaller than in the other groups. By interacting with the difference matrix we found that G_0 and G_8 belonged to the aneurysm region.

5.2. Case Study: VDS2

The vessel structure of the VDS2 can be perceived in Figure 11 (b) and (c). The first row of Figure 11 demonstrates the query of particles with large age. The inlet histogram shows that all particles with large age actually correspond to one bin, which means that they were released at a very small region on the seeding plane. From the VesselMap and the particle visualization, we observed that these particles went through the aneurysm dome region, and their paths were close to the centerline of the main vessel branch. In addition, as expected, these particles mostly have small velocity values. In the second row of Figure 11, we select the aneurysm dome region for query. We found that most particles in that region had moderate age values. This confirms that the particles with large age were

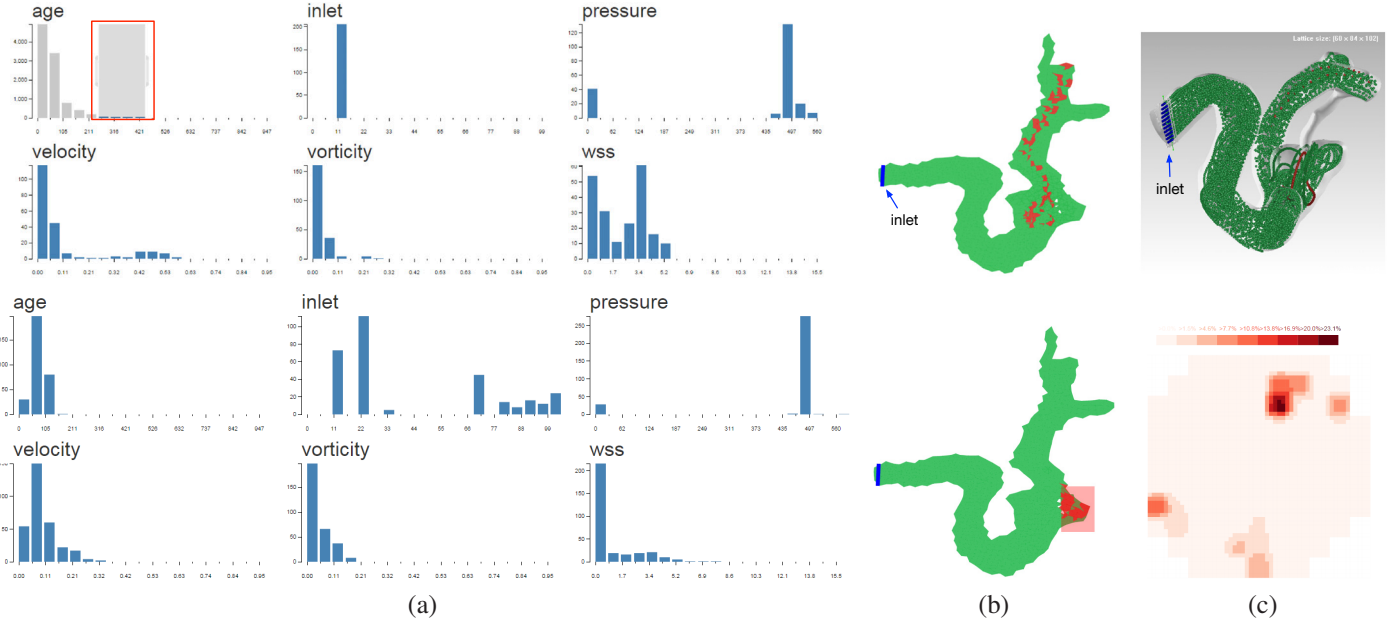


Figure 11: Visualization result using the VDS2 data set. The first row shows the query of particles with large age. The second row shows the query regarding the aneurysm. (a) shows histogram visualization. (b) shows the VesselMap representation. (c) shows the visualization of particles with large age in the first row, and the inlet heat map associated with the aneurysm in the second row.

mostly those that had already left the aneurysm dome. The inlet histogram shows that the particles filling the aneurysm sac actually came from multiple inlet positions. This is also confirmed by results presented in Figure 11 (c). However, comparing the two inlet histograms in Figure 11 (a), we observed that the particles that stayed within the aneurysm sac for a long time were mostly released from a specific location. From the velocity histograms, we found that the aneurysm dome even contained more particles with slightly higher velocity values than the particles with large age. This observation is consistent with the fact that intra-aneurysmal flow contains both high velocity jet(s) and re-circulation zone(s) with low velocity. The pressure and especially the WSS values mostly concentrate in a few histogram bins, indicating more uniform distributions.

In the first row of Figure 12, we show the segmentation result guided by different properties. In (a), using the distances to the inlet, we obtained the skeleton of the vessel structure. In (b), we can see four large groups segmented by age, which are arranged similarly as the distances to the inlet. But there are also some small groups, in which the particles are probably stuck and stay for a long time (i.e., long particle residence time). In (c), we can see that three large groups align along the main vessel branch. This indicates that there might be several major paths with different velocities along the vessel. Note that this finding is consistent with our discussion of the histograms in Figure 11. In (d), we notice that the average vorticity value varies in most of the vessel structure, but they are somewhat stable in the aneurysm dome and its immediate neighborhood. In (e), we find that the WSS varies in the region near the aneurysm sac. However, it is very stable inside the aneurysm, since the aneurysm forms a single group. This can also be perceived in the histogram shown in Figure 11.

We then investigate the relationship among groups, as shown

in the second row of Figure 12. For the groups segmented by the distance to the inlet, we find that the age distribution of groups G_{11} and G_{12} are different from other groups, as shown in (f). The age for most groups cover a large range, with their maximum values much larger than other values. But the age in G_{11} is distributed somewhat uniformly in a smaller range. Note that G_{11} corresponds to the top of the aneurysm region. This indicates that the particles in this region follow similar paths before, and have similar age. (g) shows that the velocity in G_2 and G_{12} differs from other groups. We observe that G_2 contains more particles with small velocity, which is common for an aneurysm. (h) shows that the aneurysm region corresponding to G_2 and G_{11} contains more particles with small vorticity. For the groups segmented by age, we find that the groups form basically two clusters, considering the age distributions or the velocity distributions, as shown in (i) and (j), respectively. Groups G_0 to G_4 form one cluster, and the other groups form the other one. Two groups from the same cluster have similar age and velocity distributions, but two groups from different clusters are less similar. Note that the first cluster consists of groups with more particles, containing the four major groups in (b). It is likely that the isolated small groups have different average age values from their neighbors due to velocity differences.

5.3. Case Study: VDS3

We compared the particle transport information using the VDS3. The same quantity of particles was released for the same duration of time at two different releasing locations, respectively, as shown by the blue and orange lines in Figure 9 (a). The numbers of particles over time are displayed in (b), where a curve corresponds to the inlet of the same color. Note that, in (b), segments where the numbers of particles are zero

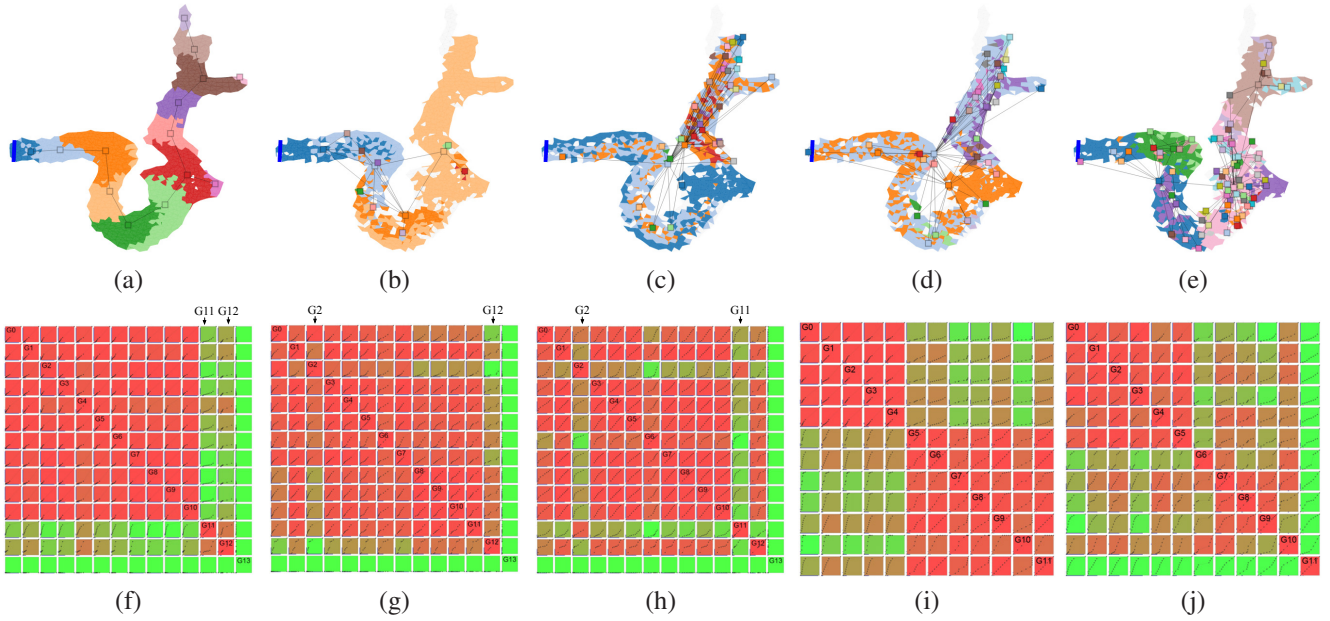


Figure 12: The segmentation result and group comparison using the VDS2 data set. First row: segmentation result guided by (a) distances to the inlet, (b) age, (c) velocity, (d) vorticity, and (e) WSS. Second row: group comparison of (f) age, (g) velocity, and (h) vorticity among groups segmented by distance to the inlet; and of (i) age and (j) velocity among groups segmented by age.

in the time-intensity curves were removed, so that those two displayed curves were aligned by the time when their first respective particles entered the selected region (see the red box in (a)). We observed, from the curves in (b), that although the orange plane was closer to the selected aneurysm sac region, the peak density was actually lower, and the washout time was shorter. Putting this into a practical perspective, we may tend to choose a location that is closer to the aneurysm sac for agent delivery. However, our result suggests that this intuition is actually incorrect.

In Figure 9 (c) and (d), the histograms corresponding to the blue and orange releasing planes are shown. These two histograms have similar shapes, indicating similar percentages of particles staying for a certain amount of time. However, the absolute numbers of particles in (d) were actually smaller, which is consistent to the finding shown in (b). From the perspective of flow physics, a higher percentage of particles released from the orange location/plane is directly diverted away from the aneurysm, i.e., directly moving toward downstream without entering the aneurysm. In short, the availability of such a tool may enable biomedical engineers to improve their releasing locations for catheter-based drug delivery in a “patient-specific” manner [36].

5.4. Case Study: VDS4

Unlike the VDS1, VDS2 and VDS3, each of which has a single CFD simulation, VDS4 comes with two CFD simulations using exactly the same vessel anatomy. Two CFD simulations represent two different flow-diverting stent configurations: 20% (Stent 1) and 30% (Stent 2) metal coverage values. This case study helps us to evaluate the performance of those two different stents in a “patient-specific” manner. In general, a stent with a high metal coverage value (i.e., lower

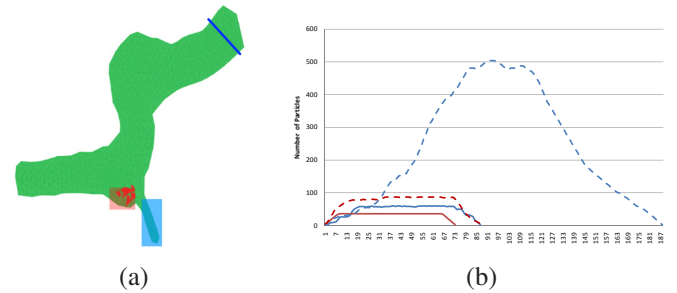


Figure 13: Visualization of particle transportation information using the VDS4 data set. (a) shows the VesselMap representation with two queried regions, indicated by the red and blue rectangles, respectively. (b) shows the numbers of particles in the two selected regions with two different vector fields over time. The red and blue lines correspond to the red and blue queried regions, respectively. And the solid and dashed lines correspond to Stents 1 and 2, respectively.

porosity) diverts more flow entering the aneurysm sac (causing stable thrombosis) but also likely blocks blood flow passing through adjacent small arteries (e.g., ophthalmic artery in this case). Now, by measuring the number of particles in the aneurysm and the small vessel branch near the aneurysm, we can evaluate the flow diverting effect while assessing how much blood flow is spared for the ophthalmic artery. In Figure 13 (a), the queried aneurysm sac region and the small ophthalmic artery are highlighted using red and blue semi-transparent rectangles, respectively. In Figure 13 (b), the number of particles over time are displayed for the red and blue queried regions using the two different stents. All particles were released from the same seeding plane for the same duration of time. It is easy to see that a significant reduction of flow entering the aneurysm sac (i.e., low particle numbers and longer wash-out time as indicated by Figure 13) has been achieved by increasing the metal coverage from 20% (Stent 1) to 30% (Stent 2). The increased




Figure 14: The VesselMap generated with different settings: (a) using only the geodesic distances for VDS3 data set and (b) using twice the geodesic distances for VDS4 data set.

metal coverage in Stent 2 still allows the blood passage into the ophthalmic artery.

5.5. Discussion

We would like to discuss three issues of the VesselMap representation: preserving orientation, preserving branching order, and the selection of target distances.

Preserving orientation. Preserving the orientation in the 2D VesselMap is almost impossible, as the 3D view can be rotated. For example, in Figure 1 (a), the left and right branches in the VesselMap correspond to the branches at the same side of the 3D view, respectively. However, if the 3D structure is viewed in an opposite direction, the orientation could be reversed. Actually, the orientation depends on the initial positions of blocks. If a standard 3D view can be given, it is likely to generate the VesselMap with similar orientation, by using the 2D projections of the blocks under the standard view as the initial positions. But, if the 3D view is rotated, the orientation could still be misleading. Therefore, unless the 2D VesselMap will be adjusted according to the 3D view, it is not possible to preserve the orientation.

Preserving branching order. Our VesselMap representation does not guarantee to preserve the order of vessel branches, although the relative branching order may be preserved due to the distance preservation of different branches. For example, if there exists one branch b_1 between two branches b_2 and b_3 , the distances between b_1 and b_2 , and b_1 and b_3 should be small, and the distance between b_2 and b_3 should be larger. To preserve the distance in the 2D VesselMap, the mapped branch b_1 should still be in between of b_2 and b_3 and the branching order is preserved. However, this may not always be the case. For example, in Figure 1 (a), the two branches on the right are symmetric at the bifurcation point, and their order cannot be captured by the distance between them. Also, similar to the orientation, the branching order may also depend on the viewing angle.

Selection of target distances. Our current approach uses Euclidean distances for nearby blocks and geodesic distances for the others. We argue that using both the two distances is necessary to obtain the desired visualization. It is clear that the different branches cannot be stretched out using only the Euclidean distance. In addition, our experiment shows that using only the geodesic distance is also not enough, as shown in Figure 14 (a). We can observe that the local shape is not well preserved, and two bumpy regions appear on the vessel. In

contrast, in Figure 9, the VesselMap representation is smoother along the boundary and preserves the curvy turns of the vessel. This is probably due to our approximation of geodesic distance is not accurate when the Euclidean distance is small. Therefore, for a small local neighborhood, it is still better to use only the Euclidean distance; and for non-neighboring blocks, using geodesic distance to stretch out the vessel branches seems to be reasonable.

However, the current approach may be improved. In Figure 14 (b), we find that the aneurysm region highlighted in the blue rectangle becomes larger and more standing-out when the target distances for non-neighboring blocks are twice the geodesic distances. This is likely due to the increase of the force to stretch out the vessel branches (including aneurysms). But the artifact is also obvious. It seems that the balance between the Euclidean distance and the geodesic distance is broken. The boundary of the vessel structure is not smooth. The width of the main vessel branch varies in an unnatural way, and serious distortion appears on the small branch, which seems to be twisting. These indicate that the improvement demands further research effort.

6. Evaluation

6.1. Empirical Expert Evaluation

Expert evaluations were done by one of the co-authors Dr. Jingfeng Jiang, who is a specialist in image-based biomechanical analysis. Dr. Jiang research interests include predicting and analyzing bio-flows for cardiovascular diseases. He has approximately 10 years experience in cardiovascular flow analysis. This evaluation aimed at providing an initial assessment on the usefulness of the proposed system instead of quantitative results. The expert was asked to perform several tasks, including querying regions with histograms, finding seeding locations that fed the aneurysm, and selecting specific properties for segmentation. These tasks were only designed to guide the expert through the workflow of our approach. During the evaluation, Dr. Jiang was informed to freely interact with our interface to explore the VDS1 according to his own needs.

After learning the features of our interface, Dr. Jiang explored the VDS1 using our interface and provided the following comments: “Overall, this is a good application and could be particularly useful to collaborative research in transnational biomedical device/technology design involving an engineer and physician partnership. With this tool, engineers and physicians can interactively explore and subsequently communicate experimental and/or pre-clinical data over the internet. In addition, this application does not significantly differ from typical web applications. Thus, for ordinary users who are comfortable with web browsing, it should be easy to learn. I definitely see its potential use for flow visualization in a clinical setting given further developments.”

“In terms of effectiveness, the VesselMap allows users to quickly select a region of interest in 2D, because the flattening process will “evenly” map the 3D volume onto a 2D plane. While in original 3D volumes, selecting a location inside the

volume is difficult because viewing of the desired location may be blocked by objects in the background or foreground. A technical concern is that the flattening process of the VesselMap may cause unrelated 3D points to look like spatial neighbors. In terms of building connection between the 2D VesselMap and the 3D vessel structure, users can brush across several branches on the VesselMap, and observe the connection from highlighted region in the corresponding 3D particle visualization. However, since users may not be immediately aware of this, they should be forced to perform this task in their learning stage. In addition, some visual landmarks can be placed on both the VesselMap and the particle visualization, so that the connection can be perceived more easily.”

“Cross and Boolean filters are very useful. Currently, everyone looks at flow data based on his or her own experience. Thus, there is a good chance to overlook some important features or correlations. In this application, once the queries are set up, users can use certain combinations to quickly identify areas of interest that meet all criteria. This helps to examine the flow characteristics within and around a cerebral aneurysm in a more comprehensive way.”

“The segmentation based on a criterion (e.g., distance, velocity, etc.) will enable the user to quickly grasp spatial distributions of the flow characteristics. Based on that, it should help the user select some important regions of interest for further analyses. Some work may be needed to better organize or cluster the segmented regions. For instance, combined with streamline visualization, segmented regions following a flow pathway will provide additional insight from the perspective of bio-transportation. The difference matrix provides the overall correlation between different regions as well as a Q-Q plot. With the color-coded matrix, if we use age for segmentation and identify groups with different velocity distributions, it is possible to guide the user to flow stagnation regions. Thus, this segmentation scheme could potentially help users identify regions which contain interesting flow features.”

6.2. User Study

We recruited six unpaid graduate students for a quantitative study: four majoring in biomedical engineering, and the other two majoring in computer science. Three of those four biomedical engineering students have (1-2 years) experience of using commercial or open-source software (e.g., EnSight and Paraview) for bio-flow visualization. Therefore, these three students were considered to be the advanced users (AU). The remaining three students are either knowledgeable in vascular structures or visualization, but they never used visualization tools for vascular flow before this study. They were considered as beginning users (BU). The user study was conducted in a lab using the same PC with only one participant presented at one time. The PC has a 27-inch monitor with 1920 × 800 resolution, where the user study interface occupies an area of 1800 × 600 pixels. Each study took about an hour to complete. The study consists of three stages: an exploration stage, a test stage, and a rating stage. The exploration stage allowed users to get familiar with our interface and the vascular structure in the VDS1, which was used in the test stage as well. In the test stage,

Table 1: Averages and standard deviations of task completion time and accuracies for different user groups. “AU” and “BU” were advanced users and beginning users, respectively. Timing is measured in seconds, and accuracy is presented in percentage.

		timing		accuracy (%)	
		VesselMap	punch	VesselMap	punch
AU	average	26.0	64.8	89.2	81.4
	std. dev.	20.2	73.5	7.8	8.8
BU	average	31.3	72.8	91.4	77.1
	std. dev.	19.6	36.6	8.0	11.2
All	average	28.7	68.8	90.3	79.2
	std. dev.	20.0	57.9	7.9	10.3

users were asked to reproduce previous selection results using both the VesselMap and the traditional punching selection. The punching operation is commonly used in clinical workstations by clinicians and biomedical engineers. It selects a polygonal region in the screen space so that all voxels projected to that region are selected. Users may rotate the scene and select again, so that only the intersection of the two selections remains, and the undesired voxels are excluded. This operation will be repeated several times so that users can gradually refine the selection. Finally, in the rating stage, users were asked to provide ratings to several statements.

The major goal of this user study is to compare our VesselMap and the punching selection in terms of the effectiveness to select a specific region. Two factors were considered: efficiency and reproducibility, measured by the time used and accuracy to reselect a set of blocks given one or two screenshots of a previous selection result. Users were informed that the accuracy was more important than the completion time. A program was designed specifically for this user study. For selection using the VesselMap, one screenshot of the previous selection would be given; and for the punching selection, two screenshots would be given: one in a manually selected direction that best demonstrated the region to be selected, and the other in the anterior direction as a standard view. For fair comparisons, the selection for both methods were performed at the block level.

We selected five regions and generated the screenshots for both the VesselMap and the punching selection. For each task, users were presented the screenshots of one selected region. Users would need to reselect that region to finish the task. Each user was required to perform three rounds of selection tasks. In each round, they would need to perform five different tasks using the VesselMap followed by five using the punching selection. Therefore, each user performed $2 \times 5 \times 3$ selection tasks in total. Note that the tasks were the same for each round, but they appeared in different orders.

Timing and Accuracy. Since the completion time and accuracy might be related, we applied MANOVA (multivariate-ANOVA) to investigate the overall difference, and ANOVA to study the difference of each individual variable. For MANOVA, Wilk’s Lambda test was used and both the completion time and

Table 2: Averages and standard deviations of ratings for the four statements.

	Q1	Q2	Q3	Q4
average	4.33	2.67	4.67	4.17
std. dev.	0.51	1.03	0.51	0.98

accuracy were considered at the same time. We used the significance level of 0.05 for all tests, and our findings were as follows.

MANOVA found significant differences of timing and accuracy for the VesselMap and the punching selection with the p -value of 2×10^{-16} . The differences of timing and accuracy using these two methods were significant using ANOVA as well, with the p -values of 2×10^{-9} and 1×10^{-13} , respectively. In Table 1, we observe that using the VesselMap, the completion time is much shorter and accuracies are higher for both the AU and BU groups. This indicates that our VesselMap improves the efficiency and reproducibility over the punching selection.

ANOVA also suggested a significant difference of completion time for the three rounds with the p -value of 0.040. The overall averages of completion time for the three rounds were 59.3, 43.5 and 43.4 seconds, respectively. The average completion time for both the VesselMap (32.6, 27.4 and 26.1 seconds) and the punching selection (86.1, 59.7 and 60.6 seconds) show the same trend. We found that the differences among the completion time during the three rounds were smaller using the VesselMap, which might indicate that it is easier to repeat a task with the VesselMap. However, no significant difference of accuracy was found for the three rounds. The overall accuracies for the three rounds were actually very closed, being 85.1%, 84.7% and 84.5%, respectively. Users were able to perform the tasks with high accuracy, but it just required more time for the first time.

We further investigated the impact of expertise using these two methods. The averages and standard errors are shown in Table 1. We did not find any significant difference of timing using the VesselMap between AU (averaged at 26.0 seconds) and BU (averaged at 31.3 seconds) groups with the p -value of 0.210. For accuracy, no significant difference was found as well (with the p -value of 0.193), and the accuracies for AU (89.2%) and BU (91.4%) groups were very close. For the punching selection, although the difference of average completion time was larger (64.9 seconds for AU and 72.8 seconds for BU), ANOVA did not suggest any significant difference (the p -value is 0.508). This was likely due to the larger variation of the completion time, as shown in Table 1. In terms of accuracy, ANOVA found a significant difference between AU (81.3%) and BU (77.1%) groups with the p -value of 0.050. This indicated that the punching selection required more experience from users to perform accurately.

Overall, the results show that our VesselMap is better in terms of efficiency and reproducibility. Furthermore, it is likely that the VesselMap has lower requirements on user expertise and learning effort.

Ratings. We observed during the user study that the diffi-

culty of using the punching selection resided in selecting the viewing direction shown in the screenshots, which might be solved by our automatic viewpoint selection. For the rating questions, we asked users to freely interact with the interface to select viewpoints using the VesselMap and rate the following statements:

- Q1. The automatically selected viewpoint is close to the ideal viewpoint.
- Q2. It is easy to repeat a viewpoint by translation and rotation in 3D.
- Q3. It is easy to repeat a viewpoint using the VesselMap.
- Q4. It is easy to adjust the automatically selected viewpoint to observe the region of interest if it is not ideal.

The rating is a five-point scale, where “5” means strongly agree and “1” means strongly disagree. The rating results are shown in Table 2. The averages of ratings for Q1, Q3 and Q4 were all higher than 4, with the highest being 4.67 for Q3. The high rating for Q1 and Q4 indicated our automatic viewpoint selection using the VesselMap produced satisfactory results. The rating for Q3 indicated that users agreed that repeating a viewpoint using the VesselMap was easy. The rating for Q2 (2.67) was low, meaning users did not agree that repeatedly selecting the same viewpoint is easy by manipulating the 3D view. We then applied the non-parametric Kruskal-Wallis (KW) test to study the rating difference between Q2 and Q3. Note that, unlike in the analysis of the timing and accuracy, we did not use t-test for the rating, since the sample size was small and the rating did not follow a normal distribution. The KW-test suggested a significant difference with the p -value of 0.005. This indicated that repeating a viewpoint using the VesselMap was easier.

7. Conclusions and Future Work

We present VesselMap, a web interface to explore multivariate vascular data sets. The web interface provides great compatibility for users working on different platforms. Interaction with histograms allows users to quickly identify regions meeting certain criteria. The 2D VesselMap representation provides not only an overview but also an efficient method to select regions of interest. In addition, local behaviors of different properties can be discovered effectively using the segmentation and group comparison features.

There are challenges to be solved in the future as well. First, although the discontinuity produced by the flattening process from 3D volumes to 2D maps is unavoidable, we might post-process the user selection and enforce the selected blocks to be spatial neighbors in the 3D volume. The same procedure can be applied when highlighting a region, so that the blocks look continuous on the VesselMap. Second, the performance of generating the VesselMap can be improved. The constraints in the linear system are enforced on every two blocks. We may use the centers of groups segmented by distances to the inlet as anchor points, and only enforce the constraints on every two anchor points and neighboring blocks. In this way, the size of the

linear system to be solved can be greatly reduced. Third, the automatic viewpoint selection can be significantly improved. The current approach uses a weighted sum of the three principal components as the viewing direction, which only considers the spatial distribution of selected blocks. A more sophisticated solution may evaluate the information revealed under a certain viewpoint, considering both occlusions among blocks and flow patterns. Fourth, we may allow the users to adjust the layout of the VesselMap. Users can simply drag anchor points to desired locations, and our application will update the positions of other blocks accordingly. Fifth, the segmentation scheme can be extended to handle unsteady data sets. For example, the average value in a block over a time span can be used to guide the segmentation. Finally, many practical features can be further developed based on the VesselMap. For example, after placing a seeding plane using the VesselMap, users may further rotate the plane for a desired orientation. Users may place labels on the VesselMap, and our system can generate a report highlighting the labeled regions in the 3D volume, together with comments. We may also allow users to measure the volumetric size of the selected region, etc.

Acknowledgments

This research was supported in part by the U.S. National Science Foundation through grants IIS-1456763, IIS-1455886, DUE-1140512, and DUE-1245310.

References

- [1] Brass LM, Lichtman JH, Wang Y, Gurwitz JH, Radford MJ, Krumholz HM. Intracranial hemorrhage associated with thrombolytic therapy for elderly patients with acute myocardial infarction results from the cooperative cardiovascular project. *Stroke* 2000;31(8):1802–11.
- [2] Molyneux AJ, Cekirge S, Saatci I, Gál G. Cerebral aneurysm multicenter european onyx (CAMEO) trial: results of a prospective observational study in 20 european centers. *American Journal of Neuroradiology* 2004;25(1):39–51.
- [3] Guan YS, Sun L, Zhou XP, Li X, Fei ZJ, Zheng XH, et al. Polyvinyl alcohol and gelatin sponge particle embolization of splenic artery pseudoaneurysm complicating chronic alcoholic pancreatitis. *World Journal of Gastroenterology* 2005;11(17):2684–6.
- [4] van Rooij WJ, Sluzewski M, Beute GN. Brain AVM embolization with onyx. *American Journal of Neuroradiology* 2007;28(1):172–7.
- [5] Hahn HK, Preim B, Selle D, Peitgen HO. Visualization and interaction techniques for the exploration of vascular structures. In: *Proceedings of IEEE Visualization Conference*. 2001, p. 395–402.
- [6] Höhne KH, Pflessner B, Pommert A, Riemer M, Schubert R, Schiemann T, et al. A realistic model of the inner organs from the visible human data. In: *Proceedings of International Conference on Medical Image Computing and Computer-Assisted Intervention*. 2000, p. 776–85.
- [7] Felkel P, Wegenkittl R, Bühler K. Surface models of tube trees. In: *Proceedings of Computer Graphics International*. 2004, p. 70–7.
- [8] Oeltze S, Preim B. Visualization of vasculature with convolution surfaces: Method, validation and evaluation. *IEEE Transactions on Medical Imaging* 2005;24(4):540–8.
- [9] Lorensen WE, Cline HE. Marching cubes: A high resolution 3D surface construction algorithm. In: *Proceedings of ACM SIGGRAPH Conference*. 1987, p. 163–9.
- [10] Låthé G, Lindholm S, Lenz R, Persson A, Borga M. Automatic tuning of spatially varying transfer functions for blood vessel visualization. *IEEE Transactions on Visualization and Computer Graphics* 2012;18(12):2345–54.
- [11] Mistelbauer G, Varchola A, Bouzari H, Starinsky J, Köchl A, Scherthner R, et al. Centerline reformations of complex vascular structures. In: *Proceedings of IEEE Pacific Visualization Symposium*. 2012, p. 233–40.
- [12] Schumann C, Oeltze S, Bade R, Preim B, Peitgen HO. Model-free surface visualization of vascular trees. In: *Proceedings of Joint Eurographics - IEEE VGTC Symposium on Visualization*. 2007, p. 283–90.
- [13] Kanitsar A, Fleischmann D, Wegenkittl R, Gröller ME. Diagnostic relevant visualization of vascular structures. In: *Bonneau GP, Ertl T, Nielson G, editors. Scientific Visualization: The Visual Extraction of Knowledge from Data*. 2006, p. 207–28.
- [14] Kretschmer J, Preim B, Stamminger M. Bilateral depth filtering for enhanced vessel reformation. In: *Proceedings of Eurographics Workshop on Visual Computing for Biology and Medicine*. 2014, p. 117–26.
- [15] Mistelbauer G, Morar A, Varchola A, Scherthner R, Baclija I, Köchl A, et al. Vessel visualization using curvilinear feature aggregation. *Computer Graphics Forum* 2013;32(3):231–40.
- [16] Borkin M, Gajos K, Peters A, Mitsouras D, Melchionna S, Rybicki F, et al. Evaluation of artery visualizations for heart disease diagnosis. *IEEE Transactions on Visualization and Computer Graphics* 2011;17(12):2479–88.
- [17] Zhu L, Haker S, Tannenbaum A. Flattening maps for the visualization of multibranched vessels. *IEEE Transactions on Medical Imaging* 2005;24(2):191–8.
- [18] Marino J, Kaufman A. Planar visualization of treelike structures. *IEEE Transactions on Visualization and Computer Graphics* 2016;22(1):906–15.
- [19] Neugebauer M, Gasteiger R, Beuing O, Diehl V, Skalej M, Preim B. Map displays for the analysis of scalar data on cerebral aneurysm surfaces. *Computer Graphics Forum* 2009;28(3):895–902.
- [20] Won JH, Jeon Y, Rosenberg J, Yoon S, Rubin G, Napel S. Uncluttered single-image visualization of vascular structures using gpu and integer programming. *IEEE Transactions on Visualization and Computer Graphics* 2013;19(1):81–9.
- [21] 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* 2013;8(2):279–89.
- [22] van Pelt R, Bescós JO, Breeuwer M, Clough RE, Gröller ME, ter Haar Romeny B, et al. Exploration of 4D MRI blood flow using stylistic visualization. *IEEE Transactions on Visualization and Computer Graphics* 2010;16(6):1339–47.
- [23] Köhler B, Gasteiger R, Preim U, Theisel H, Gutberlet M, Preim B. Semi-automatic vortex extraction in 4D PC-MRI cardiac blood flow data using line predicates. *IEEE Transactions on Visualization and Computer Graphics* 2013;19(12):2773–82.
- [24] van Pelt R, Bescós JO, Breeuwer M, Clough RE, Gröller ME, ter Haar Romeny B, et al. Interactive virtual probing of 4D MRI blood-flow. *IEEE Transactions on Visualization and Computer Graphics* 2011;17(12):2153–62.
- [25] Born S, Markl M, Gutberlet M, Scheuermann G. Illustrative visualization of cardiac and aortic blood flow from 4D MRI data. In: *Proceedings of IEEE Pacific Visualization Symposium*. 2013, p. 129–36.
- [26] Oeltze S, Lehmann DJ, Kuhn A, Janiga G, Theisel H, Preim B. Blood flow clustering and applications in virtual stenting of intracranial aneurysms. *IEEE Transactions on Visualization and Computer Graphics* 2014;20(5):686–701.
- [27] Angelelli P, Hauser H. Straightening tubular flow for side-by-side visualization. *IEEE Transactions on Visualization and Computer Graphics* 2011;17(12):2063–70.
- [28] Straka M, Köchl A, Červeňanský M, Šrámek M, Fleischmann D, La Cruz A, et al. The VesselGlyph: Focus & context visualization in CT-angiography. In: *Proceedings of IEEE Visualization Conference*. 2004, p. 385–92.
- [29] Gasteiger R, Neugebauer M, Beuing O, Preim B. The FLOWLENS: A focus-and-context visualization approach for exploration of blood flow in cerebral aneurysms. *IEEE Transactions on Visualization and Computer Graphics* 2011;17(12):2183–92.
- [30] Lawonn K, Gasteiger R, Preim B. Adaptive surface visualization of vessels with animated blood flow. *Computer Graphics Forum* 2014;33(8):16–27.

- [31] Lawonn K, Glaßer S, Vilanova A, Preim B, Isenberg T. Occlusion-free blood flow animation with wall thickness visualization. *IEEE Transactions on Visualization and Computer Graphics* 2015;22(1):728–37.
- [32] Kanitsar A, Fleischmann D, Wegenkittl R, Felkel P, Gröller ME. CPR - curved planar reformation. In: *Proceedings of IEEE Visualization Conference*. 2002, p. 37–44.
- [33] Singh G, Memoli F, Carlsson G. Topological methods for the analysis of high dimensional data sets and 3D object recognition. In: *Proceedings of Eurographics Symposium on Point-Based Graphics*. 2007, p. 91–100.
- [34] Hatcher A. *Algebraic Topology*. Cambridge University Press; 2002.
- [35] Chambers JM, Cleveland WS, Tukey PA, Kleiner B. *Graphical Methods for Data Analysis*. Duxbury Press; 1983.
- [36] Kavanagh CA, Rochev YA, Gallagher WM, Dawson KA, Keenan AK. Local drug delivery in restenosis injury: thermoresponsive co-polymers as potential drug delivery systems. *Pharmacology & therapeutics* 2004;102(1):1–15.