

Visualization of 4D Blood-Flow Fields by Spatiotemporal Hierarchical Clustering

R. F. P. van Pelt¹, S. S. A. M. Jacobs¹, B. M. ter Haar Romeny¹, and A. Vilanova¹

¹Eindhoven University of Technology, Eindhoven

Abstract

Advancements in the acquisition and modeling of flow fields result in unsteady volumetric flow fields of unprecedented quality. An important example is found in the analysis of unsteady blood-flow data. Preclinical research strives for a better understanding of correlations between the hemodynamics and the progression of cardiovascular diseases. Modern-day computer models and MRI acquisition provide time-resolved volumetric blood-flow velocity fields. Unfortunately, these fields often remain unexplored, as high-dimensional data are difficult to conceive. We present a spatiotemporal, i.e., four-dimensional, hierarchical clustering, yielding a sparse representation of the velocity data. The clustering results underpin an illustrative visualization approach, facilitating visual analysis. The hierarchy allows an intuitive level-of-detail selection, largely retaining important flow patterns. The clustering employs dissimilarity measures to construct the hierarchy. We have adapted two existing measures for steady vector fields for use in the spacetime domain. Because of the inherent computational complexity of the multidimensional clustering, we introduce a coarse hierarchical clustering approach, which closely approximates the full hierarchy generation, and considerably improves the performance. The resulting clusters are visualized by representative patharrows, in combination with an illustrative anatomical context. We present various seeding approaches and visualization styles, providing sparse overviews of the unsteady behavior of volumetric flow fields.

Categories and Subject Descriptors (according to ACM CCS): I.5.3 [Pattern Recognition]: Clustering—Algorithms & Similarity measures, I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation & Viewing algorithms, J.3 [Computer Applications]: Life and Medical Sciences—Medical Information Systems

1. Introduction

Analysis of flow phenomena plays an important role in many research areas. Over the past few decades, advancements in the acquisition and modeling of flows have led to improved descriptions of the dynamics, providing unsteady volumetric flow fields of unprecedented quality. Visual analysis of these high-dimensional flow data is increasingly challenging.

An example is found in patient-specific blood-flow fields, obtained through computational fluid dynamics models, or measured by imaging modalities. There is an increasing interest to understand these blood-flow fields. New insights may be obtained through quantitative and qualitative analysis, relating blood-flow patterns to the progression of diseases.

These insights potentially improve diagnosis and prognosis for different application areas. For instance, blood flow plays an important role in pathogenesis of various cardiovascular diseases [CCA*05, Ebb11].

Interpretation of unsteady velocity fields is exceedingly challenging. For the blood-flow application, the process involves a mental reconstruction of the anatomical structures, combined with a profound understanding of the blood-flow patterns in relation to the morphology. These patterns should be inspected throughout the cardiac cycle. Conventional slice-based inspection is ineffective for this purpose.

We present an abstract representation of time-resolved volumetric flow fields, based on spatiotemporal hierarchical clustering. Therefore, we introduce a coarse hierarchical clustering method, which efficiently clusters the unsteady velocity data. The clustering relies on a dissimilarity measure, which can be changed for different applications. We have adapted two distinct dissimilarity measures for steady volumetric vector fields for the spacetime domain.

Using the hierarchy, the desired level-of-detail can be selected intuitively. Lower levels are closer to the original field,




Figure 1: A sagittal slice of one phase of the 4D PC-MRI blood-flow velocity data, showing the acquired directions. (a) Right-to-left, (b) anterior-to-posterior, (c) feet-to-head.

and higher levels provide a sparse overview (Fig. 2). Typically, higher levels of abstraction yield comprehensive visualizations, reducing visual clutter and occlusion, while preserving important flow characteristics. We present a patharrow visualization to represent the obtained clusters. Furthermore, we have investigated several spatiotemporal seeding approaches, conveying the unsteady nature of the flow field. For the blood-flow application, an anatomical context visualization was included.

To demonstrate the clustering and visualization techniques, we present various unsteady volumetric blood-flow velocity fields. The data are mostly acquired by phase-contrast magnetic resonance imaging (PC-MRI) (Fig. 1). Additionally, we use simulated blood-flow velocity data. Although our methodology is largely motivated by the blood-flow application, no domain-specific knowledge is employed. All techniques can therefore be applied generically to other unsteady volumetric vector fields.

An overview of the hierarchical clustering and the blood-flow visualization is presented in figure 3. The results of the proposed clustering approach were validated using artificial data. We present a qualitative and a quantitative evaluation of the clustering results, assessing the parameter sensitivity and noise robustness of our method.

In summary, the main contributions of this work are:

- A coarse hierarchical clustering approach for time-resolved volumetric velocity fields, including two dissimilarity measures adapted for the spacetime domain.
- An illustrative visualization of spatiotemporal clusters, using representative patharrows. For the blood-flow application, we include an anatomical context.

2. Related work

Various flow visualization techniques have been employed to analyze unsteady vector fields, aiming to convey the essential aspects. For instance, there are various texture-based visualizations [LHD*04], which have shown to be highly effective for two-dimensional vector fields. However, in the three-dimensional domain, clutter and occlusion impair the visualization. The conveyed information is typically reduced, for instance through transparency [IG98].




Figure 2: Hierarchical clustering allows for intuitive level-of-detail selection, without a priori knowledge of the data domain.

Geometric primitives, such as particle spheres and integral lines, are also used to sparsely convey flow information [MLP*10]. Inspection is often performed locally, without a prior overview of the data. Local inspection is often based on regions-of-interest for seeding [HFS*11, vPOB*11].

Geometry-based visualizations are often used for the qualitative analysis of patient-specific blood-flow fields in pre-clinical research [CCA*05, Ebb11]. In the domain-specific literature, these visualization are used to communicate blood-flow patterns, for instance in the aorta [MDH*04, HH08], in relation to Marfan syndrome [KEW*04], or for congenital birth defects [VSSB10]. An extensive overview of applications was provided by Markl et al. [MKE*11].

Besides geometry-based visualizations, other techniques emanate from a data-driven approach, extracting features to simplify the vector fields. These features become the focus of the visualization [PVH*03]. Based on a set of features, a topological structure of the flow field can be derived and visualized [TS03]. Topology extraction and visualization for vector fields is an active field of research [PHCF12], where unsteady flows remain a challenging topic [PPF*11]. In addition, topology extraction is challenging for measured data, because the used derivative information is sensitive to noise.

Instead of detecting specific features in the vector field, an alternative approach is to cluster the vector field, according to a pre-defined measure of dissimilarity. Telea and Van Wijk presented a clustering approach for steady flow fields [TvW99]. They cluster the field in a bottom-up approach, based on an elliptic dissimilarity function. Another bottom-up approach for steady flow fields was presented by Kuhn et al. [KLG*11]. Instead of a hierarchical approach, they employ a scale-dependent analysis of integral curves, limiting the detail selection to a fixed number of scales.

Alternatively, clustering may be performed without boolean merging or splitting, using a diffusion process to enhance correlation in the cluster set [GPR*00, GPR*01]. Although untruthful delineation of cluster boundaries is avoided, a diffusion process becomes computationally expensive in the spatiotemporal domain, especially for larger data sets. Moreover, storage of the velocity fields at various scales is costly in terms of memory usage.

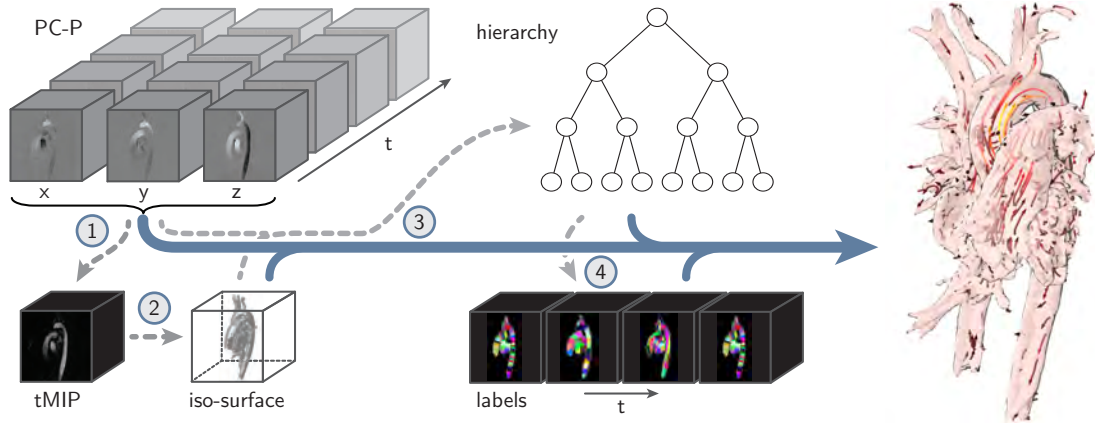


Figure 3: Overview of the visualization framework, based on spatiotemporal hierarchical clustering. The gray dashed arrows depict pre-processing steps. (1) A tMIP volume is generated, and (2) an iso-threshold captures the voxels that are clustered. (3) Next, the cluster hierarchy is constructed. (4) Using the cluster tree, labels are generated per cardiac phase. After pre-processing, the real-time visualization is generated using the available data structures, as depicted by the solid blue arrow.

To the best of our knowledge, only Yu et al. [YWM07] have presented a clustering technique for unsteady volumetric vector fields, dealing with the difficult relation between the space and time domains. In accordance to our work, they employ a hierarchical clustering, providing level-of-detail selection. However, Yu et al. have employed a single dissimilarity measure, based on the work by Telea and Van Wijk [TvW99]. Moreover, their clustering relies on an adaptive grid sampling, which is based on data-specific features. Instead, we present a generic coarse hierarchical clustering approach, independent of the application domain. Obtaining application-specific knowledge is laborious, and requires to contemplate a wide variety of cases.

In the following, we present visualizations of unsteady volumetric blood-flow information, based on spatiotemporal hierarchical clustering. In contrast to the work by Kuhn et al. [KLG*11], we do not aim to visualize the cluster boundaries, as they do not represent a physical entity. Instead, we adopt the notion of representative integral lines. These lines are seeded at the center of the clusters, and are assumed to represent the flow within each cluster. Patharrows are used to visualize the flow, based on these integral lines.

3. Clustering

Hierarchical clustering benefits from an intuitive level-of-detail selection. No prior domain knowledge is required, because the method directly employs the data characteristics. Furthermore, no a priori selection of the number of clusters is needed. The number of clusters depends on the dimensions of the data set, as well as the complexity of the flow patterns, which cannot be predicted beforehand. Hierarchical clustering is scalable to high-dimensional data, and hence suitable for time-resolved velocity fields.

3.1. Preprocessing

Before clustering, a few preprocessing steps are required (Fig. 3, steps ① and ②). First, a temporal maximum intensity projection (tMIP) is generated, providing a static representation of the time-varying flow speed [vPOB*10]. Next, an iso-threshold is applied to the tMIP volume, capturing the voxels that are considered for clustering. The velocities and the voxel mask are then used to perform the clustering (Fig. 3, step ③). The voxel mask reduces the amount of voxels that are included in the cluster process, increasing the computational performance. This approach can be applied generically to unsteady flow data on regularly-sampled grids.

3.2. Spatiotemporal hierarchy

The hierarchy is described by a precomputed binary cluster tree. A binary tree of the flow data can be constructed by a top-down [HWHJ99] or a bottom-up approach [TvW99].

The clustering is initiated from the data voxels, defined as the leaves of the cluster tree. Top-down division can be performed in various ways, leading to different results. Instead, we employ a bottom-up approach, consistently generating the desired hierarchy. Iterative merging of individual clusters leads to the tree nodes, based on dissimilarities between spatiotemporally neighboring clusters. Our optimized approach is described in subsection 3.4. In the following subsection, we present the selected dissimilarity measures (Fig. 4).

3.3. Four-dimensional dissimilarity measures

We have adapted two measures for the spatiotemporal domain. A physically substantiated coupling between the spatial and the temporal domain is nontrivial, and hence we treat the dimensions equally for our dissimilarity measures.




Figure 4: The clustering yields different results, using various dissimilarity measures. (a) An artificial vector field is clustered using: (b) the elliptical dissimilarity ($\alpha = 0.9$, $\omega = 0.9$), and (c) the local linear expansion dissimilarity.

An elementary entity to represent a cluster is the average flow speed, based on the magnitudes of the velocity vectors. However, flow velocities with deviating directions should not be merged into a cluster. Therefore, the average $\bar{\mathbf{v}}$ of all velocity vectors in a cluster seems appropriate. The dissimilarity between clusters can also be measured by an average of directions. However, this approach does not take into account position information, nor does it consider specific flow patterns. Therefore, two extended dissimilarity measures are used, combining velocity and position information.

3.3.1. Elliptical

The elliptical dissimilarity measure (Fig. 4b), introduced by Telea and Van Wijk [TVW99], includes spatial information using the positional average $\bar{\mathbf{x}}$. The measure d_{ell} combines a velocity dissimilarity d_v , and a positional dissimilarity d_p in a linear fashion, controlled by weighting parameter ω :

$$d_{ell}(\bar{\mathbf{x}}_1, \bar{\mathbf{v}}_1, \bar{\mathbf{x}}_2, \bar{\mathbf{v}}_2) = \omega d_p(\bar{\mathbf{x}}_1, \bar{\mathbf{x}}_2) + (1 - \omega) d_v(\bar{\mathbf{v}}_1, \bar{\mathbf{v}}_2).$$

The velocity dissimilarity d_v is similar to a Euclidean distance, with the exception that a higher cost is associated with flow deviations perpendicular to the average cluster flow. This anisotropic behavior is achieved by modeling the iso-error contours around the flow vector with an ellipse.

The positional dissimilarity d_p is defined accordingly. The iso-error contours describe the positional dissimilarity, modeled with an ellipse oriented along the velocity vector. The aspect ratio of the ellipse is parameterized by a parameter α .




Figure 5: Schematic depiction of a bottom-up generation of a full cluster hierarchy. Tree nodes are merged, iteratively processing the two clusters with the smallest dissimilarity.

For the time-varying data, positions are defined in space-time: $\bar{\mathbf{x}}_{4D} = (x, y, z, t)$. The dimensional disparity is solved by extending the velocity component with a constant, transforming the component $\bar{\mathbf{v}}_{4D} = (u, v, w, 1)$ into spacetime.

3.3.2. Local linear expansion

Instead of average positions $\bar{\mathbf{x}}$ and velocities $\bar{\mathbf{v}}$, local linear approximations of the flow field may be used to preserve flow patterns. Carmo et al. [CNPHY04] presented the local linear expansion dissimilarity measure, which aims for a linear model that estimates the velocities \mathbf{v}_i within each cluster (Fig. 4c). The linear model describes the relation between the voxel positions \mathbf{x}_i and the modeled velocity $\hat{\mathbf{v}}_i$ as:

$$\hat{\mathbf{v}}_i = \mathbf{A}\mathbf{x}_i - \bar{\mathbf{v}},$$

where \mathbf{A} is a matrix describing the velocity gradients, estimated by a least squares fitting over N cluster voxels. For a cluster C , the total squared error ϵ is given by:

$$\epsilon(C) = \sum_{i=1}^N \|\hat{\mathbf{v}}_i - \mathbf{v}_i\|^2, \text{ where } \mathbf{v}_i \in C.$$

The dissimilarity measure d_{lin} is defined by the cost to merge two clusters, C_1 and C_2 , into a new cluster C_{merged} :

$$d_{lin}(C_1, C_2) = \epsilon(C_{merged}) + (\epsilon(C_1) + \epsilon(C_2)).$$

To the best of our knowledge, this measure has not been applied to time-varying velocity data before. To achieve this, we extend matrix \mathbf{A} , proposed by Carmo et al. [CNPHY04], to include the spatial and temporal gradients, resulting in a 4×3 matrix that captures the relation between the spatiotemporal position \mathbf{x}_i , and the three-dimensional velocities \mathbf{v}_i .

3.4. Coarse hierarchical clustering

Processing of high-dimensional data inherently leads to significant computational cost. Hence, clustering is generally applied to small 2D and 3D examples (e.g., [TVW99]). To obtain acceptable performance for actual measured data, the computational complexity should be reduced, without downscaling the data dimensions by plain subsampling.




Figure 6: Schematic depiction of a coarse bottom-up generation of the cluster hierarchy, requiring half the number of iterations in comparison to a full cluster tree generation.

A major performance bottleneck is the repetitive search for the smallest dissimilarity. This operation has a complexity of $O(n)$, and the number of iterations also scales with $O(n)$. The overall complexity therefore becomes $O(n^2)$ for a straightforward implementation. In particular, the initial phase of the clustering is computationally expensive, because the number of clusters is still relatively large.

Therefore, we introduce the coarse hierarchical clustering (CHC) approach. Instead of merging a single pair of clusters (Fig. 5), the CHC method merges multiple clusters per iteration (Fig. 6). An incremental dissimilarity threshold d_{th} is introduced, causing all cluster pairs with a dissimilarity smaller than this threshold to merge simultaneously. For each iteration, d_{th} is increased with a factor Δd_{th} . This reduces the computational complexity by lowering the number of iterations, as well as the number of dissimilarity computations.

We have performed several experiments with varying Δd_{th} values. For the blood-flow application, the algorithm approaches a linear complexity for an increment of $\Delta d_{th} = 0.01$. This threshold yields a good balance between computational cost and cluster quality for both dissimilarity measures. The CHC cluster results were visually compared with the full cluster tree, as depicted in figure 7. We observe good correspondence in the temporal evolution of both cases.

Merging of multiple clusters per iteration affects the resulting tree, and hence provides an approximation of the true data hierarchy. However, it is fair to assume that small variations in the order of clustering hardly affect the cluster tree structure. Whenever clusters within an iteration are neighbors, small deviations will occur. These deviations depend on the user-defined threshold increment Δd_{th} .

In addition to the algorithm complexity, performance of the clustering also depends on other aspects. For instance, we assume a substantial amount of background information in the data, which is discarded prior to processing. Furthermore, implementation aspects greatly influence the performance, such as the data structure that stores dissimilarity values. The operations performed on the container include




Figure 7: Qualitative comparison of the cluster boundaries, based on a planar reformat through the volunteer PC-MRI blood-flow data at 330ms. The selected hierarchy level is 99.98%. (a) A full hierarchical clustering, in comparison to (b) the coarse hierarchical clustering (CHC) approach, with a threshold increment of $\Delta d_{th} = 0.01$.

Table 1: Complexity of specific container operations on the set of dissimilarities, used for the hierarchical clustering.

Operation	Dynamic array	Linked list	Balanced tree
Lookup	$O(n)$	$O(n)$	$O(1)$
Insertion	$O(1)$	$O(1)$	$O(\log n)$
Deletion	$O(n)$	$O(1)$	$O(1)$

insertion, deletion, and lookup of the minimum value. The operation complexity varies per container type (Table 1).

The unsorted dynamic array container is discarded, due to the relatively expensive lookup and insertion. In comparison to the unsorted linked list, insertion is more expensive with a sorted balanced tree, due to an additional search operation. However, lookup and deletion are efficient (Table 1).

We have implemented the clustering with these container types, employing the Standard Template Library (STL) for the C++ programming language. This confirmed the theoretical complexity $O(n^2)$ for the linked list, however the balanced tree did not nearly achieve the complexity of $O(n \log n)$. The best performance was obtained by our custom performance-optimized implementation of a linked list.

4. Visualization

4.1. Hierarchy selection

The visualization relies on a selected level-of-detail, using an intuitive percentage of the cluster tree. Comprehensive visualizations are obtained at the top-most percentages of the tree. The level-of-detail is the only essential parameter of the visualization, and should be explored by the user. Based on the selected level-of-detail, the spatial structure of the clusters can be reconstructed for each time-point. The clusters are represented by a labeled volume (Fig. 3, step ④).

4.2. Center point selection

The cluster boundaries is not directly visualized, as they convey no evident physical entity in the blood-flow field. Instead, we show the clusters by representative pathlines, seeded at the cluster centers. These cluster centers can be defined spatially, or in the spacetime domain.

The 3D cluster centers can be generated from the reconstructed clusters per time-point. An intuitive approach would be to compute the center of mass. However, this center may reside outside the boundaries for concavely shaped clusters.

Instead, we devise an approach that locates the point within the cluster C with the smallest distance to all N positions within that cluster. Using the Euclidean distance, the center position \mathbf{x}_m is defined by its index m as:

$$m(C) = \underset{i}{\operatorname{argmin}} \frac{1}{N} \sum_{j=1}^N \|\mathbf{x}_i - \mathbf{x}_j\|, \text{ where } \mathbf{x}_{i,j} \in C.$$




Figure 8: Simplified representation of the patharrow seeding. The horizontal axis shows the time, while the spatial domain is represented on the vertical axis. One time-frame is highlighted in yellow. (a) For 3D seeding, a spatial center is computed for each existing cluster, and short pathlines are traced interactively. (b) For 4D seeding, a spatiotemporal center is computed, and long pathlines are precomputed.

To avoid a disproportionate performance penalty, we avoid computing all mutual distances. As the visualization does not require the exact center, we perform a stochastic sampling per cluster, and then compute the center. Based on visual inspection, a minimum of hundred samples is sufficient.

The clustering is performed in four dimensions, and hence the clusters reside in the spacetime domain. However, a physically meaningful definition of the cluster centers in 4D is nontrivial, because of the disparity between space and time, in terms of units and dimensionality. To prevent an unsubstantiated coupling between space and time, we first select the temporal center by computing the median of the time spanned by the 4D cluster. Subsequently, the spatial center is computed, similar to a 3D cluster center.

4.3. Representative patharrows

The four-dimensional constructs are inherently challenging to communicate to a human observer. Moreover, the patterns in the unsteady flow data are continuous entities, which are challenging to convey, because of the lack of a clear visual demarcation. In three dimensions, the visualization is furthermore impeded by visual clutter and occlusion.

We represent the spatiotemporal clusters by discrete pathlines, traced forward and backward in time, starting from the cluster centers. By default, the lines are traced for a time period of two time-frames. The lines are generally not restricted to the cluster boundaries.

We assume that the integral lines capture the time-varying structure of the flow field for each cluster, based on spatiotemporally well-distributed seed points, imposed by the clustering. The extent to which the integral lines represent the flow field depends, inter alia, on the dissimilarity measure.

The representative integral lines are depicted as arrows in the direction of the flow, rendered in real-time. The arrows are composed from illuminated tuboid imposters, with arrowheads and end-caps (Fig. 9). A procedural texture is ap-

plied, superimposing an arrowhead pattern that indicates the local direction on the tuboids. The long patharrows do not have to be traced to the end to determine the direction. Starting from either a 3D or a 4D center point, we devise two approaches to generate the patharrows.

4.3.1. 3D-seeded arrows

The cluster centers provide an intuitive data-driven seeding of the patharrows. However, the 3D cluster centers have no explicit temporal coherence, as the spatial centers are computed for all clusters that exist at the inspected time-point (e.g., $t = 1$), without considering the temporal extent of the cluster (Fig. 8a). This results in a relatively dense seeding, in comparison to the use of 4D centers (Fig. 8b). Naturally, the visual denseness also depends on the selected level-of-detail.

At the inspected time-point, short pathlines are traced in real-time for each cluster center, and visualized as patharrows. The resulting renditions are not suitable for animation, due to the lack of temporal coherence between the 3D cluster centers. However, this approach provides a comprehensive overview of the structure of the flow at a given time-point.

4.3.2. 4D-seeded arrows

Alternatively, the sparser set of 4D cluster centers can be used to seed the patharrows in a temporally coherent way. Long pathlines are used to represent each cluster (Fig. 8b).

Based on this set of precomputed pathlines, we present several visualization styles (Fig. 9). First, the whole set of pathlines can be visualized directly using patharrows. This provides a static overview of the full spatiotemporally clustered flow data. To obtain a sufficiently sparse representation, a high level-of-detail is required. The temporal characteristics can be regained by color mapping the time component (Fig. figure 9a). Alternatively, the local speed can be encoded by color (Fig. 9b), losing temporal information.




Figure 9: Close-up of 4D-seeded patharrows. (a) A time color coding reveals the temporal character. (b) Color may also convey the flow speed. (c) The pathlines are employed for animation. The color is desaturated, except for a sliding window around the animation time-point. (d) Alternatively, small patharrows may represent the sliding window. Note the color correspondence between (b), (c), and (d).

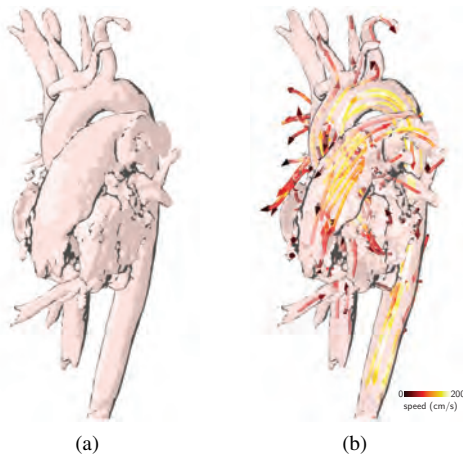


Figure 10: Blood-flow depicted by 3D-seeded patharrows with anatomical context. (a) A cel shaded surface with contours based on the surface curvature. (b) Opacity modulation of the frontfaces ensures visibility of the patharrows.

Because of the temporal coherence, the static lines may be used for animation. Therefore, we highlight parts of the lines within a certain time window around the animated time-point (Fig. 9c-d). The sliding window can be depicted as a small patharrow, omitting the remainder of the pathlines outside of the window (Fig. 9d). This sliding window is temporally interpolated, and smoothly appears and vanishes at the end points of the underlying integral line.

To inspect the structure of the flow field more closely, the full set of lines can be visualized as long patharrows, rendering them with a fully desaturated color. The parts within the sliding window are emphasized by color (Fig. 9c). During animation, this window moves along the long patharrows, highlighting flow patterns. These approaches are essentially identical, shown by the corresponding regions in figures 9c and d. The temporal interpolation also coincides.

Both animation styles comprehensively depict the temporal behavior of flow patterns. Human observers are by nature exceptionally good at recognizing such patterns. Despite the imposed cognitive load [TMB02], the animations take advantage of the intrinsic human perception of motion, facilitating understanding of the spatial relations and depth.

4.4. Anatomical context

In addition to the flow visualizations, we provide an anatomical context (Fig. 10a), based on an iso-surface extracted from the tMIP volume. To avoid visual occlusion, frontfaces are rendered using a view-dependent opacity modulation (Fig. 10b), based on the Fresnel term [GNKP10]. The backfaces are rendered with cel shading [vPOB*10]. In addition, we render contours, using the curvatures on the surface.

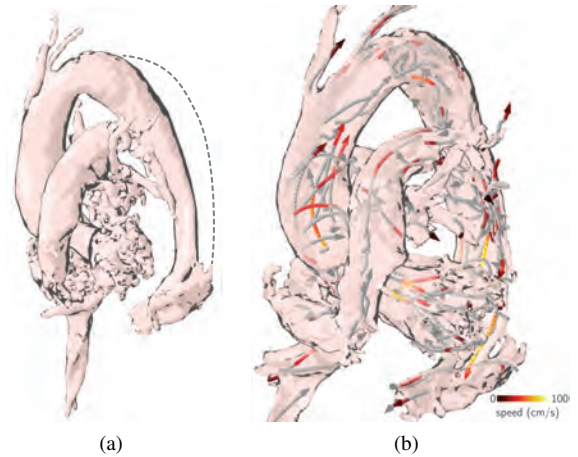


Figure 11: Aortic dissection with pathological flow. (a) Anatomical context based on a tMIP does not show the secondary system (dashed line). (b) A 4D-seeded patharrow animation highlights patterns. A window around 320 ms is shown.

5. Results

5.1. Data sets

The first data set comprises a PC-MRI blood-flow measurement of a healthy volunteer (Fig. 10b). The second data set includes pathological blood flow, used to demonstrate detection of abnormal flow patterns. The data are acquired from patient suffering from an aortic dissection (Fig. 11 and 12). The vessel wall is separated due to a tear in the vessel wall, effectuating a secondary system with slowly flowing blood. The slow flow is not captured by the tMIP (Fig. 11a), and is hence omitted by the clustering threshold. However, anomalous blood-flow patterns occur in the main bloodstream. For an application-specific analysis, the anatomical context of the dissection should be based on an anatomical scan. The patient study was approved by the local research ethics committee (study no. 08/H0809/49).

The third data set consists of a computational fluid dynamics (CFD) simulation of the blood flow in a cerebral aneurysm (Fig. 13). The temporal resolution has the same order of magnitude, while the spatial resolution is considerably better than the PC-MRI examples. The blood flow interacts with the surrounding tissues, and hence relations with the vessel wall properties are important. The patharrows are therefore shown in relation to the context surface with a color mapping of the wall-shear stress (WSS) [Day90].

5.2. Qualitative evaluation

The cluster shapes were visually evaluated for different data sets. Based on the patharrows, we impose several requirements on the clusters. The clusters need to be spatially coherent, compact, and shaped along the blood-flow direction.




Figure 12: Aortic dissection with pathological blood flow, depicted by long patharrows, based on 250 clusters using the elliptical dissimilarity. The 4D-seeded short patharrows convey the spatiotemporal structure, revealing flow patterns during animation. The notable regions are indicated.

The linear expansion dissimilarity d_{lin} yields compact clusters, which are spatially coherent, and capture patterns such as bifurcations. However, the cluster shape does not necessarily align with the flow direction. Instead, the elliptical measure d_{ell} shows compact clusters, which can be parameterized [TvW99]. The positional part ensures spatial coherence, while the velocity dissimilarity leads to clusters that are shaped along the blood-flow. The preferred parameters were empirically determined: $\alpha = 0.9$, and $\omega = 0.9$.

Using the elliptical dissimilarity, we have inspected blood-flow patterns. As expected, blood flow of the volunteer primarily shows laminar behavior. In contrast, several anomalous patterns are detected for the pathological case (Fig. 12), including high vorticity regions.

For the inspected data, the patharrows are well-distributed in space and time. The provided overviews are complementary to local inspection techniques that enable detailed pattern analysis. For the blood-flow application, local inspection was for instance presented using probing techniques [vPOB*10, vPOB*11].

5.3. Quantitative evaluation

The clustering was evaluated quantitatively, considering noise robustness and performance. The noise in the velocity data approximates a Gaussian distribution, with a signal-to-noise ratio ranging between 10 and 25. The experiments were performed on artificial flow data, varying the standard deviation of the additive Gaussian noise. The cluster results are compared with the ground truth using the adjusted Rand index [HA85]. The linear expansion dissimilarity shows a varying index, while the elliptical dissimilarity is stable.




Figure 13: A CFD simulation of a cerebral aneurysm, depicted by static patharrows, based on 4D seed positions. A high-speed in- and outflow is observed from the parent vessel. Recirculating patterns occur within the aneurysm, which are more clearly conveyed through animation.

A performance evaluation of both the clustering and the rendering was carried out. A computer system was used, with a dual core processor, 6GB memory, and an NVidia GeForce 570 GTX graphics card. For each data sets, systole and early diastole were clustered, capturing the important dynamics. Larger data sets currently lead to memory footprint issues. A tMIP volume was thresholded at 5% of its intensity range. The selected voxels were clustered using different dissimilarity measures. Table 2 summarizes the results.

For the elliptical measure d_{ell} , we observe significant improvements in the clustering computation time using the CHC approach, in comparison to a full tree generation (Table 2). The linear model dissimilarity d_{lin} is computationally more expensive, in particular for noise-prone PC-MRI data.

For an abstract level-of detail, all visualizations render at interactive framerates over fifteen frames per second, including anatomical context (Table 2). Lower levels lead to more clusters, resulting in lower framerates.

6. Discussion

The presented techniques were implemented using the C++ programming language, supported by the OpenGL library, the visualization toolkit (VTK), and the Qt graphical user-interface library. Interactive inspection and animation require high performance rendering. Modern consumer graphics hardware were exploited, using the GLSL shading language.

Clustering - The presented data-driven approach generically clusters the unsteady vector data. The domain-specific knowledge is emerging, and hence no explicit knowledge exists to identify clusters in the unsteady blood-flow data. Therefore, we rely on properties of the blood-flow field.

Table 2: Performance of clustering and visualization. For the d_{ell} measure ($\alpha = 0.9, \omega = 0.9$), a full hierarchy and a CHC approximation ($\Delta d_{th} = 0.01$) is included. The frame-rates (FPS) are based on a cluster set near the tree top (99.9%). We show the spatial dimensions, the number of time-points, and the maximum speed that occurs in the data.

	Volunteer	Dissection	Aneurysm
Dimensions	144 × 144 × 50	144 × 144 × 60	144 × 128 × 64
Resolution (mm)	2.1 × 2.1 × 2.7	2.1 × 2.1 × 2.5	0.2 × 0.1 × 0.3
Time-points	20	25	50
Iso-threshold	5%	5%	5%
Max speed (cm/s)	180	240	250
Cluster time-points	[1,12]	[1,12]	[1,10]
Cluster nodes	621.8K	744.7K	589.9K
Cluster d_{ell} full	08h:43m	12h:08m	07h:33m
Cluster d_{ell} chc	58s	67s	53s
Cluster d_{in} chc	02h:20m	02h:02m	01h:04m
Cluster count	366	890	356
FPS dynamic	23.7	19.8	16.6
FPS static	23.8	19.9	17.1
Figures	7, 10	11, 12	13

Although the cluster boundaries represent regions of coherent velocities, they do not truly exist in the continuous flow field. They were not visualized to avoid false interpretation of the boundaries as physical entities in the field. Consequently, coverage of the patharrows may be hard to derive, especially in regions with opposing arrows.

The coherence depends on the dissimilarity measure. The definition of such a measure is complicated. A non-trivial weighting between space and time is required, as the units of these dimensions cannot be compared.

Visualization - Related work employs pathline visualizations, similar to our patharrows. Clinically-oriented papers show seeding throughout the cardiovascular anatomy (e.g., [MKE*11]). However, the densely sampled seed-points are positioned randomly, and not based on the underlying data. Therefore, the clustering provides a data-driven seeding approach, for a coarse to detailed view.

An important aspect of the visualization is the level-of-detail selection. The selection is currently global, and does not provide local detailed representations on demand. Construction of labeled cluster volumes is performed when the level-of-detail is changed. With the used data structures, real-time localized level-of-detail changes need precomputed cluster label volumes, leading to excessive memory use.

The temporal behavior of the intrinsically unsteady flow data is captured using static and animated patharrows. Although animations impose a considerable cognitive load [TMB02], the human observer is sensitive to the apparent motion in flow [SI10, WSE05]. Using animated short patharrows (Fig. 9d), the patterns attract the attention. The structure of the flow field can be investigated in more detail using static patharrows, possibly with sliding windows (Fig. 9c).

7. Conclusions and future work

In conclusion, we have presented a coarse hierarchical clustering technique for four-dimensional flow data. The CHC approach achieves a near linear computational complexity, in contrast to the quadratic complexity of a naive implementation. Two dissimilarity measures were adapted for the spacetime domain.

Clusters at a selected hierarchy level were represented by patharrows. We introduced 3D-seeded and 4D-seeded patharrow visualizations, combined with an illustrative anatomical context. The 4D-seeded patharrows allow for a static representations and animations of the spatiotemporal information.

The clustering and visualization was applied to time-resolved volumetric blood-flow data, either measured through PC-MRI, or simulated by CFD. Different flow patterns could be distinguished. These patterns can be related to properties of the vessel wall.

In the future, the cluster approach can be improved. The incremental threshold for the CHC approach could be adjusted during the clustering, according to the inter-cluster dissimilarity. Alternatively, parallelization of the clustering could be investigated, to further improve the performance. Interactive level-of-detail selection, together with associated visualization techniques, could provide details on demand.

The current visualization techniques depend on camera interaction and animation to reveal the spatial relations. Different depth cues, such as shadows and halos, may be investigated. Also, time could be conveyed by illustrative techniques, and additional information could be visualized, such as the intra-cluster variance.

Moreover, it is worthwhile to investigate application-specific approaches. The generic clustering can be fine-tuned using domain-specific dissimilarity measures. This will lead to different cluster results, which may require alternative visualizations. Also for the blood-flow application, further research is needed to address the specific goals of physicians. Specific dissimilarity measures could cluster regions with high vorticity, or low coherence of the blood-flow field. The clustering results and associated visualizations should be evaluated, measuring the value for preclinical research, and potentially future diagnosis.

Acknowledgements

The volunteer and patient 4D MRI blood-flow data were provided courtesy of the division of Imaging Sciences, King's College London at St. Thomas' hospital. In particular, we express our sincere gratitude and appreciation to professor T. Schaeffter, Dr. R. E. Clough, and Dr. I. Valverde. The simulated blood-flow data was provided courtesy of the Institute of Fluid Mechanics and Thermodynamics (ISUT) of the University of Magdeburg. We therefore especially thank Dr. G. Janiga, and R. Gasteiger MSc. We sincerely thank professor C. Hansen for providing the artificial flow data.

References

- [CCA*05] CEBRAL J. R., CASTRO M. A., APPANABOYINA S., PUTMAN C. M., MILLAN D., FRANGI A. F.: Efficient pipeline for image-based patient-specific analysis of cerebral aneurysm hemodynamics: technique and sensitivity. *IEEE Transactions on Medical Imaging* 24, 4 (2005), 457–467. 1, 2
- [CNPBY04] CARMO B., NG Y., PRÜGEL-BENNETT A., YANG G.-Z.: A data clustering and streamline reduction method for 3D MR flow vector field simplification. In *Medical Image Computing and Computer-Assisted Intervention* (2004), vol. 3216, Springer Berlin / Heidelberg, pp. 451–458. 4
- [Day90] DAY M. A.: The no-slip condition of fluid dynamics. *Erkenntnis* 33, 3 (1990), 285–296. 10.1007/BF00717588. 7
- [Ebb11] EBBERS T.: Flow Imaging: Cardiac Applications of 3D Cine Phase-Contrast MRI. *Current Cardiovascular Imaging Reports* 1, 4 (2011), 1–7. 1, 2
- [GNKP10] GASTEIGER R., NEUGEBAUER M., KUBISCH C., PREIM B.: Adapted surface visualization of cerebral aneurysms with embedded blood flow information. In *EG Workshop on Visual Computing for Biology and Medicine* (2010), pp. 25–32. 7
- [GPR*00] GARCKE H., PREUßER T., RUMPF M., TELEA A., WEIKARD U., VAN WIJK J. J.: A continuous clustering method for vector fields. In *IEEE Visualization* (2000), IEEE Computer Society Press, pp. 351–358. 2
- [GPR*01] GARCKE H., PREUßER T., RUMPF M., TELEA A. C., WEIKARD U., VAN WIJK J. J.: A phase field model for continuous clustering on vector fields. *IEEE Transactions on Visualization and Computer Graphics* 7, 3 (2001), 230–241. 2
- [HA85] HUBERT L., ARABIE P.: Compacting partitions. *Journal of Classification* 2 (1985), 139–218. 8
- [HFS*11] HENNEMUTH A., FRIMAN O., SCHUMANN C., BOCK J., DREXL J., HUELLEBRAND M., MARKL M., PEITGEN H.-O.: Fast interactive exploration of 4D MRI flow data. In *Society of Photo-Optical Instrumentation Engineers* (2011), vol. 7964, pp. 79640E–1 – 79640E–11. 2
- [HH08] HOPE T. A., HERFKENS R. J.: Imaging of the thoracic aorta with time-resolved three-dimensional phase-contrast MRI: a review. *Thoracic and cardiovascular surgery* 20, 4 (2008), 358–64. 2
- [HWHJ99] HECKEL B., WEBER G., HAMANN B., JOY K. I.: Construction of vector field hierarchies. In *IEEE Visualization* (1999), IEEE, pp. 19–25. 3
- [IG98] INTERRANTE V., GROSCH C.: Visualizing 3D flow. *IEEE Computer Graphics and Applications* 18 (1998), 49–53. 2
- [KEW*04] KVITTING J.-P. E., EBBERS T., WIGSTRÖM L., ENGVALL J., OLIN C. L., BOLGER A. F.: Flow patterns in the aortic root and the aorta studied with time-resolved, 3-dimensional, phase-contrast magnetic resonance imaging: implications for aortic valve-sparing surgery. *Journal of Thoracic and Cardiovascular Surgery* 127, 6 (2004), 1602–1607. 2
- [KLG*11] KUHN A., LEHMANN D. J., GASTEIGER R., NEUGEBAUER M., PREIM B., THEISEL H.: A Clustering-based Visualization Technique to Emphasize Meaningful Regions of Vector Fields. In *Vision, Modeling, and Visualization* (2011), pp. 191–198. 2, 3
- [LHD*04] LARAMEE R. S., HAUSER H., DOLEISCH H., VROLIJK B., POST F. H., WEISKOPF D.: The State of the Art in Flow Visualization: Dense and Texture-Based Techniques. *Computer Graphics Forum* 23, 2 (2004), 203–221. 2
- [MDH*04] MARKL M., DRANEY M. T., HOPE M. D., LEVIN J. M., CHAN F. P., ALLEY M. T., PELC N. J., HERFKENS R. J.: Time-resolved 3-dimensional velocity mapping in the thoracic aorta: visualization of 3-directional blood flow patterns in healthy volunteers and patients. *Journal Of Computer Assisted Tomography* 28, 4 (2004), 459–468. 2
- [MKE*11] MARKL M., KILNER P. J., EBBERS T., MICHAEL M., PHILIP K., TINO E.: Comprehensive 4D velocity mapping of the heart and great vessels by cardiovascular magnetic resonance. *Journal of Cardiovascular Magnetic Resonance* 13, 7 (2011), 1–22. 2, 9
- [MLP*10] MCLOUGHLIN T., LARAMEE R. S., PEIKERT R., POST F. H., CHEN M.: Over Two Decades of Integration-Based, Geometric Flow Visualization. *Computer Graphics Forum* 29, 6 (2010), 1807–1829. 2
- [PHCF12] PEIKERT R., HAUSER H., CARR H., FUCHS R. (Eds.): *Topological Methods in Data Analysis and Visualization II – Theory, Algorithms, and Applications*. Mathematics and Visualization. Springer, 2012. 2
- [PPF*11] POBITZER A., PEIKERT R., FUCHS R., SCHINDLER B., KUHN A., THEISEL H., MATKOVIC K., HAUSER H.: The state of the art in topology-based visualization of unsteady flow. *Computer Graphics Forum* 30, 6 (2011), 1789–1811. 2
- [PVH*03] POST F. H., VROLIJK B., HAUSER H., LARAMEE R. S., DOLEISCH H.: The state of the art in flow visualisation: feature extraction and tracking. *Computer Graphics Forum* 22, 4 (2003), 775–792. 2
- [SI10] STEELE J., ILINSKY N.: *Beautiful Visualization: Looking at Data Through the Eyes of Experts*, chapter 19 ed. O'Reilly Media, 2010. 9
- [TMB02] TVERSKY B., MORRISON J. B., BETRANCOURT M.: Animation: can it facilitate? *International Journal of Human-Computer Studies* 57, 4 (2002), 247 – 262. 7, 9
- [TS03] THEISEL H., SEIDEL H.-P.: Feature flow fields. In *Symposium on Data visualisation* (2003), VISSYM, pp. 141–148. 2
- [TvW99] TELEA A., VAN WIJK J. J.: Simplified representation of vector fields. In *IEEE Visualization* (1999), pp. 35–42. 2, 3, 4, 8
- [VPOB*10] VAN PELT R., OLIVÁN BESCÓS J., BREEUWER M., CLOUGH R. E., GRÖLLER M. E., TER HAAR ROMENY B., VILANOVA A.: Exploration of 4D MRI blood flow using stylistic visualization. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 1339–1347. 3, 7, 8
- [VPOB*11] VAN PELT R., OLIVÁN BESCÓS J., BREEUWER M., CLOUGH R. E., GRÖLLER M. E., TER HAAR ROMENY B., VILANOVA A.: Interactive Virtual Probing of 4D MRI Blood-Flow. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2153–2162. 2, 8
- [VSSB10] VALVERDE I., SIMPSON J., SCHAEFFTER T., BEERBAUM P.: 4D phase-contrast flow cardiovascular magnetic resonance: comprehensive quantification and visualization of flow dynamics in atrial septal defect and partial anomalous pulmonary venous return. *Pediatric cardiology* 31, 8 (2010), 1244–8. 2
- [WSE05] WEISKOPF D., SCHAFFITZEL T., ERTL T.: Real-time advection and volumetric illumination for the visualization of 3D unsteady flow. In *EG/IEEE EuroVis* (2005), vol. 5, pp. 13–20. 9
- [YWM07] YU H., WANG C., MA K.-L.: Parallel hierarchical visualization of large time-varying 3D vector fields. In *IEEE Supercomputing* (2007), ACM Press, p. 1. 3