Enhancing Visibility of Blood Flow in Volume Rendered Cardiac 4D PC-MRI Data

Benjamin Behrendt[1], Benjamin Köhler[1], Uta Preim[2], Bernhard Preim[1]

[1]Dept. of Simulation and Graphics, University of Magdeburg, Germany
[2]Dept. of Diagnostic Radiology, Municipal Hospital Magdeburg, Germany
behrendt@isg.cs.uni-magdeburg.de

Abstract. Four-dimensional phase-contrast magnetic resonance imaging (*4D PC-MRI*) is a method to non-invasively acquire blood flow in the aorta. This flow is commonly visualized as path lines inside of the vessels. Direct volume rendering (*DVR*) uses a transfer function to directly render the dataset without needing a manual segmentation. Since the transfer function can be manipulated on the fly, DVR allows fast exploration of the dataset. Using a simple intensity-based transfer function, however, either the intravascular blood flow would be hidden behind the vessel's front side or the entire vessel has to be culled from the visualization. Therefore, we propose an automated mechanism that reveals the vessel anatomy by removing their front sides based on the viewing direction. This creates an effect similar to frontface culling on surface renderings. The visibility of focus objects inside the anatomy is guaranteed while spatial awareness is mostly maintained due to the presence of anatomical structures as context information. While we were able to confirm the effectiveness of our method in an interview with a collaborating radiologist, it still proved to be somewhat limited by the data quality and lack of a manual segmentation.

1 Introduction

Cardiovascular pathologies have been related to blood flow pattern changes in current research [1,2]. They can be assessed non-invasively by analyzing 4D PC-MRI data, producing time-resolved three-dimensional datasets. To study these pathologies as well as their causes and effects, it is vital to analyze and understand these patterns. In order to make these analyses more objective, blood flow patterns can also be identified automatically [3].

A common way to visualize 4D PC-MRI datasets is to show the flow as path lines inside their respective vessels, creating a visualization with multiple, overlapping features. How to deal with these features strongly depends on the rendering method. When using indirect volume rendering and working with a polygonal mesh extracted from a segmentation, detecting and removing overlapping features can be achieved by using frontface culling or employing the Z-buffer. A disadvantage of polygonal meshes is that they require a rather detailed segmentation. Therefore, the adequacy of indirect volume rendering for exploring datasets is limited.

Direct volume rendering is better suited for exploration, as the transfer function can be changed in real time and the user receives instant visual feedback. In addition, it allows a qualitative analysis of the blood flow without the need for performing a manual segmentation. Contrary to indirect volume rendering, however, the graphics hardware does not offer predefined functions to deal with overlapping.

Therefore, displaying overlapping structures using DVR requires some mechanism to dynamically hide context regions in favor of more important focus objects. There are a variety of methods to accomplish this, ranging from simply cutting out the entire context to geometric calculations identifying obstructing parts of overlapping features and removing them [4]. However, these methods seem unsuitable for depicting blood flow obtained from 4D PC-MRI using path lines. Since these path lines occupy nearly all of the space inside the vessel's context anatomy, rendering geometric calculations to find overlapping parts would be superfluous. On the other hand, simply removing the entire context may reduce the viewer's spatial awareness, making it harder to estimate position and size of the focus objects.

Csébfalvi et al. reduce the opacity of each voxel whose gradient is not orthogonal to the viewing direction [5]. This effectively hides flat surfaces and only leaves object boundaries visible, therefore opening the view for any insight feature. However, the restriction to boundaries reduces potentially useful context information by culling parts of the anatomy that are not actually covering the focus elements.

Instead of gradients, Zhou et al. use the distance of each voxel to a user-defined focal point as a factor for the opacity of this voxel [6]. This only works reliably for spherical objects, and also removes context that does not obstruct the view, e.g. regions that lie behind or next to the focus object. When visualizing blood flow data, the context anatomy wraps closely around the focus path lines. Therefore this approach would simply remove all of the anatomy. An improved version of this method by Tappenbeck et al. introduces two-dimensional transfer functions that incorporate both intensity value and distance to a previously segmented structure [7]. This approach removes the restriction to spherical objects, but requires a segmentation as well as configuration of the transfer function.

Viola et al. employ a conical culling volume to remove context in front of the focus object and leave other parts intact [4]. To determine which structures to focus on, each feature is assigned a unique "*importance*" value, whereas the culling volume only removes voxels belonging to structures with less importance. This is an effective method to ensure the visibility of important objects while retaining as much of the context as possible. It is especially useful when visualizing data sets with multiple anatomical layers, but not appropriate for complex, intertwined structures. Since blood vessels tend to intertwine with each other, removing context information in the vicinity of the streamlines may impede the perception of depth and therefore make it hard to interpret the visualization. Additionally, in our use case, the calculation of a culling volume is mostly superfluous, as cardiac 4D PC-MRI data do not contain more than two layers.

As a solution for these problems concerning cardiac blood flow data, we propose to extract and exclusively render the back sides of all anatomical structures. Therefore, we propose an automated ray-sampling algorithm to open up any context structure and show inset focus objects and achieving a similar effect to frontface culling on indirect volume visualizations. To validate our algorithm, we have created a prototypical application that is capable of visualizing blood flow data both utilizing our approach and a common one, and asked an expert radiologist to compare both visualizations.

2 Materials and Methods

2.1 Datasets and Visualization

We have tested our algorithm with a set of 20 three-dimensional, time-resolved flow images of the aorta acquired from 4D PC-MRI scans using a 3 Tesla Magnetom Verio MR with a V_{ENC} of 1.5 m/s. They were acquired from both healthy subjects as well as patients with different pathologies. The datasets consist of six images for each time step, containing flow direction and magnitude in x, y and z direction. They have a resolution of 132 × 192, with 15 to 25 slices and 11 to 19 time steps. For each dataset, a temporal maximum intensity projection (*TMIP*) was generated from the magnitude images. This TMIP image shows high intensities where blood was flowing fast at least once during the cardiac cycle, creating a high-contrast approximation of the vessel suitable for depicting the anatomical context using DVR.

The blood flow was visualized using particles with a uniform seed distribution inside the anatomy, drawing a trail behind them to form path lines as they follow the blood flow. To avoid placing particles outside of the anatomy, a mask defining valid particle positions is created by applying a threshold to the original TMIP. This is similar to the approach by Stalder et al. [8], but instead of drawing the path lines in front of a MIP of the dataset, we place them inside of the vessel anatomy. Since the threshold is linked to the transfer function used for anatomy rendering, the mask only contains voxels that are not rendered completely transparent. The composition of DVR anatomy and geometry-based path lines is accomplished by comparing the current depth during ray-sampling with the OpenGL depth buffer generated when rendering the lines.

2.2 Loopback-Based Frontface Culling

Since all particles are placed inside the anatomy, their path lines are generally occluded, as shown in Fig. 1b. The intravascular flow can only become visible if the vessel front is culled. For a geometric surface, we could simply calculate the normals for each face and remove those pointing towards the viewer (*traditional frontface culling*). To transfer this approach to DVR, we employ gradients, as high gradient magnitudes appear at structural boundaries and point towards the inside of structures with high intensity (vessels). However, this approach suffers from the fact that it simply cuts the anatomy in half instead of precisely culling anything but the back sides. Our algorithm requires an approximated binary vessel segmentation. For the screenshots presented in this paper, we used an intensity threshold on the TMIP, which is set to the median.

During ray casting, our algorithm stores two boolean values for each ray. The first (b_1) represents whether the ray has already hit an anatomical structure, the second (b_2) acts as a switch to control whether encountered voxels are to be rendered or skipped (see Fig. 1a). For each sampling point on the ray, we check if the current position is part of the anatomy by looking it up in the segmentation. This happens first at point P1, where the structure is hit for the first time. Now, b_1 switches to *true*. After that, once we encounter a voxel that is no longer part of the anatomy (point P2), b_2 is flipped to true and the sampling position is set back to a previous point on the ray (point P3). Only now the intensity values read from the sampling position will contribute to the rendered image. Due to jumping back after leaving the anatomic structure, its back side will be rendered while any other part of it gets culled. The effect of this rendering method in combination with path lines can be seen in Fig. 1c. By manipulating the distance between P2 and P3, thickness of the rendered back sides can be adjusted. To prevent holes in the visualization, this parameter should be set to a distance that equals at least one voxel in every possible viewing direction.

Our loopback-based frontface culling closely resembles a closest vessel projection [9], but instead of discarding all sampling points *after* hitting the first structure, we discard those that come *before*. To increase the robustness or our algorithm against noise and other image artifacts, it is possible to specify a minimum structural thickness. If the ray exits a structure without having traversed this given distance inside of it, the loopback will not occur and b_1 will be reset to *false*, effectively culling the entire structure. It should be noted that using a minimal structural thickness will not only remove artifacts, but also erode all anatomical structures. Therefore, in order to keep the visualization as faithful as possible, the minimal structural thickness should be at most two voxels.

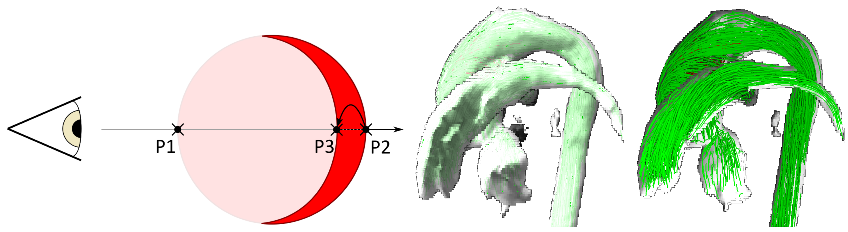

(a) (b) (c)

Fig. 1. Principle of our algorithm. P1 and P2 are entry and exit points, respectively, whereas P3 is the loopback point (a), flow visualization with loopback-based frontface culling disabled (b) and enabled (c).

3 Results

In an informal evaluation, we asked an experienced collaborating radiologist to compare our visualization against one where the path lines were simply drawn in front of a MIP, as shown in Fig. 2. The dataset in this example was acquired from a patient with a bypass and contains abnormal flow patterns in the form of vortices in the ascending aorta and aortic arch. In both visualizations, the red highlights are best noticeable when hiding laminar flow (*green*). With only the vortex visible, the lack of spatial information in the MIP visualization makes it difficult to pinpoint its exact location. Using our approach, on the other hand, the vessel structure is still clearly recognizable, allowing for easier localization of the vortex.

The expert was asked to compare the visualizations in respect to the visibility of path lines and the ability to pinpoint their location inside of the vessel. They confirmed that, in our visualization, the vessel's inner surface were clearer to see. It was also easier to judge the distance between path lines and vessel boundaries. To understand the overall vessel shape, however, the MIP visualization was more suitable, because even parts with lower contrast in the dataset were visible here. This was partially remedied by the ability to disable frontface culling at will, resulting in a more intuitive and complete model of the vessel. Still, it was noted that the ability to display the dataset using MIP should be kept as an option.

4 Discussion

We have presented a technique to visualize complex, overlapping structures in cardiac 4D PC-MRI data with minimal user interaction. It can be applied to any volume dataset as long as there is an approximate segmentation of the vascular anatomy and allows emphasizing inset features while retaining spatial information. Our method allows for a fast exploration and qualitative analysis of the dataset and containing flow patterns without the need of time-consuming

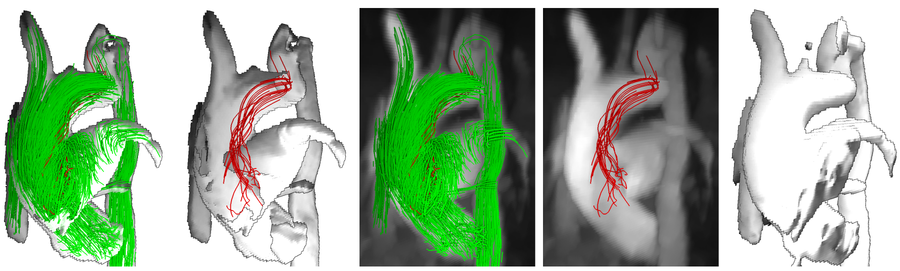

(a) (b) (c) (d) (e)

Fig. 2. Blood flow inside the aorta and pulmonary artery rendered with our method (*a*, *b*) and a MIP (*c*, *d*), in each case with (*right*) and without (*left*) laminar flow (green). Vortices are highlighted in red. A visualization of the vessel anatomy is provided for reference (*e*).

manual segmentation. Since only the front sides of vessels are removed, spatial information remains available. While we visualized flow inside of blood vessels, our method can also be used to display complex scenes including multiple layers of structures by rendering each object individually and then composing them into one final image.

Problems with our method arise from data quality as well as from using a TMIP to visualize the vessel anatomy. When two neighboring voxels have similar blood flow speeds, the TMIP does not generate a gradient between them even if they belong to different anatomical structures. Therefore, our algorithm will consider these intertwined structures as a single large structure and not properly separate them with a wall. These artifacts could be reduced by incorporating the principal flow direction into the anatomy visualization. Additionally, parts of the vessel with lower blood flow speed show up with reduced contrast in the TMIP, possibly causing them to be removed in the visualization. This could be solved by applying preprocessing algorithms to the TMIP to locally increase the contrast.

References

1. Barker AJ, Markl M, Bürk J, et al. Bicuspid aortic valve is associated with altered wall shear stress in the ascending aorta. Circ Cardiovasc Imaging. 2012;5(4):457–66.
2. Dyverfeldt P, Kvitting JE, Sigfridsson A, et al. Assessment of fluctuating velocities in disturbed cardiovascular blood flow: in vivo feasibility of generalized phase-contrast MRI. J Magn Reson Imaging. 2008;28(3):655–63.
3. Köhler B, Gasteiger R, Preim U, et al. Semi-automatic vortex extraction in 4D PC-MRI cardiac blood flow data using line predicates. IEEE Trans Vis Comput Graph. 2013;19(12):2773–82.
4. Viola I, Kanitsar A, Gröller ME. Importance-driven volume rendering. IEEE Visualization 2004. 2004; p. 139–45.
5. Csébfalvi B, Mroz L, Hauser H, et al. Fast Visualization of Object Contours by Non-Photorealistic Volume Rendering. Comput Graph Forum. 2001;20(3):452–60.
6. Zhou J, Döring A, Tönnies KD. Distance Based Enhancement for Focal Region Based Volume Rendering. In: Tolxdorff T, et al., editors. BVM 2004. Informatik aktuell. Berlin, Heidelberg: Springer Berlin Heidelberg; 2004. p. 199–03.
7. Tappenbeck A, Preim B, Dicken V. Distance-based transfer function design: Specification Methods and Applications. Simulation und Visualisierung. 2006; p. 259–74.
8. Stalder AF, Gulsun MA, Greiser A, et al. Fully automatic visualization of 4D Flow data. Proc Intl Soc Mag Reson Med. 2013;21:1434.
9. Zuiderveld KJ, Koning AHJ, Viergever MA. Techniques for speeding up high-quality perspective maximum intensity projection. Pattern Recognit Lett. 1994;15(5):507–17.